

UNIVERSITY OF LIVERPOOL

**Discovering the Process for New
Product Development (NPD) in
Machine Learning Software
DBA Thesis**

Thesis submitted in accordance with the requirements of the University of
Liverpool for the degree of Doctor of Business Administration

By Ali Naqvi

I **declare** that this **thesis** has been composed solely by myself and that it has not been submitted, in whole or in part, in any previous application for a degree. Except where states otherwise by reference or acknowledgment, the work presented is entirely my own.

A handwritten signature in black ink, appearing to read 'Ali Naqvi', with a stylized flourish at the end.

Ali Naqvi

Acknowledgements

I would like to thank my supervisor Dr. Paul Ellwood – whose guidance, direction, and support were instrumental for this thesis.

I would also like to acknowledge the second supervisor Dr. Sally Sambrook. Her feedback was extremely helpful.

I want to thank the two members of the AIAI team who participated in the research.

I would like to thank my family – my wife Shakila, my children, and my mother. Shakila's immense patience and support were instrumental throughout my DBA program.

I want to thank all my professors and the faculty of University of Liverpool who opened new ways of thinking for me. I also acknowledge the contribution of all those scholars and researchers whose work I referenced.

And finally, my thanks to the people who defined the times during which this thesis was done: the frontline workers who kept us safe during Covid and many who gave their lives to protect us. Many thanks to the doctors and nurses and all those whose sacrifices helped us survive this global crisis.

Ali Naqvi

Contents

ABSTRACT.....	1
Chapter 1 Introduction and Problem Outline.....	3
Overview.....	3
1.1 Introduction to AIAI	3
1.2 The Business Problem of AIAI.....	5
1.3 Practitioner’s Need for Research	8
1.4 Problem Statement.....	9
1.5 Statement of Purpose and Research Questions.....	13
1.6 Additional Background and Context	14
1.7 Research Approach	16
1.8 Rationale and Significance:	18
1.9 Findings	19
1.10 Next Steps	20
Chapter 2 Literature Review.....	21
2.1 Overview	21
2.2 Assumptions.....	22
2.3 Existing ML Methodologies	27
2.4 Systems Development as a Social Paradigm.....	39
2.5 Section 4: Synthesis.....	50
2.6 Section 5: Conceptual Framework	59
2.7 Building the Basic Conceptual Framework.....	67
2.8 The Difference with Traditional NPD	69
2.9 Next Steps.....	75
Chapter 3 Methodology	76
3.1 Introduction and Review	76
3.2 Research Approach	76
3.3 Research Sample	81
3.4 Overview of Information Needs	84
3.5 Research Design	86
3.6 Action cycles.....	86
3.7 The assumptions for choosing the method	88
3.8 The investigation process.....	89
3.9 Data Collection Methods.....	89

3.10 Data Analysis and Synthesis.....	94
3.11 Ethical Considerations	99
3.12 Issues of Trustworthiness	99
3.13 Limitations and Delimitations.....	100
3.14 Summary.....	101
Chapter 4 Findings and Cycle of Inquiry.....	102
4.1 Introduction.....	102
4.2 Cycle 1: <i>How should we generate ideas about a new ML product?</i>	103
4.3 Cycle 2: How to define and clarify a new ML product once its idea has been generated?	119
4.4 Cycle 3: Learning how to design the ML system?	131
4.5 Cycle 4: Learning how to understand the market potential or value creation of the system?	139
4.7 Cycle 5: Learning to Develop the ML Product	146
4.8 Cycle 6: Learning how to deploy the ML product.....	154
4.9 Customer Input and Verification	160
4.10 Competitive Advantage for AIAI	161
Chapter 5 Discussion	163
5.1 Introduction.....	163
5.2 The AIAI Model – A Review.....	164
5.3 The AIAI Experience	167
5.4 Reflections on the Social Paradigms of Machine Learning.....	173
5.5 Contribution to Knowledge	176
5.6 The Applied Examples of the Framework.....	179
5.7 The Ontological Insights	185
5.7 The Epistemological Insights	189
5.8 Application in Practice	191
5.9 Sensemaking	194
5.10 Summary	195
Chapter 6 Conclusion.....	197
6.1 Introduction.....	197
6.2 Limitations	199
6.3 Ideas for Future Research	202
6.4 Future of NPD in ML.....	204
Chapter 7 Reflections of a Scholar-Practitioner	206
References.....	211

List of Figures

Figure 2.1 Adopted from (Mariscal et al., 2010).....	28
Figure 2.2 Adopted from Fayyad et al., (1996).....	30
Figure 2.3 Adopted from Wirth (2000).....	32
Figure 2.4 Adopted from Rollins (2015), Methodology for Data Science	34
Figure 2.5 Four Paradigms Adopted from Hirschheim and Klein, 1989	43
Figure 2.6 Conceptual Framework Paradigms.....	65
Figure 2.7 Conceptual Frameowrk Application: Recommendation System	66
Figure 2.8 Conceptual Framework Application: Political Campaign System	67
Figure 2.9 Basic Conceptual Framework.....	69
Figure 4.1 Department of Homeland Security RFP.....	105
Figure 4.2 The Health and Human Services RFP (US).....	106
Figure 4.3 The Housing and Urban Development RFP (USA).....	107
Figure 4.4 From HUD RFP Response.....	108
Figure 4.5 HUD RFP Front Section	114
Figure 4.6 USAID RFP	115
Figure 4.7 The Discover Cycle.....	119
Figure 4.8 The Define Cycle	130
Figure 4.9 The Design Cycle.....	139
Figure 4.10 Adopted from Chapman et al. 2000.....	145
Figure 4.11 The Value Description Cycle	146
Figure 4.12 The Development Cycle	154
Figure 4.13 The Deployment Cycle	160
Figure 5.1 The AIAI NPD Framework.....	166
Figure 5.2 STEP 1 Analysis for Tesla.....	180

List of Tables

Table 2-1 Differences between Conventional and ML Systems	26
Table 3-1 Sample (People).....	83
Table 3-2 Information Sought.....	85
Table 3-3 Action Cycles	87
Table 4-1 Cycle 1 Data 1	111
Table 4-2 Meeting Notes Analysis Cycle 2.....	122
Table 4-3 Discovering Emerging Concepts	123
Table 4-4 Model to generate ideas about Products/Services.....	124
Table 4-5 AIAI Identifying its Product	125
Table 4-6 Identifying the Governance Needs	126
Table 4-7 Input, Output, Processing Data	132
Table 4-8 The SALA Model	135
Table 4-9 Value Creation Model	141
Table 4-10 Value Creation from ML (VCML)	142
Table 4-11 The Development Process	148
Table 4-12 Comparison with Conventional	150
Table 4-13 Exploring Deployment Paths	157
Table 4-14 The Deployment Process for AIAI	158
Table 5-1 Generating Ideas for Automation	170
Table 5-2 Identifying Paradigms in Existing Systems	175

ABSTRACT

American Institute of Artificial Intelligence, my employer, is seeking to launch a machine learning (ML) based product. The new product launch constitutes as a strategic initiative which the board has determined as necessary and critical for the survival of the business. ML is a rapidly expanding branch of artificial intelligence (AI).

A generic New Product Development (NPD) process is composed of two subprocesses: the concept design subprocess and the product development subprocess. At present, data mining methodologies (for example, CRISP-DM, KDD, SEMMA) from the late 1990's are being widely used to develop ML products – including for NPD purposes. Recent research indicates that those methodologies are being viewed as too narrow, unsuitable, and incompatible with the growing needs of the machine learning's rapid adoption and that practitioners are searching for alternatives. Recent findings also indicate that the failure rate in ML products and services is high. To make the product launch successful, AIAI must deploy a reliable and functional framework. Launching a new product without a methodology or framework will be irresponsible and using the existing methodology will be too risky for the organization. Given the strategic nature of the product launch, AIAI has determined that there is a need to explore what would constitute as an appropriate NPD framework for ML products and services. This research was launched to achieve that goal.

The fundamental question addressed by this study is: **what is a new product development framework for designing and developing machine learning products?**

Many well-developed and time-tested frameworks exist in the conventional (non-ML) information systems development. Information systems development

(ISD) paradigms in the IT field have a renowned status as many researchers have pointed out that all approaches and information systems methods (for example, Waterfall, Agile, Rapid Application Development, and others) are directly or indirectly linked to the paradigms and approaches. This study also explores how conventional ISD paradigms can help define or discover an ML NPD methodology.

This inquiry is undertaken as *action research* in the active setting of launching the AIAI's ML product. Whether conventional paradigms and approaches apply, or new ones are discovered, it is hoped that the research will contribute to a deeper understanding of ML systems design and development and advance the knowledge of designing new ML technologies.

To have a successful product launch, the Institute must proceed with this research. As such, the research and the project timings will be synchronized. It is expected that the actions related to the product launch will help drive the research and the research will help drive the actions to achieve the goals of the project. Product launch involves other colleagues and partners and collaboration between various groups is essential. In addition to addressing the critical business problem and making contribution to knowledge, it is hoped that through this experience I will develop myself as a practitioner-researcher, as a leader, and most importantly, as a human being.

Chapter 1 Introduction and Problem Outline

Overview

This chapter establishes the context and background of the research. The twofold approach to frame the problem captures both the specific AIAI needs as well as the research and knowledge gap that exists in the industry. That is followed by the problem statement, purpose, and the research questions. The chapter also covers the research approach, introduction to the researcher (me) and the participants, my perspective and assumptions, and the proposed rationale and significance of study. The chapter concludes with introducing the next steps.

1.1 Introduction to AIAI

The American Institute of Artificial Intelligence (AIAI) is a Washington DC based private institute that specializes in education and research in artificial intelligence (AI). Founded in 2016, American Institute of Artificial Intelligence is the world's first AI business school. AIAI has achieved international recognition and revenue growth from its course offerings. Due to a recent business change, AIAI's board has mandated AIAI to transition from an education provider to a product company. Specifically, AIAI management has been tasked to use its AI knowledge to successfully design, develop, and launch a machine learning (ML) product.

The Organization

The Institute is composed of various professionals with different backgrounds. The educators serve as advisors, trainers, or researchers. The technology team supports the internal technology infrastructure. For instance, ML research products, website, and marketing and customer relationship management.

In addition to the technical team, the Institute is being helped by various business development professionals.

The primary Action Research team was composed of three people – including me. I had the dual role of practitioner-researcher as I had to ensure that AIAI's new product development continued as the research progressed. I was the lead researcher while two experienced members of AI team were the research participants. Feedback from the AR participants served as the data for determining the practical aspects of the research and for driving action. It also kept me grounded and pragmatic, and it enforced the discipline and timeline necessary for the success of the AIAI project.

The Culture

AIAI is an academic institute as well as an entrepreneurially run firm. AIAI maintains a culture of intellectual openness and research orientation. AIAI is composed of many thought leaders in multiple fields and is supported by professional technology and sales staff. The unique blend of research and teaching is being tested to transform the Institute into a Silicon Valley style software company. The key components of AIAI's culture are integrity, openness, intellectual honesty, and science.

I have more than 25 years of experience in business. I serve as the CEO. In addition to conducting this research to create value for AIAI, I sought to develop my leadership skills. For this research, I received guidance from my research supervisor via frequent progress check and feedback sessions.

AIAI has established a charter for values and strictly follows the following values:

Scientific Knowledge

We believe in the supremacy of responsible science and technology to solve human problems. Our responsibility is to enable all humankind to benefit from the advances in science and technology. We produce and teach science, technology, engineering, and math to improve lives and reduce human suffering. Our research will be based upon responsible science and technology.

Human Benefit

We believe in maintaining the supremacy of humans over machines. We will not use technology and science in areas that will increase human suffering, diminish human dignity and privacy, and allow our research and technology to hurt human interests.

Better World

We believe in maintaining and preserving the supremacy of all biological lifeforms over machines. We will create a better world with AI. A better world ensures that biological lifeforms thrive and can exist in a safe and healthy environment for all.

Ethical Service Leadership

We believe in the supremacy of justice, empathy, and service to create a collective human consciousness for the good of all humankind. We will work tirelessly to improve human lives, approach all decisions with the mindset of uncompromising ethics, and serve diligently. We exist to serve. We lead by example.

Multi Stakeholder Value Creation

We believe in the supremacy of collective value creation that serves all our stakeholders. We will not serve the interests of a single constituency (for example shareholders) at the cost of other stakeholders (e.g. customers, humankind, climate).

1.2 The Business Problem of AIAI

Despite the novelty and success, I suspected that AIAI's growth will be short-lived. The business of education, these days, has low barriers to entry. Online copycat courses can rapidly cannibalize the market position and revenue stream of educational institutions. Copies of online courses can be shared illegally. With

the massive material now available on YouTube, consumers expect education to be free. As the revenues declined nearly 30% after a steep increase in the early years, AIAI management is under pressure to increase sales. While AIAI's revenues declined on the education side, they temporarily increased on the research side. Companies paid to fund research projects via grants. Unable to compete with the research by low-priced graduate and undergraduate students from universities, AIAI had to drop its research prices. Companies expected AIAI to conduct sophisticated research but paid the rates of university students. Hence, while that part of the business could have flourished, the profit margins from that business also diminished by nearly 60%

With both Research and Education business segments hitting an impasse, AIAI needed a new source of revenues. Urgent intervention was needed. AIAI could not have survived without a new strategy and business model. A decision was made to leverage AIAI's knowledge of the AI field and to transition the institute to a software product company. This meant to launch a product that is developed with ML. ML is a subbranch of AI. The concerns about product diffusion and absorption were alleviated by the market knowledge that the Institute possessed. The Institute extensively covered various areas of business (for example, AI in Finance, AI in Supply Chain, AI in Marketing, etc.) and hence possessed market information to assess the potential product gaps and opportunities.

Companies have been launching new products for hundreds of years and companies have been launching IT products for decades, so AIAI's initial expectations did not assume any major complications. AIAI had decided to launch a new machine learning based product that can automate a business function. We did not know what specific product we will build as those questions will be answered during the NPD process. However, as soon as the initial planning work began, a rather disturbing statistic became the source of major alarm and concern for the firm. Nimdzi Insights Pactera (2019) reported that

75% to 85% machine learning projects were failing to achieve their milestones and are failing to impress chief information officers (Nimdzi Insights Pactera EDGE, 2019). This created a major dilemma. AIAI was about to undertake a project where it was known from the very latest statistics that the chances of success are merely 25%. Suddenly, the idea that 'launching an ML project will be a trivial undertaking' disappeared and was replaced by a much more cautious approach.

The more conservative approach entailed the company needing a reliable AI and ML centric NPD framework for launching the new product. As soon as it was recognized that a framework will be needed, a very high-level business search for the framework began. The search for the framework led to the consideration of the following insights:

1. A generic NPD framework can be applied for developing the AIAI's machine learning product.
2. AIAI can also search for and identify an existing framework that was specifically designed for developing ML products and artifacts.
3. While recognizing that ML is different than conventional information systems, AIAI can also explore existing methodologies and frameworks that are applied in conventional information technology as useful or usable to build machine learning systems.
4. Finally, it was understood that merely having a methodology or framework was not enough, as the methodology or framework must have the potential to achieve project success. It should be able to help architect a competitive advantage for firms (Brockhoff, 2003; Kahn et al., 2006; Cooper and Sommer, 2018). Since it was expected that the industry would have used some frameworks to design and develop ML products, the failure rate of 75-85% was viewed as an indicator that the existing methodologies or frameworks were not achieving the desired results. Therefore, the

possibility that AIAI may have to discover a methodology for developing and designing ML products was kept open.

The above led to the recognition that a broader study will be needed to identify a framework that can help address the AIAI problem of undertaking a project without a reference framework.

While the first three areas of inquiry indicated a search for an existing framework and could have been business consultant type investigations, it was recognized that the fourth point – discovering the methodology that creates success for AIAI – would require more than just a business consultant. AIAI realized that the discovery of a NPD framework could include elements from existing AI development methodologies, from conventional systems methodologies, and elements may not exist in the first three areas pointed out above. Such elements would need to be developed, while others removed or replaced. Developing those elements implied that research needed to move beyond simple business practitioner research. With research acquiring a broader scope, it was recognized early in the process that the project itself will contribute to the framework as project's dynamics can provide rich source of data that can be developed, tested, and then findings applied. Building upon AIAI's existing experience in various areas of AI/ML, it was expected that AIAI has a unique understanding of product gaps in various areas of business and hence some expected product diffusion knowledge existed within the firm – but since I was trying to develop an NPD framework, I wanted to explore such factors as part of the process and not make any upfront assumptions. The process wheels were already set into motion. The research was expected to influence the activities taking place in the project.

1.3 Practitioner's Need for Research

From a practitioner's perspective, the research encompasses ensuring that a methodology is applied to increase the likelihood of new product launch success. I, when approaching the problem from a practitioner's perspective, was

concerned about the high failure rate of machine learning projects. I reasoned that with three out of four projects failing implied either a methodology (or even a framework) was not used, or if used, it did not achieve results. Both options were problematic from the project perspective. The high failure rate indicated that this research would require realignment and reassembling of various methodologies, and probably the discovery of a new framework.

As the CEO, I was aware of the importance of finding a framework that works. Not having a framework or having an underdeveloped or wrong framework could cost the firm its product and its future. The product launched needed to be more than an offering, it was expected to establish a competitive advantage for the Institute. The NPD process was also necessary to decipher the elements of competitive advantage. Furthermore, the NPD process would lead to a product that will deliver sustained revenues for AIAI and help position AIAI as more than just an academic institution.

1.4 Problem Statement

The AI revolution is a recent phenomenon and even though many firms are being launched in this domain, little is known about new product design and development frameworks in AI. NPD carries tremendous risk and AIAI's plans to launch a new product are imperiled by the absence of a comprehensive, relevant, and compatible framework to design and develop ML products. To close the gap created by not having a framework, I formulated the following specific problem that I explored during the study:

The fundamental question addressed by the study was: **what is a new product development (NPD) framework for designing and developing machine learning products?**

One derivative of the above research question is: how is ML centric (AI) business model and product launch different than the traditional models and launch? In

this context, “traditional models” implies conventional methods applied to launch nonintelligent digital technologies (regular software).

This research question acquires significance because even at the outset of the research, certain facts are known about the differences between ML and conventional information systems:

- 1) Non-intelligent ICT was based upon the ability to capture, organize, process, analyze, and retrieve data. Intelligent automation or ML automation is based upon using data to help machines learn (Domingos, 2015). This creates a fundamental difference between the functional performance from the two types of systems: one performs as a deterministic machine that processes data; while the latter accumulates experience, learns, and improves its performance with experience.
- 2) The associated economics of AI centric transformation are greatly different than non-AI centric technological transformation (Agrawal et al., 2019). Consider the question like: should the parking lot capacity of new building construction be remodeled based upon autonomous vehicles? Autonomous vehicles will most likely be always on the road, as cabs/taxis, and will only stop to get recharged. This means, new buildings will need less parking space.
- 3) The business processes related to ML automation will require fundamental rethinking (Dirican, 2015). For instance, the process of online sale (which can be viewed as a digital business process) was greatly altered when intelligent recommendation engines began inducing people to buy more based upon their previous history and other data. Behind the scenes, intelligent pricing engines churn to maximize the price based upon various supply and demand factors, and the intelligently automated robot-controlled warehouses scramble to

perform fulfillment and rapid delivery functions. Automation of such complicated and interdependent processes necessarily requires the automation of tasks that require complex decision-making and intelligence.

- 4) The organization of the intelligently automated world is a modern, never-seen-before phenomenon (Schwab, 2015). Humans working next to digital workers, sharing thought process and data, and competing with robots for jobs are only some of the issues. Other profound issues are of humans getting analyzed, their privacy invaded, their every action recorded, and even their thoughts and actions predicted. The organizational factor is of great importance to this thesis.
- 5) Digital automation has been approached from the concept of reductionism. Business is viewed as a simple or a complicated system that can be reduced to simpler, discrete, linear system whose parts can be automated by machines. In reality, business is a complex system and therefore a Newtonian mechanical view of business may not be helpful (Stacey, 2011) in intelligent automation.

Accordingly, we are operating in an era that is remarkably different than the times of simple digital technologies or nonintelligent machines. While the AI field has been around since the 1940's, recent innovations in technology have enabled the extensive widespread production and adoption of AI and ML products (such as autonomous cars, personal assistant such as Siri, recommendation systems, and others), leading to what Klaus Schwab calls the Fourth Revolution (Schwab, 2015). For the purposes of this thesis, it is necessary to draw a clear distinction between the information systems of the non-intelligent digital era and the information systems of the intelligent and ML automation era. Non-intelligent information systems can perform strictly in accordance with what their programmer has instructed them to do via line-by-line code of program instruction. Intelligent systems are those that can learn, generalize learning,

adapt, accumulate experience, and improve their performance (Vieira et al., 2020). These significant distinctions between intelligent machines (products) and all other products and services demands exploration of NPD frameworks and methodologies relevant and applicable for these new generation of products.

Despite the profound promise and novelty of the *intelligence era*, the business and operating dynamics of the intelligence era have not been fully deciphered (Hajkowicz et al., 2019; UK, 2017; SCAI, 2019). The human civilization had barely begun to understand the business (and other social, political etc.) dynamics and implications of the digital era (computers and software), and now it has already entered the *intelligent era* (Schwab, 2016). The novelty itself may be a prelude to statistics such as 85% of AI projects are failing (as shown above). It is possible that without an AI/ML NPD framework, the development of business models and AI products could be haphazard.

However, NPD frameworks are not just technical in nature, they also help achieve better co-operation, co-ordination, and communication amongst the product development teams (Shepherd and Ahmed, 2000). Specifically, software development is a social process and the role of social paradigms has been identified as an effective way to architect the governing and underlying philosophies (Isaias and Issa, 2015) and information systems development methodologies are derived from or are part of social paradigms and approaches (Hirschheim and Klein, 1989). Research shows that methodologies and frameworks are essential for complex engineering projects however both conventional and data mining methodologies fall short of supporting ML development and engineering (Martinez-Plumed et al., 2019; Schelter et al., 2018; Marbán et al., 2009). The dearth of a ML design and development paradigm and approaches is not only an AIAI specific issue, but as previously cited research also indicates, it is an industrywide problem. The knowledge gap is

profound, and this gap has a material financial impact on the industry, and it also exposes AIAI to substantial business risk.

1.5 Statement of Purpose and Research Questions

The purpose of this research was to investigate what is an NPD framework for machine learning systems. It was anticipated that the research will focus on three main areas of:

- Identification of a baseline skeleton for a generic NPD framework
- Specific design and development frameworks in machine learning
- Social paradigms from which the design and development frameworks in conventional information systems are extracted.

To support the primary research question, the main questions explored by the research were:

- What is a generic NPD framework that is suitable for building an NPD for ML products and services?
- What are the social paradigms from which design and development frameworks in conventional systems were extracted? In what ways they can be applied in ML?

Having a baseline skeleton generic NPD model (Wang, 2016; Shepherd and Ahmed, 2000) would shed light on how to connect the business requirements side (Concept Design Subprocess: *what should the firm do*) to the engineering side (Product Development Subprocess: *how should the firm do it*).

The strategy employed in this research first explored the engineering approaches (Gruner, 2010, 2011) of ML to investigate the Product Development Subprocess. It addressed: What are the prevalent or current design and development frameworks for ML? Are those frameworks effective, and do they meet the requirements of modern-day ML? Since the research is about ML process and because software systems are known to be based upon social and scientific

realities (Hirschheim and Klein, 1989), the research then focused on exploring the concepts of scientific and social realities in ML development. Combining the two – the research establishes a basic approach for ML product development and attempts to identify one possible NPD framework for ML.

It was anticipated that the research findings from the paradigms and engineering process can be knit or stitched together to propose a conceptual framework. This framework can then be enriched from practice based fieldwork to help develop an NPD framework. Through this exploration, it was expected that the research will lead to the discovery of the supporting framework for AIAI's ML product launch. Furthermore, the study expected that there could be a more generalized application of the discovered framework that could be suitable for industrywide application.

It was likely that the insights from the research would also contribute to the literature by exploring if social paradigms used in conventional information systems development (ISD) can be used to improve ML development and possibly discovering new *ISD paradigms* that are more suitable for ML. This research employed Action Research methodology. The primary investigation techniques were qualitative.

1.6 Additional Background and Context

ML has recently become pervasive in business (Makridakis, 2017; Wright and Schultz, 2018). Digital transformation, often associated with the advent of AI and ML, is being considered a paradigm shift and is even being described as a revolution of its own (Schwab, 2016, 2015). AI has become the top priority of countries and companies (Naqvi and Munoz, 2020; US Government, 2019). Despite being positioned as a game changing technology, recent practitioner research reveals that 85% of AI projects are not meeting the expectations (Nimdzi Insights Pactera EDGE, 2019). One of the reasons offered for such a high failure rate is that AI and ML are new technologies, the data science field has evolved,

and practitioners do not have a clear understanding of how to design and develop these technologies (Martinez-Plumed et al., 2019; Marbán et al., 2009).

The design and develop methodologies in information technology are not a new phenomenon. Originating from general engineering, many methodologies go back to 1950's (Isaias and Issa, 2015). It is also known that the core process of ML development is substantially different than the development process of conventional IT systems (Rahman et al., 2019). NPD for products with data driven features has been analyzed (Li et al., 2019) but that does not provide a framework.

Initial research suggests that software development in information technology has long relied upon design, development, and engineering methodologies (Isaias and Issa, 2015; Eason, 2016; Stoica et al., 2016; Nugroho et al., 2017).

However, does the same process apply for developing intelligent systems? Would a system capable of programming (and even engineering) itself will require a different approach for NPD? Would customers know what their needs are when customers may lack the knowledge about the possibilities and potential of what ML can achieve? Would customers (or even designers) know or address the governance and ethical issues of creating intelligence? For example, if ML is being used to make decisions that affect people's livelihoods or if the product is making decisions itself (e.g. hire or not hire, promote or fire), lives (justice department sentencing or healthcare pathways), brake or run over (autonomous car), or privacy – should governance and ethics be included in design consideration? All of those were important considerations but the methodologies and frameworks designed for non-intelligent products and services generally do not address those issues. The obvious difference is that intelligent machines are a new phenomenon for humankind and therefore may need a corresponding change in NPD.

Current literature suggests that ML products are being designed with conventional IT methodologies, or with data mining methodologies, or with a combination of the two (Amershi et al., 2019; Langford and Ortega, 2012; Dåderman and Rosander, 2018; DiMasi et al., 2009; Martinez-Plumed et al., 2019). In conventional IT, humans use information to make decisions and take actions – but the artifacts themselves do not possess intelligence to make decisions or take actions.

From a practitioner's (my) perspective, the problem of operating without a framework to launch a new product required immediate and material intervention. I was convinced that it would have been irresponsible and reckless to proceed without some type of a framework. It had become a source of great risk for AIAI.

1.7 Research Approach

Various approaches to conduct this research were considered. The Action Research (AR) approach was selected as the preferred methodology for this research. Action researchers “see the development of theory or understanding as a by-product of the improvement of real situations, rather than application as a by-product of advances in ‘pure’ theory.” (Carr and Kremmis, 1986, p.28) “This is a means to generate ideas (theory) that are relevant locally – to the people who are involved in the research, and to the environment in which it has taken place.” (Willis, 2010, p.167).

With the approval of the ethics board at University of Liverpool, two leaders of AIAI were contacted and briefed on the research and their participation was requested. Upon receiving the proper disclosures and communications related to the research, the participants voluntarily opted to participate in the research. No pressure of any sort was placed upon the participants to engage in the research. It is my understanding that participants viewed this research as an opportunity to

help the firm and hence were intrinsically motivated. Even though I am the CEO, the AIAI culture empowers our employees to act as peers and equity owners rather than just employees. The research followed the directives by University of Liverpool on data collection during the Covid19 crisis.

The AR approach calls for research in action on action for action – specifically the underlying principles offer participatory character, democratic impulse, and simultaneous contribution to social science (knowledge) and social change (practice) (Carr and Kemmis, 1986). Hence, the research is primarily driven by action. In that regard, cycles of actions were set up – and within each cycle were several *action* → *reflection* cycles. In the absence of an alternative path, the action cycles were assumed to coincide with the natural progression path of the product launch. It was recognized that the research's influence on the progression path of the product launch could be significant, and it is possible that the research findings can alter the cycles or lead to a different path. AR methodology takes into account the organizational context, political issues, social constructions, and other such practical concerns that are often ignored in some forms of research (Raelin and Coghlan, 2006). The researcher embeds himself (herself) into the research and observes data from various vantage points (Eden and Huxham, 1996). AR also provides an avenue of reflexivity and helps the researcher discover his/her own biases while improving the researcher in various ways. As will be explained later in Chapter 3 (Methodology), such reflexivity was formally added in the research design.

In AR, actions become the primary source of data. The researcher gets many opportunities to collect data from different vantage points. The primary sources of data included meeting notes, emails and texts, narratives captured from calls, literature review, structured questionnaires, field notes, and AIAI documents. Certain factual information (for example, capital investment going into the AI field) was picked up from credible or widely used practitioner websites and

reports (for example, Gartner). Data was generated from the actions undertaken during the action cycle and the post-activity meetings. The research was approached as qualitative research.

Appropriate consideration was given to establish trust, avoid bias, and enhance the credibility of the analysis. Wherever possible attempts were made to back the findings with checks with the participants and the relevant literature. This provided an added measure of refinement and crosscheck (triangulation). To analyze qualitative data, the concepts and exercises were developed based upon my understanding of the field and the literature review. Evaluation and reflection were backed by the relevant literature review or established process details. Caution was exercised to make sure standard conceptual frameworks were used. The questions posed for participants were mostly descriptive and explorative (for example, an observation leading to asking to describe the attributes of existing ML applications) and interpretive (for example, exploring the participant's interpretation of an observation).

1.8 Rationale and Significance:

The rationale for this study emanated from my employer, AIAI, launching a new product for which AIAI needed a genuine and necessary investigation. Since the business problem was not related to existing knowledge or that could have been solved by simple survey research, a formal investigation was deemed necessary. Additionally, the problem of not having such a methodology at an industry level implied that the study could have also helped formulate a framework for an industry level ML methodology. However, while I hoped that certain broad guidelines may emerge, this research, even if it did result in formulating an integrated framework, did not necessarily claim or seek generalizability.

When the statistic of 85% failure rate is combined with the observation of multiple researchers that a ML specific framework and methodology are vital, the need for an industry level framework becomes evident. Paradigms, approaches,

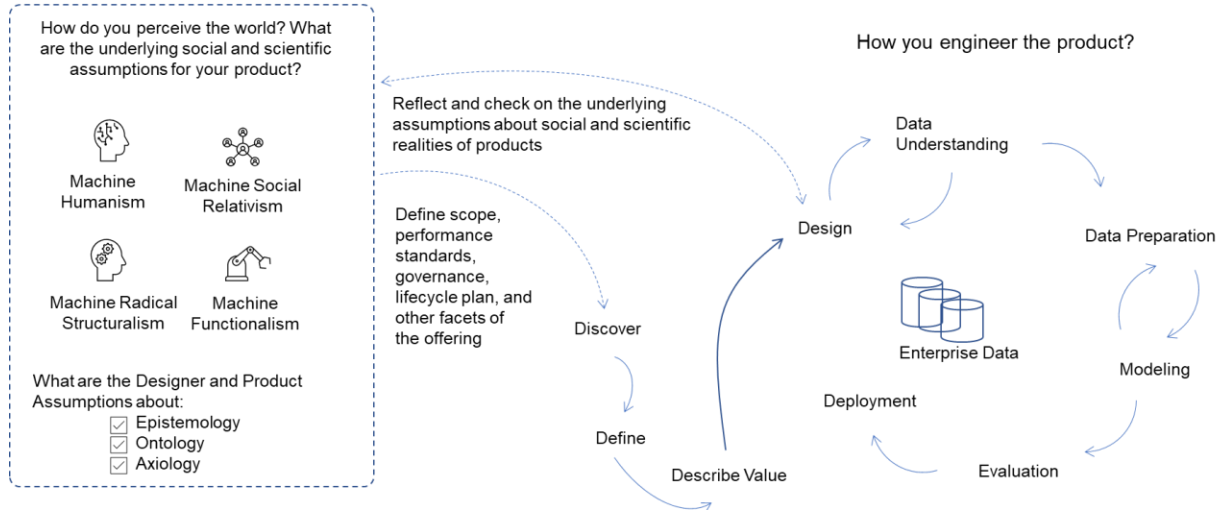
and methodologies are meant to improve outcomes (Isaias and Issa, 2015). The amount of new venture investment going into AI increased 350% between 2013 and 2017 and was estimated to be nearly \$7 Billion in 2017 (Shoham et al., 2018). To generate a rough estimate of value at risk, one can multiply the 75% to 85% failure rate (Nimdzi Insights Pactera EDGE, 2019) with the capital investment to determine the value at risk for the investment to be at least \$5 Billion. This implies that improving the ML development process can have a meaningful financial impact on the entire industry.

For AIAI, the research was expected to create value by significantly reducing the new product launch risk. A framework could also help AIAI launch future products. The framework, if generalizable, could help other firms in the industry and governments. A literature review was developed and is presented in the next chapter.

1.9 Findings

The research findings led to the development of an ML NPD framework composed of a process where product concept and design teams (including product marketing, lifecycle planning, etc.) operate with an awareness about the social and scientific realities of the product and follow a recommended action sequence for product idea-to-launch process which is composed of 12 action sequences and over 30 tasks. At all parts in the idea-to-launch journey new product development teams are encouraged to constantly check their social-scientific assumptions – including the underlying epistemological, ontological, and axiological vantage points for products and human users. The underlying assumptions about the nature of the product, as determined by the social and scientific beliefs and their sub-assumptions, are likely to produce very different products to solve the same problem and are likely to influence the subsequent process steps in the NPD process.

THE AIAI ML NPD FRAMEWORK



1.10 Next Steps

Once the research need for AIAI was identified and proper approvals were obtained, the next step was to With no frameworks available to support my practitioner needs, I replaced my practitioner hat with my researcher hat and conducted extensive literature review which is presented in the next chapter.

Chapter 2 Literature Review

2.1 Overview

The purpose of this research was to develop a new machine learning (ML) based intelligent product in an educational organization which had never launched a ML product before. The management challenge was to identify or discover a new product development (NPD) framework for machine learning products. The research goal of identifying or discovering a framework required literature review to gain insights that are not available in practitioner centric research and that were deemed necessary for the product launch success.

The broad strategy for research is based upon identifying a barebone, skeleton, generic NPD process and then enriching it via action research.

The literature review strategy was centered upon first investigating the dominant data mining (ML related) methodologies and then exploring the underlying social (ontological, epistemological, and ethical) foundations of conventional information systems development (ISD) paradigms. To explore the fundamental concepts of ISD – concepts that clarify how people and organizations perceive and approach information systems – I focused the second part of the search on information systems development paradigms. Finally, toward the end, I analyze some of the NPD literature to give meaning to the conceptual framework.

In accordance with the traditions of information systems literature, systems in this research were viewed as a combination of technology, data, design methodologies, information, social norms, values, cultures, and human imagination (Iivari, 2017b; A. S. Lee et al., 2015). This also implied the presence of oppression, politics, power, and domination in ISD (Iivari and Kuutti, 2018). A conceptual framework was developed from the research that was not only referred to during the field work but also became a critical part of the discovered final framework. To develop the

theoretical foundations of this research, support the research choices, carve out a framework, and gain insights about the state of the knowledge, the critical review presented in this chapter was conducted.

The following two major bodies of literature were critically analyzed:

- a. The data mining methodologies that are being widely used in ML.
- b. The conventional information systems development (ISD) frameworks and paradigms.

The critical review is divided into five sections of: **Section 2.2** sets the stage for the analysis by clarifying assumptions and orienting reader to some fundamental concepts. **Section 2.3** covers the critical review of existing ML engineering methodologies (data mining) that are being used in practice as well as analyzing their shortcomings; **Section 2.4** covers the social paradigms used in conventional ISD and critically analyzes their applications for ML; **Section 2.5** synthesizes and integrates the critical review from the previous three sections; **Section 2.6** proposes a conceptual framework which forms the scaffolding for the methodology of this research. **Section 2.7** describes the final conceptual framework. **Section 2.8** covers discussion about the conceptual framework in light of the new product development literature

2.2 Assumptions

The study was undertaken with the following three Assumptions and two fundamental concept explanations.

Assumption 1: The critical analysis is based upon seminal works, however, given the nature of novelty of the field some referenced research is too recent to have a seminal status (in terms of citations). The value of such research is determined via cross validation.

Assumption 2: This is management research and even though the research deals with important scientific and engineering issues, the review did not include complex mathematics or engineering related scientific jargon since such an inclusion would have been both counterproductive and outside the scope of the research.

Assumption 3: This research takes a position that humankind today is already using significant AI and ML technology (Schwab, 2015). Work today is performed by collaboration between humans and machines (Lauterbach and Bonim, 2016; Brynjolfsson and McAfee, 2011, 2015; McAfee and Brynjolfsson, 2016), and significant automation is already underway (Frey and Osborne, 2017).

In addition to the assumptions, two important technical concepts set the stage for the critical review discussion that follows. First fundamental explains the difference between conventional information systems (IS) and ML systems. Second, provides a very high-level, non-technical, introduction to machine learning techniques. Both are explained in nontechnical, business-friendly language.

Machine Learning and Conventional IT systems

In 1959 Arthur Samuels (considered a pioneer of ML) coined the term machine learning (IBM) (Samuel, 1959). By paraphrasing Arthur Samuel, researchers attributed to him the definition of ML as computer programs that learn “without being explicitly programmed” (Moser, 1990, p.10). As a field, ML aims to answer the question “How can we build computer systems that automatically improve with experience, and what are the fundamental laws that govern all learning processes?” (Mitchell, 2006, p.1). Tom Mitchell explains that we can think of ML as when a machine learns a particular task T , with a performance metric P , and type of experience E and that the machine reliably improves its performance P to perform task T . The various combinations of P , T , and E , Mitchell clarifies, produce different learning methods (Mitchell, 1997). Other have provided

simpler explanations. Vieira et al. (2020) define ML as “area of artificial intelligence that is concerned with identifying patterns from data and to use these patterns to make predictions about unseen data”(Vieira et al., 2020, p.4); and ML can be seen as science and art by which computers learn things by themselves (Domingos, 2015). Bishop views pattern recognition and ML as two facets of the same field (Bishop, 2006). In terms of extracting reliable patterns from data, ML is a blend of computer science and statistics, but also covers mathematical areas such as algorithms and human or animal learning such as from psychology and neuroscience (Mitchell, 1997). In ML, a human does not provide step by step instructions to the computer. Throughout this dissertation, machine learning systems will be addressed as ML systems. In summary, ML systems are goal-oriented learning systems that develop from data and that can navigate through uncertainty.

In contrast, a conventional information system is information technology that results from precise instructions given by a human (programmer) that can be executed by a computer to perform a specific task (Rochkind, 2004). In conventional systems, computer can only perform within the strict confinements of the instructions provided to it and does not possess any ability to learn, accumulate experience, or improve performance. A conventional information system does not morph into existence from data; it does not evolve. It is formed with programmed instructions (instructions can also be viewed as rules); it processes and stores data; and it constitutes as a system to provide information. Throughout this dissertation, traditional (non-AI) information systems will be addressed as conventional systems. Information systems are defined as “a computer-supported system which provides a set of people (users) with information on specified topics of interest in a certain organizational context” (Iivari and Hirschheim, 1996, p.552). Citing Heinrich (1999), Gruner and Kroeze (2014) establish the following baseline definitions (Gruner and Kroeze, 2014, p.1,2; Heinrich, 1999):

- *An information system is “a hybrid human/task/technics-system for the acquisition, production, storage and usage of information (including communication) for the purpose of satisfying some users' information demand”;*
- *Information is clarified as “action-guiding knowledge about past, present or future states of reality, or events in reality” – that transpires specially in “commercial enterprises or administrative organizations”;*
- *this concept of information in the IS field is thus “strongly related to the world of human life and must therefore not be conflated with the abstract technical notion of information (quantified in the unit of Bit) as in Shannon's mathematical information theory”;*
- *“Human is understood specifically as the bearer of some task” and task is viewed as “one of the organizational or business processes to be supported by such an information system”; and*
- *Technics (not to be confused with technology) is understood as the means, methods and devices by which such a system is materially implemented.*

This research accepts all of the above definitions as applicable in this research. Specially, the concepts of *task and work* are important since intelligence automation or ML attempts to automate work (Frey and Osborne, 2017). Alter (2008) also includes reference to work in IS definition and defines an information system as a work system that supports an organization's networks of information creation, gathering, processing, or storing (Alter, 2008).

Based upon the above definitional constructs, it is possible to extrapolate the differences between conventional systems and ML systems as presented in Table 2.1:

PROPERTY	CONVENTIONAL	ML
Solution space	Deterministic	Statistical/Stochastic
Programmer	Human programmer provides specific instructions	System learns by itself; Machine as the programmer
Development process	Code written in a programming language with specific instructions	Programming language used to train algorithms
Nature	Does not learn, adapt, or accumulate experience.	Learns, adapts, finds patterns, and accumulates experience.
Performance potential	Static performance	Performance improvement possible
Construction	For data	From data
Structure	Rules based	Algorithm based
Role of Human Creator	Human programmer trained in programming language	Human creator trained in ML methods (data science) and programming language
Related Fields in addition to Computer Science	Discrete Math, Data Structures, Programming language	Statistics, Calculus, Linear Algebra, Psychology, Cognitive Science, Learning Theory, Neuroscience
Role	Information creation, gathering, processing, or storing	Learning from data

Table 2-1 Differences between Conventional and ML Systems

Machine Learning Techniques

To help the business reader develop the context for ML, this subsection briefly describes the three main ML techniques. ML is composed of three broad techniques: supervised, unsupervised, and reinforcement learning.

Supervised learning happens when machines are provided examples of inputs and related outputs. From these examples, machines learn to classify objects or concepts and to predict continuous data.

Unsupervised learning is when machines learn from identifying patterns based upon how data clusters. In unsupervised learning, no examples are provided. The algorithms discover patterns in the various attributes of the provided data. It is possible that unsupervised learning can discover patterns that humans cannot explain.

Reinforcement learning happens when algorithm learns by reward and penalties.

2.3 Existing ML Methodologies

The goal of this section is to review the current methodologies that are being used in ML, identify their shortcomings and limitations as reported in the literature, and perform critical analysis. The section is divided into Review of the Current Methodologies used in ML, Failures and Limitations of Current Methodologies, and Critical Analysis.

Review of the Current Methodologies

The ML field uses methodologies that were originally developed for a subfield, and a significantly narrower area, of ML known as data mining (Martinez-Plumed et al., 2019). Since 1993, the field of data mining has been inundated with many methodologies (Kurgan and Musilek, 2006).

These methodologies are shown in Figure 2.1 – and the two main ones KDD (which appears as the starting root node on the left of the diagram) and CRISP-DM (a root node in the middle of the figure) are explained in detail later in this section. The main point expressed in Figure 2.1 is that most methodologies seem to be either derived from KDD or CRISP-DM. The diagram represents the genealogy of methodologies used to build intelligent systems. Each box

represents the name of the methodologies, cites their authors and year the study was done. The arrows show the parent-child relationships depicting the lineage of the data mining methodologies (intelligent systems). 5A, KDD, and Six-Sigma methodologies, as represented in hexagons, were unique and did not have any children. The rectangular ones were derived from KDD. The oval are derived from CRISP-DM. In practice, CRISP-DM is the most widely used in the industry (Piatetsky, 2014).

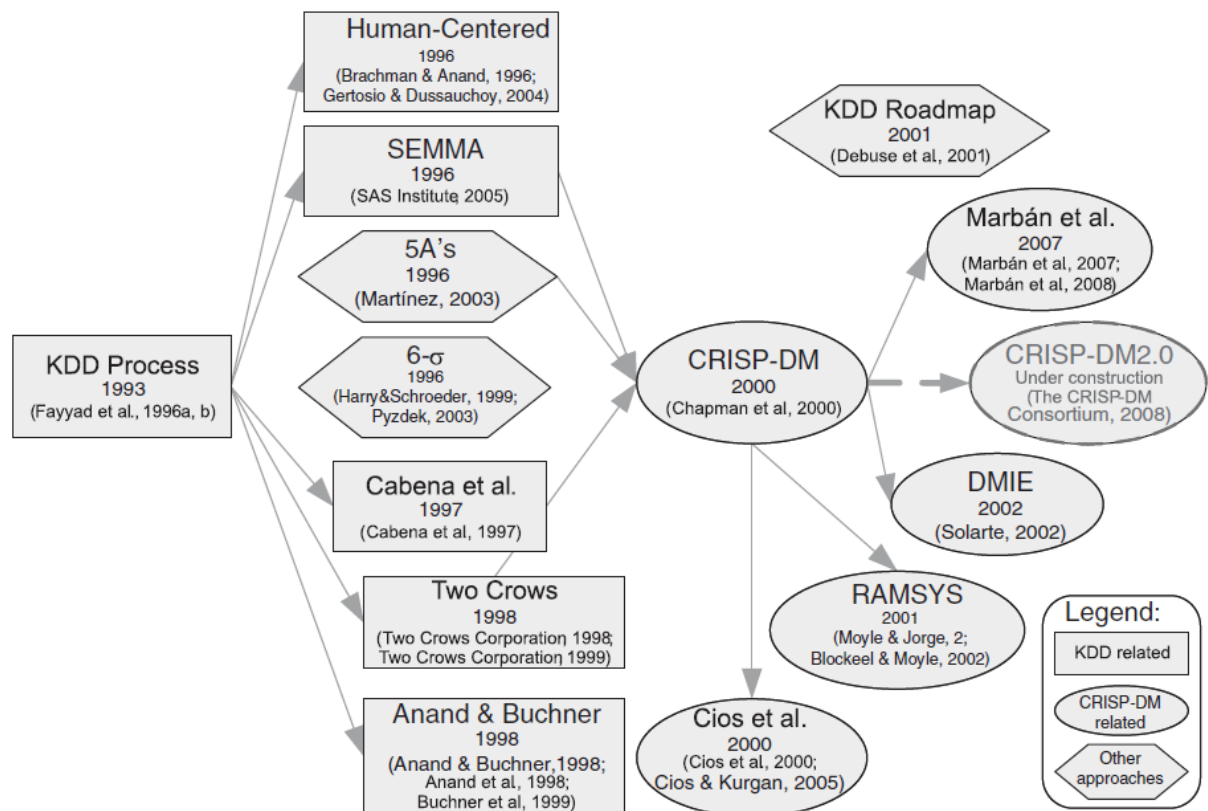


Figure 2.1 Adopted from (Mariscal et al., 2010)

Mariscal et al. (2010) performed an extensive survey of the data mining methodologies (Mariscal et al., 2010). Figure 2.1, adopted from Mariscal et al. (2010), reveals the progression of the methodologies in the data mining field. Keeping in perspective that this research is about management and not computer science, the idea of this section of the critical review is not to dwell into each

methodology from a technical perspective, but to provide an overview to demonstrate the utility and progression of various methodologies. The best way to do that is to track Figure 2.1 and make the following visual observations:

- 1) The above diagram shows the inception of computerized data mining methodology was in 1993 when Fayyad et al (1996) introduced Knowledge Discovery in Databases (KDD) (Fayyad et al., 1996).
- 2) KDD led to various companies, institutions, and academics developing and proposing their own methodologies (for example, SEMMA, Anand & Buchner and others) that were related to KDD. Two 5A's and Six Sigma were developed independently of KDD (Mariscal et al., 2010).
- 3) All the previous efforts – both KDD based and others – were consolidated and merged into a single methodology known as CRISP-DM (Wirth and Hipp, 2000).
- 4) Subsequent to CRISP-DM, five other approaches were derivatives of, or related to, CRISP-DM, and one was developed independently.

From the visual inspection of Figure 2.1, it can be observed that data mining methodologies were initiated by KDD and found their pinnacle with the CRISP-DM. The adoption statistics of methodologies also corroborates the visual as recent statistics report that CRISP-DM is still (as of 2014) the most widely used data mining methodology in the world (Piatetsky, 2014). Given the seminal status of KDD, CRISP-DM, and SEMMA, the methodologies are briefly explained below:

KDD – Knowledge Discovery in Databases

In the 1990's Fayyad et al (1996) identified the problem that data was being converted to knowledge manually (Fayyad et al., 1996). Recognizing the role played by machines, they proposed that a framework is necessary to drive efficiency and effectiveness in data mining processes. They defined KDD as “the nontrivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data” (Fayyad et al., 1996, p.40). Fayyad et al. (1996)

clarified that *data* meant a set of facts, and patterns describing a subset of data or a model applicable to the subset. Thus, extracting a pattern meant “fitting a model to data; finding structure from data; or, in general, making any high-level description of a set of data” (Fayyad et al., 1996, p.41). The pattern, they argued, must be understandable. As shown in Figure 2.2, the KDD process is composed of nine (5 process steps plus 4 results such Target Data, Preprocessed Data, Transformed Data, and Patterns) steps. In those steps, Fayyad showed that the first step is to select data from some given datasets (Target Data). Then to process it such that one is left with data which can be fed into a machine for training. Then to transform the data to make it ready for ingesting in a machine (Transformation). And then applying Data Mining techniques to extract Patterns. This leads to Knowledge creation via Interpretation and Evaluation.

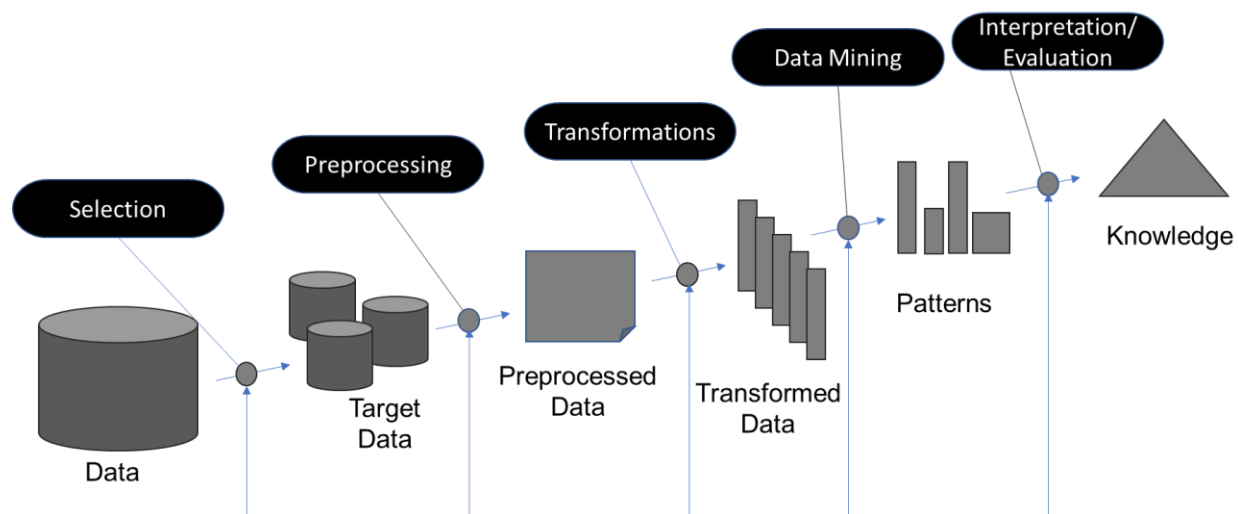


Figure 2.2 Adopted from Fayyad et al., (1996)

CRISP-DM

During the mid-1990's companies and institutions began formalizing data mining and knowledge discovery process. Data mining was becoming more pervasive the benefits of data mining were being recognized. With an EU sponsored initiative, a new standard was launched (Wirth and Hipp, 2000). This was known as CRISP-DM (Figure 2.3) which stands for CROSS Industry Standard Process for Data

Mining. Widely used, this methodology is applied to develop intelligent systems. At the center of this methodology is data (as shown in the middle of Fig 2.3) depicting that intelligent systems are developed from data – as explained in assumptions part of this chapter. Data is then taken through a process of steps that lead to the development of a system. Note that this process is very different than programming-based software development. The process is invoked by first defining the business goals (i.e. the reason why a new system is being created), and is followed by various system construction steps that include data understanding, data preprocessing, training model development (algorithm training), evaluation (testing), and then deployment. We can observe that from the perspective of this research, the first step reflects the process of discovering, defining, and designing a product or a service, whereas steps from data understanding to deployment are the system construction (engineering) steps. CRISP-DM method is composed of the following six steps:

Phase 1: Business Understanding: Phase 1 is based upon understanding the project objectives and requirements for the business, based upon which developing a data mining problem definition, and establishing a project plan. The business problem is formulated in accordance with the available data (Phase 2).

Phase 2: Data Understanding: In this phase initial data is collected, data quality is assessed, insights into data are developed, and hypothesis for hidden information is established.

Phase 3: Data Preparation: Data is prepared as a dataset for ingestion into the model. Several activities are conducted to prepare the data. These activities may include improving data quality, standardizing data, understanding the informativeness of data etc.

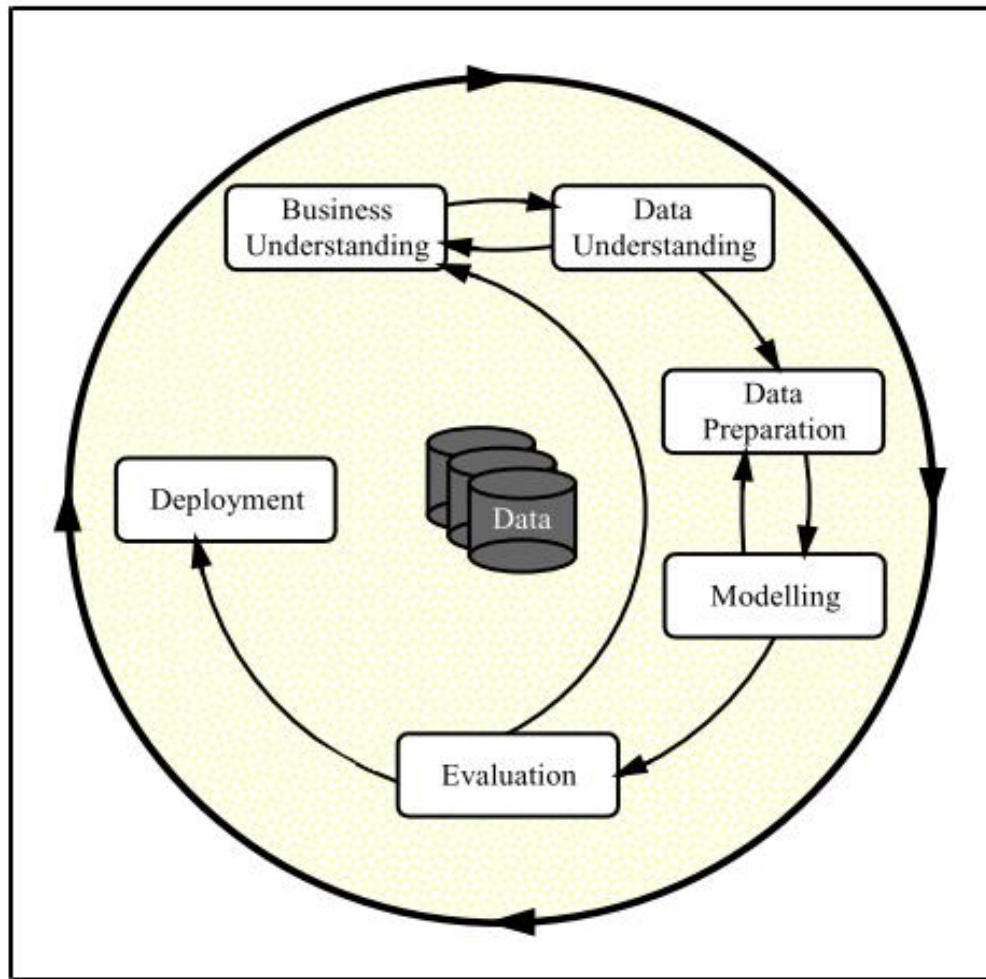


Figure 2.3 Adopted from Wirth (2000)

Phase 4: Modeling: In this phase data is used to train the algorithms. Data from Phase 3 is fed into a mathematical model and as the model gets trained, a mathematical function develops that represents the data. Various modeling techniques are chosen and applied. Different techniques are tried and their parameters are calibrated.

Phase 5: Evaluation: Model is evaluated to check if it achieved the business objective. The use of the results of the data mining can be determined. It can be viewed as testing the learning capability of a model.

Phase 6: Deployment: Once the model works it should be implemented so users can use it.

CRISP-DM is the most widely used in the industry (Piatetsky, 2014).

Other Methodologies

Besides the above discussed, many other models were developed. For example, SEMMA process was developed and proposed by a software company known as the SAS Institute (SAS, 1997). IBM developed its own model known as ASUM and many domain specific CRISP-DM methodologies were developed for sales and marketing , for web mining, finance, and others (Martinez-Plumed et al., 2019). Additionally, data mining methodologies were enhanced to fit into various disciplines including engineering (Wiemer et al., 2019), industrial (Huber et al., 2019), signal processing (Dåderman and Rosander, 2018), and others.

Other methodologies proposed, however, looked very similar to the CRISP-DM methodology. For example, see the IBM Data Science methodology in Figure 2.4 (Rollins, 2015). The key point made here is that IBM's and other companies' methodologies are all patterned after CRISP-DM. The Figure in 2.4 shows the IBM methodology and it is essentially similar to CRISP-DM shown in Figure 2.3. IBM simply breaks down the main steps of CRISP-DM into further sub-steps. For example, Data Understanding is parsed into Data Requirement, Data Collection, and Data Understanding. IBM model also expands the CRISP-DM model by adding a feedback loop at the post-deployment stage (Fig 2.4) which depicts that once an intelligent system is deployed it can continue to improve and train the model based upon its experience. This is how an autonomous car continues to improve its performance from new learning. One important point to consider from the perspective of this thesis is that while IBM expanded upon the engineering or system construction side of the CRISP-DM, the business understanding side appears as a single element, just as in CRISP-DM. In other words, all methodologies were considering that business requirements and

understanding as somehow given and not requiring additional steps or explanation to clarify how does one get to that point? How does one determine what business understanding means in terms of developing a system? By exploring the process as an NPD framework, this thesis will attempt to broaden the *business understanding* part and also reexplore the construction and engineering parts.

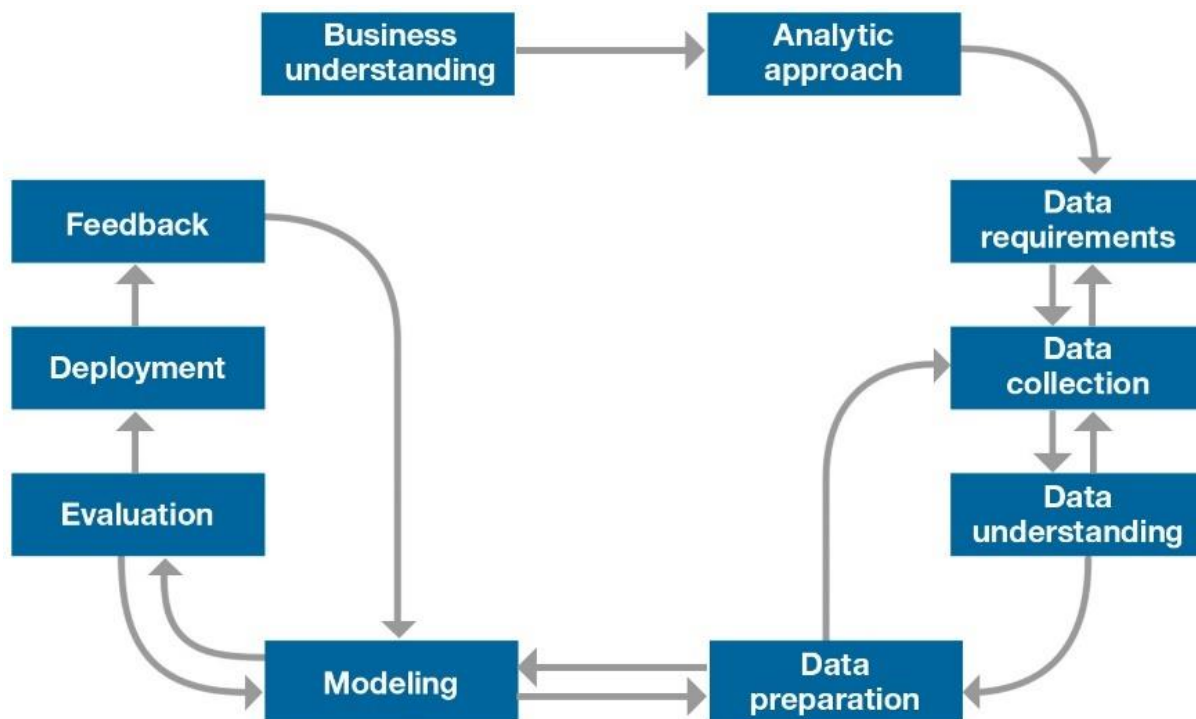


Figure 2.4 Adopted from Rollins (2015), *Methodology for Data Science*

Failures and Limitations of the Existing Methodologies

Analyzing the above methodologies, Martinez-Plumed et al. (2019) determined that all of the methodologies in one way or another are either essentially reflected in the CRISP-DM model or are derived from CRISP-DM (Martinez-Plumed et al., 2019). This finding corroborates with the visual observations (Figure 2.1) and from the statistics presented above. However, more importantly from the perspective of this research, Martinez-Plumed et al. (2019) declared that all such

methodologies “**do not fully embrace the diversity of data science projects.**” (Martinez-Plumed et al., 2019, p.1)

Even though twenty years later, CRISP continues to be the de facto standard of data mining, more recently, however, the explosion in the sources, processing, modalities, formats, volumes, and variety of data has broadened the demand for data centric solutions. Data driven products are being designed (Li et al., 2019; Hesenius et al., 2019). At the center of this revolution is ML. As new business models are being architected, what was once *data mining and knowledge extraction* have now expanded to become an entire field under a new name of: *data science* – and data science is far more than just data mining (Donoho, 2017).

Donoho (2017) defines data science as “This coupling of scientific discovery and practice involves the collection, management, processing, analysis, visualization, and interpretation of vast amounts of heterogeneous data associated with a diverse array of scientific, translational, and inter- disciplinary applications.” (Donoho, 2017, p.4). Donoho (2017) also quotes Provost Martha Pollack of University of Michigan to demonstrate the significance of data science: “Data science has become a fourth approach to scientific discovery, in addition to experimentation, modeling, and computation” (Donoho, 2017, p.4).

The limitations of data mining methodologies were compounded with the rise of data science (Rollins, 2015). These problems led some scholars to spot the inadequacy of CRISP-DM and demand that a replacement is needed for CRISP-DM (Piatetsky, 2014). Comparing the current crises to 1968’s period known as software crises when no standard engineering methods existed for software, Marban et al. (2009) pointed to the increasing demands on the data mining methodologies, and pleaded that “while CRISP-DM was an improvement on the earlier state of affairs, the process model is not perhaps yet mature enough to deal with the complexity of the problems it has to address” (Marbán et al., 2009,

p.88). Pointing the problem of 75 to 85% project failure of ML, Studer et al. (2020) say “one reason is the lack of guidance through standards and development process models specific to ML applications” and explains that “due to the lack of a process model for ML applications, many project organizations rely on alternative models that are closely related to ML, such as, the Cross-Industry Standard Process for Data Mining (CRISP-DM)” – and then goes on to explain the two major shortcomings of CRISP-DM: “First, CRISP-DM does not cover the application scenario where a ML model is maintained as an application. Second, and more worrying, CRISP-DM lacks guidance on quality assurance methodology” (Studer et al., 2020, p.1). Hesenius (2019) also pointed to the inability of CRISP-DS to apply to complex software development (Hesenius et al., 2019). More recently, in 2019, Microsoft engineers produced a paper in which they acknowledged that the current methods are not suitable for ML while offering a potential solution to suggest expanding the traditional CRISP-DM methodology (Amershi et al., 2019). An Amazon (company) team, also warned about the need to approach ML development differently than the traditional software (Schelter et al., 2018).

The shortcomings of data mining are becoming even more obvious due to the backdrop that ML systems require significantly different challenges than conventional software systems. That is what led a team of researchers to declare that “ML systems have all of the problems of non-ML software systems plus an additional set of ML specific issues” (Wan et al., 2019, p.1). A different research team also pointed out that “Software systems that learn from data are being deployed in increasing numbers in industrial application scenarios. Managing these ML systems and the models which they apply imposes additional challenges beyond those of traditional software systems” (Schelter et al., 2018, p.5). This also led Amershi et al. (2019) to highlight the three aspects of the AI domain that make it fundamentally different from prior software application domains (Amershi et al., 2019, p.291):

- 1) *“discovering, managing, and versioning the data needed for machine learning applications is much more complex and difficult than other types of software engineering,”*
- 2) *“model customization and model reuse require very different skills than are typically found in software teams,” and*
- 3) *“AI components are more difficult to handle as distinct modules than traditional software components — models may be “entangled” in complex ways and experience non-monotonic error behavior.”*

This finding was also shared by Rahman et al. (2019) who stated “Software engineering for machine learning applications has distinct characteristics that render most traditional software engineering methodologies and practices inadequate” (Rahman et al., 2019, p.2).

Critical Analysis for ML/Data Mining Methodologies

As discussed in the previous section, the leading voices from across the industry are concerned about the limitations and inability of data mining methodologies to cope with the ever-increasing requirements of data science and ML; it begs the question why these methodologies are not compatible with the needs of modern-day data science.

CRISP-DM, the central data mining methodology, was developed with a limited and narrow scope to solve a single data mining problem. Looking at the introductory paper for CRISP-DM, it can be observed from the language when the author states about the first phase that it “focuses on understanding the project objectives and requirements from a business perspective” (Wirth and Hipp, 2000, p.5). The emphasis on “project objectives” shows that the methodology was not designed for the strategic transformation of a business, it was designed to meet the needs of a project. It was a linear approach, a list of steps that progress in a waterfall style. CRISP-DM did not envision a world where products and services will be created from data (Hesenius et al., 2019; Koschwitz

et al., 2018), where interconnected ML products will perform complex work, where network of intelligent artifacts will form work chains to create competitive advantage for business (Porter and Heppelmann, 2014, 2015), where digital workers will work collaboratively with human workers (McAfee and Brynjolfsson, 2016), where machines will make decisions about humans (McClure, 2018), where entire sectors and the core fabric of the society will be changed (Schwab, 2015; Perez, 2016). The limitations of the simple six step model, as pointed out by many (discussed above), when combined with the absence of a methodology in ML points to a state of affairs that a methodology needs to be more than just a linear process to help develop a project-oriented software.

This means that CRISP-DM (and others) did not touch upon the issues of organization, of business strategy, of a networked economy, of ethics, of governance, or of considerations about the underlying philosophy of information systems development. CRISP-DM was an engineer's take on a world composed of humans, human interaction, social and political dynamics, organizational issues, and everything else that the conventional information systems development recognized as important considerations (discussed in the next section). In that engineer's take, they modeled the reference methodology from a beginning point which did not consider any of the organizational dynamics, to a middle area which remained completely oblivious to the power and political dynamics in an organizational context, and an end that assumed mechanistic positioning of the software via deployment while ignoring the organizational complexities. Thus, data mining methodologies not only lacked the technical bandwidth, an issue that was addressed by Microsoft engineers' paper and other studies cited in the previous subsection but were also greatly deficient in addressing any managerial or organizational issues.

2.4 Systems Development as a Social Paradigm

The inadequacy of the ML methodology applied explained in the previous section has led researchers to compare the current situation in ML to the 1968 software crises when software was being developed with no suitable methodologies (Marbán et al., 2009) and because the entire process lacked/lacks philosophical grounding (Gruner, 2010). It is this crisis that draws attention to the search for solution not as another methodology or engineering solution or a technician's approach, but instead a philosophical approach. Gruner (2010) reminds that it was the crisis of software engineering—not its success—that forced people to engage in its meta-scientific, methodological and philosophical reflections (Gruner, 2010). The information systems industry did not find its solutions only in methodologies, it found its answers in theory. As the meta-theoretical exploration of the structural nature of theory in information systems identified the four dimensions of: *domain questions, structural or ontological questions, epistemological questions, and socio-political questions*, theory plays a fundamental role in information systems (Gregor, 2006). It was claimed that exploring the role of stakeholders, human affairs, the emergence of interpretivist vs. positivism, politics, power, prestige, ethics, and morality were all part of information systems theory.

Unlike ML, conventional information systems frameworks are mature and have a long history of successful application. In this context maturity means they do include organizational and social perspectives. Iivari (1991) clarifies that IS development refers to the analysis, design, technical implementation (construction), organizational implementation (institutionalization) and subsequent evolution (enhancement maintenance) of information systems (Iivari, 1991, 2017b). Iivari (1991) cites Gustafsson and Karlsson (1982) to clarify that the term information system (IS) is used to refer to formal system which provides its users in a certain organizational context with information about a set of topics (Bubenko Jr, J. A. Gustafsson and Karlsson, 1982; Iivari, 1991).

ISD Methodologies, Approaches, and Paradigms

A review of conventional shows that information systems development (ISD) literature made a distinction between *methodologies*, *approaches*, and *paradigms*. Citing Avison and Fitzgerald (1995) and Jayaratna (1994), Iivari et al. (2000) elucidated that over 1000 information systems development methodologies existed at the turn of the century and that it was known as “methodology jungle,” which appears as “a seemingly impenetrable maze of competing ideas and notions” (Iivari et al., 2000, p.180; Jayaratna, 1994; Avison and Fitzgerald, 1995). This overwhelming proliferation, the authors commented, was creating confusion in practice. Partly to make sense of this baffling variety of methodologies, and partly to provide a theoretical foundation on how to think about methodologies, ISD researchers have made a distinction between ISD methodologies, approaches, and paradigms. Approaches and paradigms were invoked by the interest in exploring the underlying philosophical assumptions in ISD (Iivari et al., 1998; Iivari and Hirschheim, 1996; Iivari et al., 2000; Orlikowski and Baroudi, 1991; Klein and Hirschheim, 1987a; Hirschheim and Klein, 1989; Iivari, 2017b).

Methodology can be viewed as the “an organized collection of concepts, methods, beliefs, values and normative principles supported by material resources” (Iivari et al., 2000, p.186). Information Systems “ultimately provide the support for an organization’s networks of information creation, gathering, processing, or storing” (Isaias and Issa, 2015, p.1). Iivari et al. (1998) offer the following narrower and more practice focused explanation whereby a (ISDM) methodology is “codified into a set of goal-oriented procedures guide the work and cooperation of the various parties (stakeholders) involved in the building of an IS application” (Iivari et al., 1998, p.165) and that “these procedures are usually supported by a set of preferred techniques and tools, and activities” (Iivari and Hirschheim, 1996, p.560). Clarifying even further, techniques and tools are “well-defined sequence of

elementary operations which permits the achievement of certain outcomes if executed correctly” (Iivari and Hirschheim, 1996, p.186).

Paradigm scholars observed that ISDM (methodologies) shared many common features and hence could be grouped together. The term *approach* was defined as “ISD approach (ISDA), on the other hand, is interpreted as a class of specific ISDMs that share a number of common features. More specifically, we define an ISDA as a set of related features that drive interpretations and actions in information systems development” (Iivari et al., 2000, p.186). Information Systems Development Approaches ISDA therefore can be viewed as classes of ISDMs. It was clarified ISDAs can exist as independent classes with no member ISDMs where they serve the role of general templates – as using the templates to design new or future ISDMs (Iivari et al., 1998, 2000). ISDAs can also be extracted from existing ISDMs by studying and clustering the attributes of ISDMs (Iivari et al., 1998, 2000). Therefore, approaches were viewed as higher level of abstractions of the practice models. Practice models, directly or indirectly, intentionally or unintentionally, were architected as belonging to these approaches. The approaches, therefore, formed the primary structures from which practice methodologies emanate from or belong to.

Another level of abstraction in ISD is *paradigm* (Iivari et al., 2000). Paradigms, whether defined by Burrell and Morgan (1979) as *meta-theoretical assumptions about the nature of the subject of study* (Burrell and Morgan, 1979) or by Kuhn’s (1970) classic conception of paradigms as “universally recognized scientific achievements that for a time provide model problems and solutions to a community of practitioners” (Kuhn, 1970, p.viii) translate into some commonly shared beliefs of a professional community. Paradigms serve as the underlying philosophical concepts that drive the practice methodologies. As the paradigmatic difference between how natural sciences are approached and how social sciences can be approached was pointed out (Burrell and Morgan, 1979), the ISD

researchers observed that software development is a social process. Thus, it was paradigms that provided the necessary scaffolding for social and managerial factors to be included in what otherwise would have been purely a technical methodology. Paradigms became the superstructures or the meta-structures underneath which *approaches and methodologies* materialized.

For that reason, ISD paradigms turned into the primary unit of analysis for this section.

The seminal paradigm in ISD

An intellectual revolution materialized in the ISD world in 1989 when Hirschheim and Klein's (Hirschheim and Klein, 1989) groundbreaking and seminal work pointed out the presence of paradigms. This work received wide acclaim as demonstrated with thousands of citations as of the writing of this thesis. Adapting the Burrell and Morgan's work on paradigms in science (Burrell and Morgan, 1979) the authors created a corresponding view of paradigms for the ISD field. Burrell and Morgan (1979) had challenged the application of well-established Kuhn's scientific paradigm as too limiting for social sciences. Kuhn's (1970) model may work for natural sciences, but social sciences are different, they claimed. Burrell and Morgan explained that the intersection of two dimensions of subjective-objective and order-conflict dimension offers a more realistic model for social sciences. The subjective-objective dimension represented assumptions about science. The order-conflict about society. When combined, they represented the paradigms for social sciences. The intersection of the two dimensions produced the four spaces (Figure 2.5), each representing a paradigm – paradigm of functionalism (objective-order); paradigm of social relativism (subjective-order); paradigm of radical structuralism (objective-conflict); and paradigm of neohumanism (subjective-conflict).

Hirschheim & Klein (1989) suggested that a corresponding model can be applied in ISD. The authors recognized that ISD is as much a social process as it is

technical. Designing and developing information systems implies that the system is being designed by a human, has users who will use the system, is likely part of some organization, and is in response of some human need (Alter, 2008). Information itself is a source of power. Thus, to view ISD as a technology from a science lab is inconsistent with on the ground realities of IS. Unlike other machines or technologies, which may involve the political and social considerations after their deployment, ISD must consider those factors at the inception. The recognition that ISD is a social process led to the recognition that the positivist scientific paradigms designed for natural sciences are too limiting for ISD.

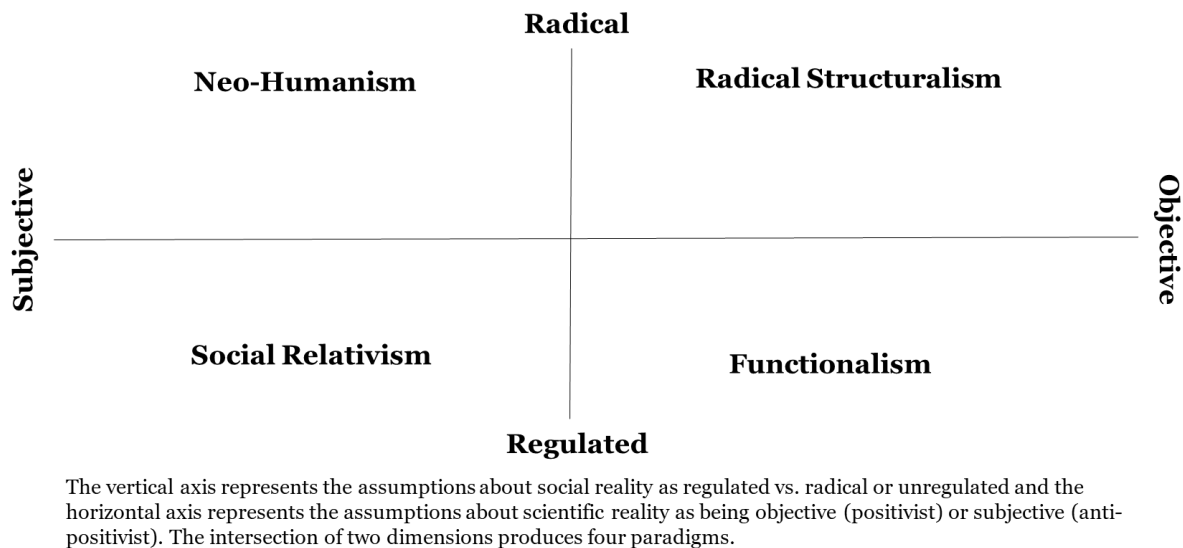


Figure 2.5 Four Paradigms Adopted from Hirschheim and Klein, 1989

Hirschheim and Klein (1989) acknowledged that the real challenge was to demonstrate: *how the paradigms are actually reflected in ISD* (Hirschheim and Klein, 1989, p.1202) and that “the paradigms are largely implicit and deeply rooted in the web of common-sense beliefs and background knowledge which serve as implicit “theories of action” (Quine and Ullian, 1979; Hirschheim and Klein, 1989, p.1202).”

Patterned after the Burrell and Morgan model, Hirschheim and Klein develop their model from the integration of objective-subjective and order-conflict dimensions. The integration leads to four paradigms as follows (Hirschheim and Klein, 1989):

Functionalism (objective-order): In this paradigm the world is viewed as objective and ordered. Systems are developed by experts (akin to Platonic philosopher king). The experts develop systems in accordance with specific methodologies and in a structured way. The requirements are determined by the management teams who approach them in strict accordance with business needs of shareholder value creation. The epistemological roots of this paradigm are of positivism and the tools and methods used are considered as rational – implying tools are capable of extracting patterns of reality that are assumed to exist. The information system is defined by elements such as people, hardware, software, rules as physical or formal objective entities.

Social Relativism (subjective-order) development is led by a catalyst or facilitator. The system is in response to subjective understanding of the organization and includes factors such as cultural sensitivity as internal forces of evolutionary social change are reflected through systems. Shared meanings, metaphors, symbolic structures, and sensemaking play important role. Ethnographic approaches can be used to acquire knowledge about design requirements.

Radical Structuralism: Unlike the above two paradigms, radical structuralism does not ignore the elements of social conflict, albeit it approaches them from the objective criteria and knowledge orientation set by the economic theory. The classic conflict between management and workers (as played in Marxian or Hegelian tradition) is reflected. Here the analyst can assume one of the two roles: representing the interests of the management team or representing the interests of the workers.

Neohumanism: While presenting this section Hirschheim and Klein (1989) write that “whereas the others can be observed in actual systems development cases, this story is hypothetical to a large degree and that it has been constructed from theory” (Hirschheim and Klein, 1989, p.1207). They describe this area as concepts of work, mutual understanding, and emancipation – the three fundamentals around which society and social organizations are arranged. To further clarify the concept, they make a distinction between human work as being the first area of knowledge. Developing mutual understanding from symbols and shared meaning as second area of knowledge. However, they point out that both work and shared meaning from mutual understanding do not protect humans from bias, injustice, and quality of human condition, or lack of truth. That is where emancipation becomes the third knowledge area whose purpose is “the establishment of truth and justice as the norm to regulate all human affairs” (Hirschheim and Klein, 1989, p.1208).

While the Four Paradigm model was seminal from the theory side, much work was needed to study its practical assumptions. Basing his analysis on (Hirschheim and Klein, 1989), and while expanding their model, Iivari in 1991 (Iivari, 1991) showed that the paradigms were applicable in practice and ISD covers all the areas of lifecycle of information systems and ISD implementation is also viewed as an instance of organizational development (OD) (Iivari, 1991).

Building upon the previous literature related to paradigms, Iivari (1991) developed the four dimensions for analysis for paradigms and based them upon Tornebohm (1976) and Burrell and Morgan (1979) (Törnebohm, 1976; Burrell and Morgan, 1979). In 1996, five years after Iivari’s (1991) paper and seven years after Hirschheim and Klein (1989) paradigm framework, Iivari and Hirschheim joined forces to revisit the topic of paradigms (Iivari and Hirschheim, 1996). There was a clear vacuum in two areas: First, despite Hirschheim and Klein (1989) paradigm framework, a thorough treatment of the relationship between

the social and organizational context and ISD approaches was still lacking. Second, further analysis was needed to show the link between paradigms and approaches – with approaches being the next level in hierarchy of ISD analysis after paradigms. To confront these challenges, Iivari and Hirschheim (1996) placed their emphasis on the two underlying assumptions of approaches: (1) the assumed organizational role of information systems, and (2) the view of information requirements. In the case of the first assumption, it distinguishes three alternatives: a technical view, a sociotechnical view, and a social view. In the case of the second assumption, they explored three alternatives: an objective view, a subjective view, and an intersubjective view. With Iivari and Hirschheim, the efficacy of the four paradigms was well-established (Iivari and Hirschheim, 1996).

In 1998, Iivari and the two original developers of Hirschheim and Klein (1989), came together to analyze several other approaches in light of the paradigms. Going beyond just functionalism paradigm centered approaches, the three authors analyzed contrasting approaches that stemmed from the remaining three paradigms (Iivari et al., 1998). In 2000, the three authors again came together to present their four-tiered model which linked four successive tiers of paradigms, approaches, methodologies, and techniques in a hierarchical relationship where each tier inherited properties from the above tiers above it – with paradigm being at the top (Iivari et al., 2000).

The seminal models developed during the 1990's and early 2000's continued to influence future research on multiple fronts (Iivari, 2017b; A. S. Lee et al., 2015). With their pivotal place in the literature, they were recently used to blend Design Science Research (DSR) and Behavioral Science Research (BSR) with ISD (Iivari, 2017a; Iivari and Kuutti, 2018; Iivari, 2019; Friedrich et al., 2017). They were also deployed to analyze the cultural norms and practices in ISD (Friedrich et al., 2017).

Critical Analysis of Social Paradigms in ISD

Iivari's (1991) application and expansion of Hirschheim and Klein (1989) leaves the door open for questioning how the paradigms apply in ML. ML, like any software development, is a social process and therefore ISD paradigmatic inquiry, whether by Hirschheim and Klein (1989), or Iivari, should be applicable to ML as well. However, the peculiarities and subtleties of ML (see Assumptions in this chapter) makes it different than conventional IS and therefore the differences between the two types of systems cannot be discounted. As Section 2.2, Section 2.3, and parts of Section 2.4 have extensively covered the differences between conventional and ML system, the key question is in what ways the four paradigms apply to ML.

Revisiting the four paradigms, the following can be observed in accordance with ML:

Functionalism: Functionalism from Hirschheim and Klein (1989), assumes order and objectivity. Developing ML is a complex process which deals with significant data, statistical modeling, and algorithm-based optimization. The process of finding a solution is not deterministic. A designer must deal with questions such as: does the firm has the relevant data for the problem being solved? Is it even possible to model the problem? Would a certain algorithm work? Would the algorithm perform to have an acceptable solution? Thus, the engineering process of data science (ML) is not deterministic and therefore unlike typical engineering projects, the design elements cannot be engineered with assurance. Thus, raw objectivity and order do not contribute to the process. Secondly, if machine is viewed as part of a social network tasked with performing work, then it is not sufficient to argue from the vantage point of a human. Exploring machine objectivity and subjectivity assumptions would become critical factors for analysis.

Social relativism: Social relativism also suffers from the same constraints as functionalism. The assumption of order is unrealistic given the dynamics of the systems. Machine's reality is shaped by data and algorithms. This reality may evolve constantly and may not be easily classified as objectivity and subjectivity. A machine may not know the difference.

Radical Structuralism: The objective-conflict representation may appear more compatible with ML, however, the application of this paradigm – whether in conventional or ML system – suffers from a major problem. This problem comes from the inability of the designer to objectivity assess, and then do something about, the conflict. The objective role, as Hirschheim and Klein (1989) clarify, is for the designer to decide if he or she will represent the interests of the managers (powerful) or the workers (weak, oppressed). If social institutions are viewed as formal or informal rules that constrain individual behavior and to make rational choices individuals must form expectations about their own and other's behavior (Knight, J., & Jack, 1992), then the use of information becomes critical to form those expectations. Denying such an information to a certain party implies that such a party will be placed in a position of disadvantage. Thus, objectivity assumption when combined with conflict produces a system's view where the designer is always in a position of betraying one of the two parties in the conflict. With ML, where information providing is combined with immensely superior predictive ability, decision-making capacity, and action-taking – the power differential between parties can exceed than that of conventional systems. With the exception of suggesting that the designer will have to take one or the other sides, Hirschheim and Klein (1989) offer no help in resolving the conflict.

Neohumanism: Hirschheim and Klein (1989) present this paradigm and clarify the need for emancipation – the third form of knowledge needed for truth and justice – but they also acknowledge that they have not seen such an application of neohumanism (a story of emancipation) in practice. The greatest problem with

this paradigm is that despite the tremendous and powerful success of the information systems sector, what appears to be the most important aspects of human life – truth and justice – are not reflected in any real stories.

Section 2.3 analysis demonstrate that approaches and methodologies are not independent of paradigms, as they inherit their properties and features from paradigms. These abstractions of paradigms are manifested in practice and the underlying goals, principle, and assumptions filter down to practice level (Iivari et al., 1998). Iivari et al. (1998) recognize the possibility there could be approaches that may not have methodologies associated with them and in those cases, they serve as templates for future methodologies. The underlying assumptions are based upon the designer's understanding of the reality and how to obtain knowledge about that reality (Hirschheim and Klein, 1989).

Hirschheim and Klein (1989), and Iivari et al. (1998) after that, do not assume a world where information systems are intelligent and capable of perceiving, learning, and reasoning. The paradigmatic assumptions are based upon the view of the reality from the perspective of a human. A human can be objective or subjective, can view order or disorder, but that all happens from the cognitive frameworks of humans. The challenge becomes when it is acknowledged that intelligent machines are now an integrated part of the business fabric. Since machines have different cognitive frameworks and a different view of the reality, their reality must be explored from their viewpoint and not from a human viewpoint. This, of course, is being argued on the backdrop of the fundamental focus area of this research that: in ML machines program themselves. Referring back to the definition offered in *the Assumptions and Fundamentals* section of this review, ML is when computer programs learn “without being explicitly programmed” or that ML programs are “from” data and not “for” data.

This creates a paradigmatic predicament. This predicament, that can be termed as the “intelligent machine predicament”, which refers to the problem of studying

and defining a paradigmatic structure that can be suitable for both – humans and intelligent machines. Stated another way, human-machine paradigm recognizes that humans and machines form an interactive and interconnected social network, that both can display certain degrees of intelligence, and that both can experience different realities. Since information system is viewed as having an ontological, epistemological, and axiological structure (Allen and Varga, 2006), an information system that is developed by machines or that takes shape by ML implies that it has an ontological, epistemological, and axiological structure. Since machines have a different cognitive, physical, and evolutionary structure than humans, the perception of machines cannot be the same as of humans (Russell and Norvig, 2016). Attempts to anthropomorphize machines do not imply that the machine reality becomes similar to the human reality (Araujo, 2018). Besides the biological reasoning, the fact that machines approach knowledge differently than humans (as shown in the definition of data science in Section 2), the reality of machine and the reality of humans could not be treated as same. While human and intelligent machine realities cannot be treated as equal or similar, it could be possible to develop a paradigmatic framework that captures the two realities and coalesces that into an integrated framework that expresses them as a collective human-machine paradigm.

The above analysis point to the enhancements that will be needed to realign the conventional paradigms with ML. This is further developed in the Synthesis section.

2.5 Section 4: Synthesis

Building upon the critical reviews performed in Section 2.2, Section 2.3, and Section 2.4, this section begins by summarizing the analysis from the previous three sections. The two thematic structures of practice centric ML methodology (Section 2.3), paradigms and critical analysis (Section 2.4) are integrated and synthesized into a single theme. That theme is carried forward to formulate an

integrated framework of paradigm and ML methodology, thus eliminating the *paradigm to ML methodology* gap pointed out in Section 2.4 (critical analysis). Substantial and original theory-building and analysis are undertaken to form the basis of the conceptual framework (next section). The conceptual framework section provides the foundation for formulating parts of the framework upon which action research is conducted.

To summarize the insights from the Section 2.2, 2.3, and 2.4, we can make the following four observations:

First, the Literature review points to the dearth of methodologies for ML. The existing data mining centric methodologies appear too narrow and incompatible with the growing requirements of ML.

Second, introducing more methodologies to the plethora of existing methodologies does not appear as a helpful option. Instead of methodology proliferation, learning from, and applying, the industry's experience from the conventional ISD can help bridge the gap. What worked in the past was discovering and focusing on paradigms. Approaches, methodologies, and techniques can inherit features from paradigmatic structures. Paradigms have the underlying philosophical assumptions and unlike engineering methodologies that exhibit technical features, paradigms include broader human concerns, organizational issues, sociopolitical perspectives, and other elements that recognize that technology is a social process.

Third, it was recognized that the existing ISD paradigms were developed for conventional information systems. The literature review indicated that ML is a significantly different technology than conventional software – which raises the question of what changes will be needed to align conventional paradigms to the ML methodology.

Fourth, in the critical review section for Section 2.4, it was suggested that using the paradigms developed for unintelligent machines will not be helpful when the new business reality is based upon an interactive intelligent machines and humans relationship. This requires building a paradigm centric model for intelligent machine paradigm.

As was pointed out, both Burrell and Morgan (1979) and Hirschheim and Klein (1989) models view the reality as composed of scientific reality (objective vs. subjective) and social reality (regulated vs. radical). Our conceptual framework building will also use the two dimensions of scientific and social realities.

However, the assumption of scientific reality being objective or subjective or social reality being regulated or radical necessarily requires some underlying epistemological, ontological, and axiological assumptions. For example, if you believe that reality is objective then you are assuming that ontology is based upon realism, epistemology is based upon positivism, human nature is based upon determinism, and methodology based upon nomothetic model. But if you assume that scientific reality is subjective, then ontology is based upon nominalism, epistemology based upon anti-positivism, human nature on voluntarism, and methodology based upon ideographic model. Similarly, on the social sciences side, Burrell and Morgan (1979) identify forces of cohesion, solidarity, consensus, reciprocity, stability, co-operation, integration, and persistence on one end of the spectrum and the forces of conflict, lack of regulation, coercion, disunion, hostility, and change on the other (Burrell and Morgan, 1979) .

Analyzing the underlying assumptions are deemed necessary since this literature review has established: 1) the overwhelming presence and pervasive inclusion of AI in human situations in modern day business (and personal) environments; 2) intelligence, by definition, requires epistemological and ontological assumptions; 3) and since both epistemological and ontological assumptions necessarily imply

some concept of reality (Burrell and Morgan, 1979), intelligent machines must be based upon assumptions about social and scientific “reality.”

Hirschheim and Klein’s (1989) model was derived from the Burrell and Morgan (1979) four paradigm model. The Burrell and Morgan model has dominated the social sciences field and even the critics acknowledge the model’s legendary following and elegance (Deetz, 1996). As pointed out in this review, Hirschheim and Klein (1989) developed their model for ISD when systems were viewed as deterministic, non-autonomous, and were deployed to meet the human information needs. However, the advent and adoption of AI/ML require a revision of, or improvement in, the underlying assumptions. From the inception, this research recognizes that the modern society is composed of a spectrum where intelligent machines play a role in human life and that role can be of a collaborator, coworker, supporter, or helper on one end (for example Siri in iPhone) to a fully autonomous role on the other end (algorithmic trading, autonomous car) – this was discussed in Assumption 3.

This area of the review is primarily patterned after Iivari’s 1991 paper where he applied the Hirschheim and Klein’s (1989) four paradigm model to study seven information systems development approaches (ISDA) (Iivari, 1991). That paper made two major contributions. First, it was a pragmatic attempt to link existing approaches with paradigms to demonstrate that approaches do in fact represent paradigmatic properties. The seminal nature of the paper was confirmed as seven years later in 1998 both Hirschheim and Klein, the original developers of the paradigmatic approach, joined hands with Iivari to expand the Iivari analysis to five more approaches (Iivari et al., 1998). Second, unlike Hirschheim and Klein (1989) who did not provide deeper analysis of ontological and epistemological assumptions, Iivari (1991) filled the gap by giving a detailed account of those assumptions as well as expanding the framework to include greater depth from including *methodology* and *ethics* (explained below). The approach used in this

synthesis is to briefly analyze each underlying philosophical assumption in accordance with the conceptual framework of machine-human we are attempting to develop in this research.

In 1991 paper, Iivari's elegant characterization of Hirschheim and Klein (1989) extracted the four features of the paradigms: epistemology, ontology, methodology, and ethics. Iivari was less concerned about the intersection of the regulated-radical and subjective-objective and more concerned about clarifying the underlying philosophical assumptions. He took each of the features and then expanded it to analyze the underlying assumptions. I followed a similar approach – albeit with application in ML.

THE ONTOLOGICAL ASSUMPTIONS FOR MACHINES

From an ontological framework perspective, Iivari (1991) identified the five perspectives of: view of information/data, view of information/data system, view of human beings, view of technology, and view of organization and society (Iivari, 1991). In the discussion below, each of the underlying philosophical assumption is analyzed in accordance with the backdrop of ML (note first two are integrated as one):

Data and Data System: Data plays a key role to shape the reality of a machine and hence is central from an ontological perspective. This assumption retains the factual and constitutive positions of data since humans are actively involved with a machine, but the machine functions as data agnostic. Machine is not assigning factual or constitutive values to data – it is simply assigning the mathematical value it observes from the data, the context in which machine is learning, and the mathematical function (algorithm) that represents the data. If a machine is fully autonomous, its factual vs. constitutive classification is a value agnostic position – as a machine does not know the difference. Machine ontology therefore is dependent upon the data that is fed to it. It does not question its reality, for example based upon values or virtues, as humans do.

Human (and machine): One way to view a human from an ontological perspective is to observe them in terms of determinism vs. voluntarism positions as argued by Burrell and Morgan (1979). Both machines and humans can also be viewed as necessarily “controlled by or controlling” each other. The term control includes factors such as influencing and supporting in roles such as trainers, collaborators, coworkers, and monitors. However, how machines perceive humans is limited to viewing humans as data and objects of action. For instance, when an autonomous car sees a human on the road views that as data (source of data) and then stops to avoid an accident (object of action).

Technology determinism: Technology determinism implies that humans retain control over machine decision-making. In autonomous machines, machines have full control of their decisions, and they may exhibit self-learning, evolutionary learning, and reusable learning. ML and adaptive systems is an emerging area of research (Farhi et al., 2014; Quin et al., 2019; Jamshidi et al., 2019).

Organizational: The last ontological assumption recognizes humans and machines will form collaborative clusters and networks of social interaction and that humans will retain their positivist and anti-positivist orientations. Machines will function in a network that includes humans and machines but overwhelmingly include working with other machines. These relationships are expressed as part of Swarm Intelligence (Aydin and Fellows, 2017, 2018).

THE EPISTEMOLOGY OF MACHINES

In the paradigms offered by Burrell and Morgan (social) and the one proposed by Hirschheim and Klein (social + ISD), epistemological structures were based upon the dichotomy of the two opposing models of positivist and antipositivist (Iivari, 1991). In the positivist tradition, reality is assumed to be hard, immutable, and independent of the observer. Knowledge, in that tradition, is acquired as factual and is derived from experimentation. The process is based upon the standard

scientific method of hypothesis generation, data collection, model selection, and testing.

ML comes from the combination of statistics and computer science and while it keeps a bit of both, it is not exactly statistics and it is not exactly computer science (Wheeler, 2016). The biggest difference between statistics and ML is that unlike statistics ML does not assume the presence of a human defined hypothesis. Working through billions or even trillions of possible variables and oceans of data constitute a different type of problem solution relationship – a relationship for which epistemology was not ready. Epistemology was still struggling with the inclusion of statistics, as Wheeler (2016) pointed out, not exactly sure how to assess the role of probability in philosophical inquiry.

ML can be viewed as accelerated and broad scientific process which can both validate science and also discover what was not sought in the inquiry. This self-inspired learning and discovery is unique to ML.

Thus, in organizations where machines are contributing to knowledge and sensemaking (M. K. Lee et al., 2015; Krush et al., 2013; Shoham et al., 2018; Brynjolfsson et al., 2017) and as organizations are viewed as composed of the interaction of humans and machines or solely of machines, the epistemological structure would need to be revised or upgraded.

On the human and machine end, the positivist and anti-positivist standards are still applicable. For the same reasons as argued by Burrell and Morgan (1979) and Hirschheim and Klein (1989), the presence of humans in the mix would necessarily imply that humans will view the social structures in terms of positivist and antipositivist knowledge orientations.

The epistemological assumptions are derived from the underlying process of ML, as well as from the fundamental assumptions of Burrell and Morgan (1979).

Research Mode: Provost Martha Pollack that “Data science has become a fourth approach to scientific discovery, in addition to experimentation, modeling, and computation” (Donoho, 2017, p.4). On the autonomous end, machines operate independently and with active and unrelentless exploration (Martinez-Plumed et al., 2019).

Uncertainty: The uncertainty embedded in human-machine system is approached with the goal of resolving it such that a state of certainty is approached. On the autonomous end, uncertainty is exploited.

Progression: Progression implies the basic functioning of the epistemological process. On the human-machine end, knowledge is discovered through experimentation where hypothesis and model must precede finding. On the autonomous end of the spectrum, knowledge results from exploration. In that case, knowledge also results from serendipity and randomness and at times no theoretical explanation exists to understand the basis for the link between input and output.

THE METHODOLOGY OF A MACHINE

Methodology implies the underlying representation model of epistemology. Iivari’s ISD framework is based upon three methods of constructive, nomothetic, and idiographic (Iivari, 1991). However, Iivari treats them as mutually exclusive. Constructive development, a subcategory of constructive method, Iivari explains, implies that it does not describe any existing reality but instead tries to create a new one. Iivari explains that idiographic methods can be approached both as positivist or antipositivist.

From a machine methodology perspective, a machine models reality based upon a mathematical model – and therefore by design it is nomothetic. However, it is nomothetic in the sense of how its model absorbs the data or how data is ingested into the model. Underneath the nomothetic structure, data can be of any form –

hence pictures, voice, text, videos, IoT, and other forms of unstructured data are converted to numerical representation. This implies that the data from any type of methods mentioned by Iivari (1991) – including idiographic research methods (for example case studies, action research) and constructive research methods can be processed by a machine. In terms of methodology therefore, we can conclude that *autonomous machines* and *human + machine* constructs can both support these various methodologies.

THE ETHICS OF A MACHINE

Burrell and Morgan (1979) began their analysis by pointing out that society can be viewed as ordered or in conflict. Perhaps to make the terminology gentler, they used the terms regulated vs. radical/unregulated. Hirschheim and Klein (1989) chose to use the term conflict. However, the concept of social struggle does not get fully addressed by Hirschheim and Klein (1989). Iivari (1991) boiled down the ethics to three things in information systems (IS) science as means-end oriented, interpretive, and critical. As second subcategory under ethics, Iivari also included values of IS research as organization/management, user-oriented, and others.

Analyzing the nature of society in the autonomous machines' era, the core elements of order and conflict can be retained but only if they can be expanded to include some important issues that have been ignored in social conflict but that can resurface with ML.

A detailed study of ethics in ML is out of scope for this research, however, the following assumptions can be made for ethics:

- 1) The social conflict –would not disappear with the emergence of ML. In fact, there are reasons to believe that ethics would become one of the most important concerns in the field of ML and AI (Hagendorff, 2020). Trust will be an important issue (Hengstler et al., 2016a, 2016b)

- 2) Governance in AI will be an extremely important consideration from two angles: 1) governance of firms by using AI; and 2) governance of AI itself (Nobre, 2012; Lauterbach and Bonim, 2016; O'Meara, 2012; Tonn and Stiefel, 2012).

The above synthesis leads to the development of the conceptual framework in the next section.

2.6 Section 5: Conceptual Framework

In light of the discussion from the Synthesis and the supporting analysis from Sections 2.2, 2.3, and 2.4, a conceptual framework will be developed in this section. A conceptual framework helps in the development of the methodology (Chapter 3), by outlining and boundaries of what will be investigated in the action research. It gives structure to the research and provides a scaffolding to support the field research. In addition to providing support for the research methodology, a conceptual framework can also exhibit the novelty of a research study. To build the conceptual framework, the following concepts developed in Sections 2.2 to 2.5 will be used:

- 1) A human-machine paradigm will be developed bottom-up to signify the social structure composed of humans and machines. The term bottom-up implies that it will be developed based upon the lower level philosophical assumptions as pointed out by Iivari (Iivari, 1991).
- 2) The data mining methodology framework, which is being widely used in ML, is enhanced in accordance with the paradigmatic assumptions.

Just as in Burrell and Morgan (1979) model, I conceptualized two dimensions. One that depicts the assumptions of science. Second that shows the assumptions of society. The science dimension is composed of subject-objective categorization. The subjective assumption characterizes ontology based upon nominalism, epistemology based upon anti-positivism, human nature as voluntarism, and

methodology as ideographic. The objective assumption characterizes ontology as realism based, epistemology as positivism, human nature as determinism based, and methodology as nomothetic.

From an ML perspective, the above assumptions extend beyond humans and include intelligent machines. In other words, machines can be viewed as purely subjective and purely objective based upon their learning models.

The second set of assumptions are about society and include society of regulation and society of radical. The society of regulation includes factors such as status quo, social order, consensus, actuality, order, stability, integration, and functional coordination. The society of radicalism is based upon conflict or coercion and includes change, structural conflict, modes of domination, potentiality, and disintegration.

The intersection of the two dimensions – assumptions of science and assumptions of society – produces four paradigms for ML (Figure 2.6). They are explained below.

Machine Functionalism

The underlying assumption of machine functionalism is that the designer is assuming a regulated world that is free of conflict and that the machine can approach its reality with objectivity. An ML product designed with these assumptions will model reality with a specific and limited problem-specific dataset and expect consistency in the underlying data distributions and the environment. Furthermore, once its learning is complete, it will no longer require additional learning since the world is assumed to be regulated and unchanging. In a predictable and regulated world, the product does not require governance since at the time of its creation, all such assumptions would be included in its design. For example, recommendation engines are intelligent software designed to recommend products to customers. A recommendation engine that is

developed under the Machine Functionalism will be designed based upon limited set of features (group's preferences, purchase history) and make recommendations based upon that. Since it is assuming that people are objective, and that its own model of reality is objective, it will base its recommendation criteria on objective assessment of preferences and previous buying behavior. Based upon rational (objective) consistency, it will also group people into broad categories – for example, it will group movie customers into those who like horror, drama, comedy etc. It will expect these groups to stay stable and the behaviors within these groups to stay consistent. It will recommend horror movies to those who fall into horror group and comedies to those who fall into comedy. A recommendation engine designed with functionalism paradigm will exhibit feature of an intelligent software designed and deployed for a world that is stable and its reality can be ascertained objectively. Consistency and stability also imply that the groups identified by the engine will not exhibit broader external influences upon the members of the group. From a ML perspective, such product designs could be based upon classical ML algorithms and supervised learning.

Machine Radical Structuralism

Here the assumption for ML software development is that while the reality can be objectively modeled but it is more complex than what a limited problem specific dataset can capture. This is because the designer assumes that the reality is being shaped by social (or other) forces that are powerful and influential. This means that the software designer acknowledges that he/she cannot model the system without including the broader social factors – however, he/she also assumes that it is indeed possible to objectively analyze and include these forces within the design. For example, when a designer conceptualizes a recommendation system, he/she assumes that the individuals for whom recommendations are being made are not necessarily representing their own preferences. Their opinions result

from the surrounding forces. Those forces can be external firms or other individuals with power and influence. They can be political campaigns or social trends. They can be marketing campaigns and ads. However, the designer believes that these factors are real influencing factors and that they can be objectively identified and determined. This has significant implications when designing recommendation engines. The assumption that the reality of system is now shaped by objectively determinable external powerful or influential forces implies that broader datasets are needed for a recommendation engine to make recommendations. It must capture the reality of broader forces – for example, of marketing campaigns, of advertisements, of broader social trends. The software is no longer concerned about the group preferences as a collection or a set of members with similar tastes and preferences – but instead it assumes that the external influences shape the opinions and preferences. To recommend the preferences of a group, the engine must first track the influences upon the groups. This implies that broader datasets about social (or other) trends or forces will be needed to train algorithms. However, since the paradigm assumes objectivity, the influences are assumed to be stable and consistent. This means that ongoing learning will not be needed. The patterns of influences discovered and learned by the algorithm will sustain and prevail.

Machine Social Relativism

The assumptions under this paradigm are that the reality is subjective to individuals and there is no external influence or conflict. The system designed to model such a reality assumes that the reality will vary and that different versions of reality may exist simultaneously. However, while the subjective dimension gives us the reality as dynamic and changing, this paradigm maintains that the dominant powers do not shape the reality for systems (or individuals). This means that the power structures with a society or organization do not influence the reality being modeled by the system. For example, a recommendation system

designed to help people make decisions will attempt to categorize individual likings based upon their own mood, circumstances, experiences, emotional states, and behaviors. For that system, the idea will not be to model a shared reality across groups of individuals – for instance dividing humans into groups who prefer different types of movies such as drama, horror – but instead the system will consider each individual as a unit and work with the underlying emotional structure of the *individual* and recognize the emotional states of a person when he/she will prefer a drama vs. a comedy movie. The subjectivity assumption, in this case, gives each individual a customized and unique place in determining his or her reality – and for the system to approach reality as shifting, changing, variable, and transitioning. For instance, data on a person's moods, emotional states, geospatial location, weather – all become input to make a recommendation. The person (customer) here is not being grouped into a logical segment of objective decision-makers. His or her subjectivity is being acknowledged as the primary contributor to decision-making. Similarly, an autonomous car designed under this paradigm will view driving conditions subject to the human's emotional states and may take the longer (suboptimized) route home which will consume more energy and take longer time to complete the journey but will pass through a scenic route that the human rider prefers. The car will learn to observe human preferences in accordance with human subjective states. Here the ML system models the world from a subjective reality basis but assumes that the individual's own data is sufficient to model. Data used for this type of modeling will be multidimensional, broad, and could include significant unstructured data. Furthermore, the algorithm would need to continue to learn – since subjectivity is transitional.

Machine Humanism

This is a unique paradigm where the designer assumes that the system will view reality as subjective and will attribute that to dominant broader external forces. If

the reality is being modeled for humans – for example in a recommendation system – the system will assume that human needs and wants will constantly shift due to human subjectivity (moods, emotions, etc.) however such an emotional preference structure is being influenced by external dominant forces such as political, social, or economic stressors and factors. This assumption implies that if the job of the recommendation engine is to make movie recommendations, it will base that decision upon not only human subjective states but also upon broader *influence forces*. For instance, people could have developed a very different movie genre liking during the Covid lockdown (external influence). In this paradigm the system assumes that human decision-making is based upon dominant factors which give rise to human preferences. A system designed upon those assumptions will model reality as constantly shifting, influenced by broader dominant factors, and driven by human emotions. For example, a system that tracks political viewpoints will use the data from dominant power structures in the society and will use that to model decision-making where it will create (even make up) stories that will cater to certain human moods and emotions, but the themes will be derived from dominant power structures and narratives.

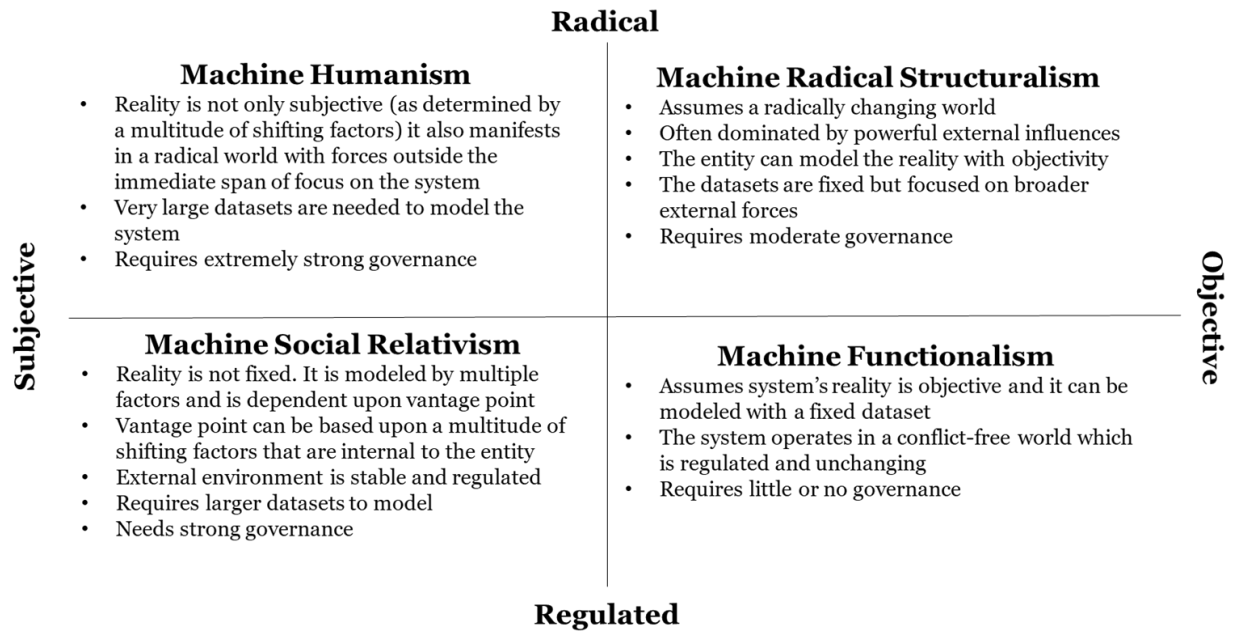


Figure 2.6 Conceptual Framework Paradigms

In order to explain the application of the mode, I present two examples of how the conceptual framework can lead to product concepts (ideas) that can produce significantly different pathways in the subsequent NPD process.

Example 1: An intelligent recommendation system

Retailers like Amazon and Walmart are developing and deploying ML based recommendation engines. Since the systems are intelligent, the application of the four paradigms will lead to very different types of products (Figure 2.7). A product with machine functionalism assumption will be based upon data collected and categorized from groups with similar interests. This assumes objectivity and regulation. This product will require no governance since the underlying assumption is that it is operating in a stable and conflict-free environment. The other manifestations of the recommendation engine have been explained above.

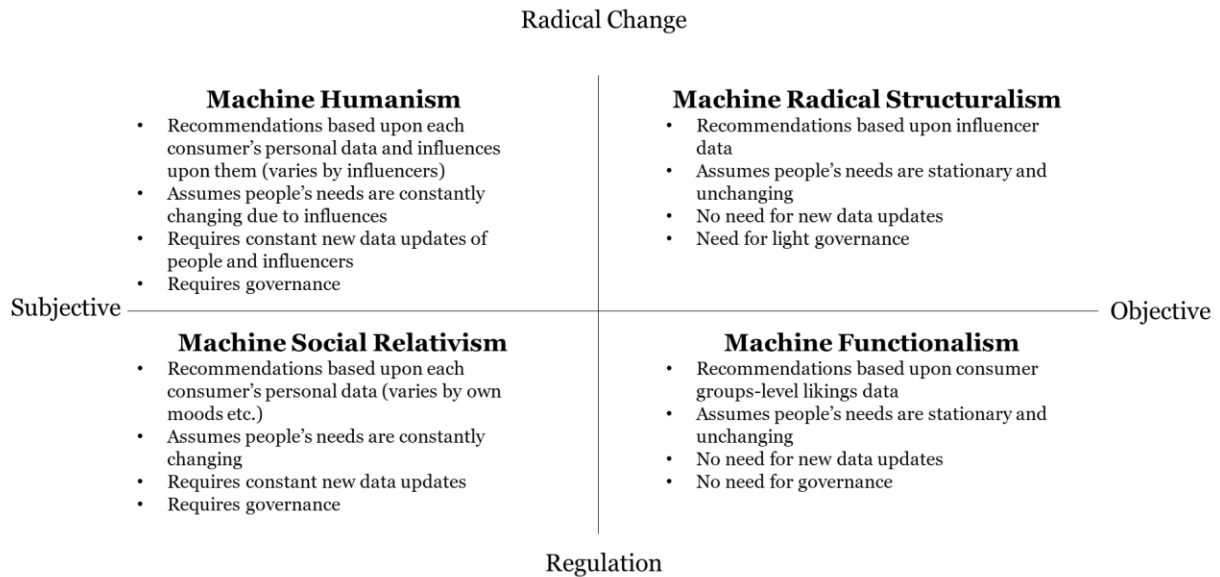


Figure 2.7 Conceptual Framework Application: Recommendation System

Example 2: A Political Campaign Management System

Another example will be of a political campaign influencing system that targets citizens with intelligent content. The basic assumptions about people's and system's underlying reality (paradigms) will greatly influence the scoping, ideation, conceptualization, and designing of the system. The assumptions are shown below in Figure 2.8. If we assume the product will be machine functionalism paradigm, the product will classify people into categories based upon rational criteria for decision-making – such as demographics, education, and other similar campaign factors. If the assumption is machine radical structuralism, the product will classify people based upon broader social, political, and economic forces and interests (conflicts and dominant forces). If machine social relativism, the product will monitor behaviors and moods, and will classify people accordingly. If machine humanism is the underlying assumption, then political, social, and economic forces will be used along with behavioral information. In this case the product can create fake stories, images,

or fake videos to influence behaviors. This shows how the conceptualization of the product will be greatly influenced by the underlying paradigms.

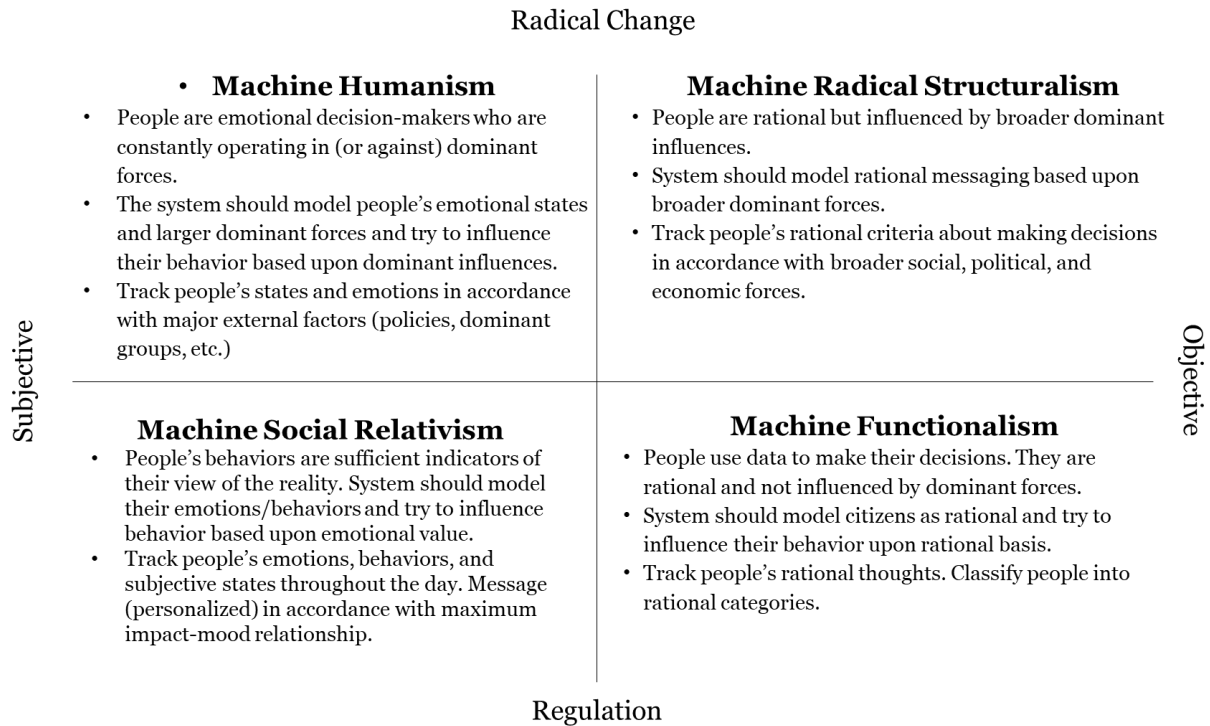


Figure 2.8 Conceptual Framework Application: Political Campaign System

As shown in the above examples, ML NPD is no longer confined to a traditional or conventional first step of Business Understanding – as it requires the designer or conceptualizer to first question his or her own sense of scientific/social reality and also question the scientific/social reality of the artifact that he/she is conceptualizing. While this inquiry does not necessarily change the configuration or sequence of the subsequent steps, it greatly changes the activities undertaken in them.

2.7 Building the Basic Conceptual Framework

By combining the paradigms and the ML solution development framework from Section 2.2, I was able to construct a Basic Conceptual framework (Figure 2.9). This model assumes that developing an ML product will be done in two steps. In

the first step, the cognitive reality of the product or offering will be explored. The cognitive reality comes from understanding the underlying assumptions about the social and scientific realities of the product, the users, and the designer. The diagram shows that process in the box on the left (Figure 2.9). The designer makes a determination about his or her understanding of the scientific and social reality of the world and also about the underlying social and scientific reality of the product being modeled or conceptualized. Answering these questions may require exploring the epistemological, ontological, and axiological assumptions – and they are depicted at the bottom of box on the left.

Once the social and scientific reality is deciphered, the product design team can now move into the more traditional intelligent system design process. This process is shown outside the box on the right in the figure. The problem element and the idea generation for a new product can be viewed as a result of the reality exploration (as shown in examples of recommendation system and political campaign management). Other ideas about the scope, performance standards and expectations, governance standards, functions, features, and lifecycle plans of the product can also be influenced or brainstormed using the paradigms. Idea generation leads to business understanding, which triggers the process of product construction as shown in the figure. This process is centered upon data – and that is why data is shown in the middle of the figure 2.9. Note that the *business understanding* step will also require reflecting back on the reality assumptions as different assumptions of reality will lead to different business models. Once the business part is clarified, in the conceptual framework, a product construction part is triggered which is composed of data understanding, data preparation, modeling, evaluation, and deployment (these steps have been explained in Section 2.3). While the *paradigms* part appeared to have rigor, the business understanding step seems to lack important details. For example, how does one understand the business requirements? How to generate or discover a new product idea? From an NPD perspective, the “Business Understanding” step

seems highly underdeveloped in its current form as it assumes that a business has already discovered, defined, designed, and conceptualized a product. It prematurely jumps into the product construction process. From the NPD perspective, this is an evident gap in the conceptual framework. Hence, in order to have a working NPD framework, I would need a far more comprehensive and detailed process. To build that part, I will field test the conceptual framework via action research to extract and enrich the model so it can achieve the objectives of the research.

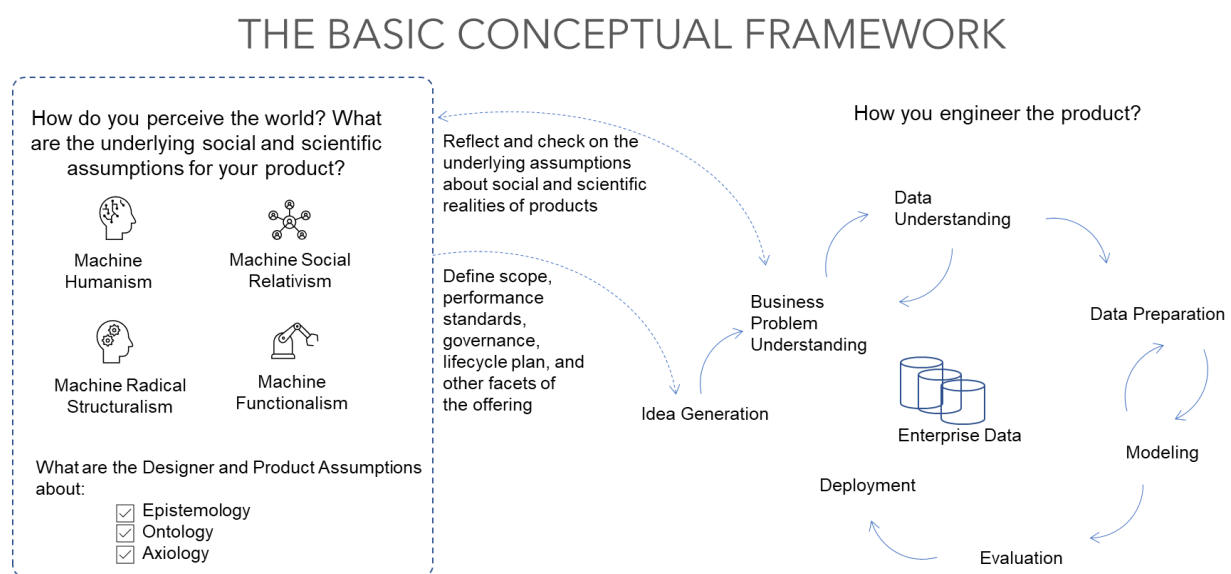


Figure 2.9 Basic Conceptual Framework

2.8 The Difference with Traditional NPD

Even though the basic conceptual framework has not been enriched from the field research, the Basic Conceptual Framework lends itself to be analyzed from an NPD perspective. As recently as 2014 Cooper, the creator of the Stage-Gate[®] process (a well-recognized discovery-to-launch process NPD model that is based upon moving projects through stages), acknowledged that for several decades his Stage-Gate model stayed as the industry standard but that was now changing (Cooper, 2014). He recognized that the genesis of the model was based upon in-

depth study of intrapreneurship in major corporations. While the model, Cooper argued, has had a positive impact in general, many firms pointed out its limitations and were adapting to different models. Referring to criticisms by Becker (Becker, 2006) and Lenfle and Loch (Lenfle and Loch, 2010), Cooper described the limitations as:

- The world has changed (faster, more competitive, and global).
- Too linear model, too rigid, and too planned (inflexible) to handle dynamic and innovative projects
- Not adaptive; lacks experimentation
- Gates are too structured; too bureaucratic and controlling
- Not context-based (one size fits all)

The above limitations, Cooper argued, have made the traditional Stage-Gate ® model open for new adaptations. Next generation process for idea-to-launch system requires agility, adaptability, and acceleration (speed), he claimed.

Cooper (2014) suggested that spiral model (build, test, feedback, revise) and risk contingency models are such adaptations of Stage-Gate. In certain cases, Stage-Gate can be shrunk, and in other cases tailored to project requirements.

Furthermore, the model can be based upon financial criteria and portfolio management concepts.

Traditional NPD models – whether linear or other adaptations – do not assume products will exhibit intelligence and will operate in an evolutionary interactive social structure where products will interact and learn from humans and vice versa. Thus, the starting point of NPD methodologies and frameworks is often the customer need or want – which is translated into a market opportunity and based upon which the subsequent steps of NPD are invoked. Since products are not expected to have a cognitive dimension, conventional NPD models are not

concerned about the epistemological, ontological, or axiological existence of a product.

The conventional NPD process, therefore, appears as limited in the sense that it is designed for corporate consumption. Cooper acknowledged that his model developed in large companies and his arguments for various adaptations are also derived from large company examples (for example Corning, GE, Honeywell, LEGO) (Cooper and Sommer, 2018; Cooper, 2014). This limitation is perhaps necessary to make NPD a pragmatic corporate process – but because of that goal, the process typically does not address more recent factors such as self-learning and self-evolution of products, self-governance, and social presence of products. The traditional NPD models view products as responses to market opportunities and often differentiate models from each other based on how new opportunities are identified (Shepherd and Ahmed, 2000; Grönlund et al., 2010). Traditionally, products are conceptualized as outside-in and via human led intentional process of discovery, whereas ML can autonomously discover products from inside-out (data-to-discovery) process via intelligent systems.

When viewing NPD as a corporate process, more theoretical aspects, such as the metaphysical essence of a product or production process, are rarely analyzed or challenged (Houkes and Vermaas, 2009; Koskela and Kagioglou, 2005). Aspects such as that products are a manifestation of market needs and human creativity are typically explored via surveys of needs and preferences or the insights and practices of the entrepreneurs. They are also viewed as the trigger or the initial steps for an NPD process. The underlying assumptions of product designers' or the underlying assumptions of products' own social and scientific realities are not considered. In this research, I argue that intelligent product discovery will require a different approach.

The prevalence of intelligent products, as explored in this research, gives rise to a new approach. In this approach, the epistemological, ontological, axiological, and

metaphysical attributes and essence become the first aspects of the idea-to-launch process. In this approach the cognitive operating model of the designer and the underlying assumptions about the reality perception of the product are considered as the drivers through which very different products can be discovered, designed, and developed to solve the same problem or to automate the same process. Just as four human beings can approach a problem from four different perspectives – based upon their experiences – intelligent machines can also address the same problem from different angles. Such angles impact the various elements of NPD – including the idea and design elements – and the scope and function of the product could turn out be remarkably different across various manifestations. In this research, such differences in product orientations are determined from the assumptions about the underlying social and scientific realities – and such a discovery process of “orientations” is not like marketing-based discovery. It models human and machine realities.

To model the realities – of both human designers and machines – we question their assumptions about social and scientific realities. Since designers are designing systems with cognitive capabilities, their own perception or assumptions of science/social reality will lead to very different products. Since the product itself has cognitive capabilities (ability to learn), the product’s underlying perception of reality social/scientific will have significantly different outcomes in how the product learns, functions, performs, and is governed.

At the outset, the model presented in this research introduces a preemptive step to determine the scientific/social reality and perception of the designer and the product prior to invoking the traditional or conventional NPD process (See Figure 2.9). Subsequent process steps of NPD will also depend upon identifying the relevant data and the same data (raw material) can produce very different products. It is expected that operating models of companies which result from

offerings organized around intelligent products will derive their competitive advantage based upon the various manifestations of social and scientific realities.

With intelligent products from ML, the assumptions of the reality of the designer and of the artifact itself become the focus of interest and analysis.

Just as the scientific process is based upon data-to-experiment-to-result configuration, ML process is also similar. ML algorithm uses data to learn by exploring all the possible relationships between data elements and finding the most optimized one. This is akin to scientific discovery where relationships are established between different variables. Hence the ML product is not a deterministic solution to a problem but instead an approach to identify a preferred solution among many solutions. It explores possible relationships between variables and identifies the optimized one. Based upon the learning, the product therefore is nebulous and exhibits flexible structure with evolving features. Examples of such products include discovering new molecules using ML in pharmaceutical industry or AlphaGo ML algorithm that beat Korean and Chinese grandmasters of the ancient game Go (Makridakis, 2017; Wirth, 2018; Siau, 2018; Jarrahi, 2018).

When covering information systems, Cooper (2014) defines them as products of reductionism where software can be segmented into millions of lines of code. He sees other products – such as food products, machines, medicines, etc. – as different since they require more complex engineering. First, his categorization of information systems as “million lines of codes” is inaccurate for ML systems (Cooper, 2014, p.25). ML systems are not like programmed software that form information systems. ML are computer programs that learn “without being explicitly programmed” (Moser, 1990, p.10) and the field of ML aims to answer the question “How can we build computer systems that automatically improve with experience, and what are the fundamental laws that govern all learning processes?” (Mitchell, 2006, p.1). Secondly, in both of Cooper’s categorization of

products – information systems and other products – he does not entertain the possibility of a product whose form is derived from underlying data and distributions, a product that develops from data and not lines of code, that is based upon learning, that adapts and accumulates experiences, that develops itself, and that can grow and evolve throughout its lifecycle. ML cannot be segmented into infinitesimal segments – just as human brain cannot be segmented into neurons to study how thoughts materialize. ML product can also not be approached like a food product or a machine that requires mixing of various ingredients or parts and is a result of deterministic human engineering based upon reductionism. An ML product takes in datasets – which means different features (variables) and vectors (records or values of features) in datasets – and then uses them to train one or more algorithms. Such products are mathematical constructs that learn to perform work – and they are not lines of codes.

For that, these algorithms can themselves produce or learn to produce ideas and even develop products that Cooper's model is attempting to pioneer. In other words, an ML software can write software, design machines, create food products, and invent new medicines. Hence the main problem of an NPD framework for ML is not whether the framework would be linear, staged, spiral, agile, or complex. The main problem is whether we would need to rethink the fundamental concept of an intelligent product as distinct and different than all other products ever developed by humankind and therefore reorient our thinking and best practices about NPD of intelligent products.

NPD requires planning (Salomo et al., 2007). Best practices are important for successful NPD (Kahn et al., 2006). There are various types of NPD models (Shepherd and Ahmed, 2000; Grönlund et al., 2010). It requires managerial sensemaking (Christiansen and Varnes, 2009). But in an era where machines are

developing capabilities to learn, there is a need to expand the traditional NPD approaches.

2.9 Next Steps

The conceptual framework proposed here is a theoretical framework which has been developed based upon theoretical assumptions. While its core aspects of social paradigms remain theoretically intact – based upon the continuation of the Burrell and Morgan (1979) and Hirschheim and Klein (1989) frameworks, it gave me a baseline methodology that can now be field tested. It is used as the baseline and tested in the field to further develop it from the practitioner's perspective.

In this research, without tying the AIAI ML product development to any one model (Cooper's or otherwise), a field study was conducted to explore the conceptual framework identified in this chapter. The methodology of that research is explained in Chapter 3, and the study and analysis are captured in Chapter 4.

The research methodology proposed in the next chapter took into account the conceptual framework developed in this chapter. While the literature review continued throughout the research, at the end of the substantial literature review performed in this chapter, I was ready to put the practitioner hat on and jump back into the practitioner-researcher role of an action researcher.

Chapter 3 Methodology

3.1 Introduction and Review

To close the gap created by not having a framework, I formulated the following specific problem that I explored during the study:

The fundamental question addressed by the study was: **what is a new product development framework for designing and developing machine learning products?**

This chapter begins with introducing the research approach and giving a rationale for selecting the approach. The next section will cover the process for research sample selection. It will be followed by identifying the information needs for the research. The next section will cover research design, followed by data collection methods. The next section will cover data analysis and synthesis. The following section will discuss ethical considerations which will be followed by issues of trustworthiness. Finally, after discussing limitations and delimitations, a chapter summary is provided.

3.2 Research Approach

As covered in Chapter 2's literature review (synthesis, and conceptual framework sections), this research asserts that the shortcomings of data mining methodologies cannot be overcome by engineering-methodology proliferation. While literature points to significant shortfalls of existing data mining methodology, the suggested remedial steps were all engineering centric. Even though many methodologies were offered by some of the leading companies (SAS, IBM, Microsoft), they lacked differentiation as they all seemed to mimic CRISP-DM (Martinez-Plumed et al., 2019) and they all seemed to approach methodology from an engineering side. The perspective developed in this research was that engineering methodologies without their counterpart scientific/social

paradigmatic frameworks are inadequate to address broader and growing information systems development (ISD) needs. As conventional ISD has shown, paradigms give structure to methodologies, and they improve the methodologies. Making the distinction between engineering-centric approach and social/scientific paradigm-centric approach, it was argued that information systems development paradigms take shape in social/scientific contexts. Software design and development, it was suggested, is a social process. In accordance with established traditions, paradigms materialize from analyzing the philosophical assumptions related to social structures (Burrell and Morgan, 1979; Hirschheim and Klein, 1989).

For the research question, the research approach required exploring human behavior and organizational decision-making – specifically how new ML products are designed and developed from idea-to-implementation. Many choices are relevant to the experiences of the people and contextual situations may also influence those choices. A quantitative approach could have been applied to take a broad population of engineers and quantitatively determine their usage patterns of various methodologies and the business performance (for example shareholder value created or stock value improvement). Such a research approach would have made sense from an engineering perspective or a financial perspective where the research could have explored the relationship between ML development methodology choice and financial performance. But the literature review has suggested that methodologies at the basic level tend to be similar (follow the CRISP-DM pattern) and that they are applied in narrow project contexts. This means that simply identifying methodologies from a quantitative perspective would have been an inconclusive exercise since all engineers would essentially use the same methodology. Secondly, the business performance would not have been comparable because the automation projects across different companies would have had different scopes and different performance expectations. Since literature review points out approaching methodologies from

a paradigmatic (social reality) perspective, qualitative approach appears to be a better match.

The second and equally important reasons for using qualitative approach is that the research is related to a business problem and a process of AIAI. AIAI offers deep and rich information about artificial intelligence. The institute operates in the field and therefore can address many issues that firms operating in the industry would not know about. This information resides with the experts who are knowledgeable about the industry wide problems. As the research approaches inquiry from the behavioral and organizational angle, the observing of social dynamics, searching for meaning, sensemaking, and understanding experience-in-play would become extremely important. Based upon the reasons discussed above, a choice was made to approach the research as qualitative research.

Within the qualitative research approach, many research methods were available. For example, this research could have taken a case study approach, grounded theory, phenomenology, ethnography, narrative, or other qualitative traditions. However, as previously mentioned and stated in the goals of the study, the business problem corresponding to this study was of launching a new product. Launching a new product requires a series of actions, a team, overcoming challenges, coordination and communication, organizational energy, operational dynamics, and tremendous momentum. All of those factors manifest in the work setting and they constitute as powerful hunting grounds to identity rich sources of data. In those scenarios, research can progress through actions and actions can benefit from research. AIAI's product launch process took place in an active work setting and provided an extraordinary opportunity to conduct this research. Therefore, after a review, a methodology known as action research (AR) was selected.

In a recent analysis Verna J. Willis argued that after analyzing dozens of definitions of action research online, she selected two excerpts from different

sources that she felt offer most straight-to-the-point descriptions of action research and its researchers (Willis, 2010, p.167):

Action research is inquiry or research in the context of focused efforts to improve the quality of an organization and its performance. It typically is designed and conducted by practitioners who analyze the data to improve their own practice (North Central Regional Educational Laboratory, www.ncrel.org).

Action researchers “see the development of theory or understanding as a by-product of the improvement of real situations, rather than application as a by-product of advances in ‘pure’ theory.” (Carr and Kemmis, 1986, p. 28, cited also in Wikiversity Action Learning article). This is a means to generate ideas (theory) that are relevant locally – to the people who are involved in the research, and to the environment in which it has taken place. (Wikiversity, en.wikiversity.org)

The above definitions applied to our problem. The research was being conducted by practitioners and theory was being developed as a by-product of real situations. Furthermore, the goal was to generate theory for local relevance and to the people who are involved in the research and to the environment local to the research.

AR is a well-established approach with decades of history of application. In accordance with the Action Research approach, knowledge creation is possible and helpful via a pragmatic approach to learning where basic handling and tackling business issues produce usable knowledge (Mode 2) (Starkey and Madan, 2001). Unlike the traditions of positivist research, AR offers actionability and reflexivity and enables research from the three voices of 1st, 2nd, and 3rd person (Coghlan and Brannick, 2014). Within the broad AR discipline, IAR (Insider Action Research) involves conducting the research in an entity in which the researcher is a part of as an employee or team member (Roth et al., 2007). This gives most clear picture of the organization, its politics, concerns, people, and issues. However, it also creates

the risk of role duality whereby it becomes hard for the practitioner-researcher to separate himself (herself) from the experiment in which he (she) is invariably part of (Trondsen and Sandaunet, 2009). The desire to achieve that neutrality produces a sense of self-awareness necessary to gain the benefits from the application of AR in business (Raelin and Coghlan, 2006). That is why reflexivity is an important part of AR.

The benefits of applying AR for this study are abundant. AR affirms the inherent human potential to contribute to knowledge based on people's personal lives and experiences. It is valuable because pragmatic wisdom, practical reasoning, and tacit knowledge are critical to form a multidimensional view of the business problems and solutions (Carr and Kemmis, 1986). Practical reasoning and socially constructed insights can be obtained. Significant literature is available about AR and presenting the AR introduction here is not to educate the reader on AR but to clarify the reasons behind selecting the methodology.

AR is not a single methodology and can be considered as a family of practices, and hence there are many ways to structure a study in AR (Reason and Bradbury, 2008; Coughlan and Coghlan, 2002; Chandler and Torbert, 2003). I chose action research not only because the participants were actively engaged in the project that is the focus of this research but also because the research outcomes were expected to help the entire firm. The survival of AIAI would ensure that participants would benefit from it. AR benefits from the collaborative action and reflection where inquiry manifests in experience and social history (Reason and Bradbury, 2008). Questions raised and problems addressed are relevant and significant for the participants – and participants, action, and research are integrated (Chevalier and Buckles, 2013). While PAR is often approached to give a voice to the subjects of the research (Reason and Bradbury, 2008), in my case it was adopted to capture the experience and social history – while hoping that emancipation and empowerment

features of PAR would be realized when ethics and social justice concerns can be integrated with ML software development.

As indicated in the literature review, past research from the quantitative and operational side (engineering of the ML methodologies) continues to be a source of disappointment, it was expected that qualitative research about behavioral and social issues would reveal new insights. The workplace provides the arena where actions happen and where meaning is assigned on the basis of internalized notions of norms, roles, values, traditions, behaviors, executive decisions, and crucial contextual variables. The setting of the research plays a role – for example, AIAI's product launch schedule and goals contributed to the pressure on participants to proceed with a certain urgency and research leads to action and action drives research.

This research was conducted with the viewpoint that thoughts, feelings, and action require interpretation and that I did not want to impose my view of the world on the participants. Due to the field research setting, it was expected, that the research will be able to draw meaning behind business processes and events.

3.3 Research Sample

One way to conduct this research would have been to define the population as composed of a broader set of managers from multiple companies. However, both the scope and the goals of the research are related to the business of AIAI. Due to being an internal project, the population from which a sample could be drawn included the employees and affiliated parties of AIAI. One way to extract a sample would have been to randomly select a set of participants from all employees for this research. However, the relevance, information richness, previous experience, and knowledge as related to the problem are not equally distributed across the entire workforce. People who were part of the new product launch team had far more experience and knowledge about the subject matter. Secondly, the research had a contextual element that is specific to AIAI's product

launch. While it was hoped that the insights obtained would be generalizable and transferable to other firms and the industry, the contextual constraint did not allow defining the population as composed of knowledgeable people from other companies or outside of AIAI. Therefore, in accordance with the definition of the Action Research presented in Section 3.2, the population was defined as the employees of AIAI and the following criteria was deemed important to draw a sample:

- Involved in the project and at the frontlines of action. This condition was deemed important because of the nature of the study: i.e. Action Research. People who serve in passive roles in AIAI or who are not directly involved in the product launch would not have constituted as a good sample. This does not mean their insights are not important. It simply means they are not part of the problem domain that is being addressed in this research.
- Knowledgeable about ISD methodologies: Since one of the objectives is to determine the issues with the data mining methodologies as they are applied in ML, knowledge about the ML and ISD methodologies was deemed important.
- Knowledgeable about the goals and of AIAI. The contextual knowledge about the strategies, processes, values, and business issues of AIAI was considered important.

The three conditions of being at the frontlines of action in projects, relevant experience, and knowledgeable about contextual information of AIAI were all viewed as relatively objective measures. In accordance with the Action Research definition presented in the previous section, it was understood that while the sample size is small (two), the depth and relevance of participant experience were of utmost importance for the research.

Based upon that, two participants David and Randy were selected as the sample.

Sample	Project Involvement in Frontlines of Action	Experience and Knowledge about AI/ML	Contextual Knowledge of AIAI
David	Yes	Technology expertise, general	Employed by AIAI since the inception of AIAI
Randy	Yes	Technology expertise, AI + general	Employed by AIAI since the inception of AIAI

Table 3-1 Sample (People)

It was expected that this process of sample selection will yield insights that are not available outside of the participants. It was understood that two action research participants will participate in research via providing data through dynamic dialogue, assisting in data gathering, participating in learning sets, analyzing and collecting documents, implementing recommended actions and collecting feedback, and collecting process knowledge from the field and reporting it back in the form of data.

The sample selection process was in accordance with established practices of purposeful sampling used to identify information-rich participants in action research. These participants exist at the frontline of action and are involved in action – and hence meet the requirements of involved researchers “as natives and actors, immersed in local situations generating contextually embedded knowledge which emerges from experience” and for action research being “research in action, rather than research about action” (Coghlan and Brannick, 2014, p.6). Such in-depth information rich insights are not available from random sampling. The researcher was cognizant that making such a choice can frame who and what are deemed important as data, however the alternative of random sampling could not have worked in the context of action research setting of this research.

Two potential participants were approached in accordance with the requirements of the University of Liverpool's ethics approval. After explaining the research goals, their roles, and answering their questions, both participants agreed to participate in this research. Their acceptance was voluntary, and no pressure of any kind was placed upon them.

It was expected that through their direct involvement in the project, they will provide rich sources of data to meet the objectives of the research.

In addition to the participants, a sample of documents were selected for this research. One sample included the entire population of key operational documents of AIAI including RFI (Request for Information), RFP (Request for Proposal), and documents produced for the product launch. In that case, the only condition that was deemed necessary was that the documents were related to the product launch and were related to operations. The term operations in this context excluded legal or intellectual property related documents.

Another set of documents were included in the research. This was the set of processes and methodologies produced by AIAI. In that category, the sample included only those documents that were authored and produced by me and are publicly available.

This research was conducted in a work setting. The entire communications and meetings related to this research happened virtually. All meetings were conducted in compliance with the Covid19 related instructions issued by University of Liverpool. The online participation of various people did not impact the study – however, it did limit the observation of body language.

3.4 Overview of Information Needs

The information needs of this research were recognized based upon the underlying assumption that reality is also socially constructed. The primacy of subject matter was recognized vs. the primacy of methods. The research dealt with several

moving, fast changing, complex, interconnected, and difficult to measure variables. I was immersed in the research and with the help of the participants was attempting to search for meaning by contextualizing, exploring, and understanding the perspectives of the participants.

To accomplish those goals, three types of information were deemed important for this research: Contextual information, perceptual information, and theoretical information (Table 3.2). Within these classes of information, the most relevant and important to search for meaning was perceptual information and that will be explained in detail after briefly introducing the remaining two.

Contextual information: Contextual information includes organizational goals, strategies, processes, schedules, project documents, and values. This was used as reference information for the research.

Theoretical information: While significant literature review effort was made in the initial stages of the research, literature research continued to be helpful and emergent during this research.

Perceptual information: The most relevant information related to this research is the perceptual information. Perceptual information provides an inside view of meaning related to processes, events, behaviors, and actions. This research was about exploring the meaning behind ML NPD as assigned by practicing professionals and doing so within the context of their actions and daily work.

Exploration	Perceptual	Contextual	Theoretical
The NPD process for Machine Learning	Dynamic discussions from learning sets (participant sessions) about the product launch process and how various issues relate to the data mining methodology. Active discussions, action planning, and critical reflection.	The information about the strategies, goals, processes, and existing methods of AIAI for new product launch.	Two types of theoretical information: 1) Theory about the research (as discussed in Chapter 2) and the Conceptual Framework 2) Theory and research about the product being developed by AIAI (for example, Audit).

Table 3-2 Information Sought

3.5 Research Design

This research was designed as an action research. The literature review (Chapter 2) provided a scaffolding to help guide through the research design. The experimental and quantitative research from the past and the engineering approaches to data mining methodologies were unable to alleviate the limitations of the ML methodologies. The methodology proliferation only contributed to the problems and even methodologies that were proposed by firms such as IBM, Microsoft, and SAS came short of addressing the problems (Section 3, Chapter 2). This research approached the problem from a social, organizational, and human perspective.

The theoretical and conceptual framework (Figure 2.9) developed in the previous chapter was used to guide the research data analysis and interpretation phases of the research. The review of the literature was maintained throughout the research and the related research was updated regularly. Literature review can provide the scaffolding, but it does not generate new information or provide the supporting evidence. To develop that, the research design was composed of an overall approach, design, and methods. They are described below.

3.6 Action cycles

Action Research is conducted in action cycles where action cycles represent actions and their evaluation and reflection. Our research resulted in six action cycles and several mini-micro cycles within each cycle.

Action Cycles	Relationship with Conceptual framework	Action Steps Explanation
Cycle 1	Idea Generation, Business Understanding	Kickoff – determine the product area
Cycle 2	Business Understanding	Define a specific problem
Cycle 3	Business Understanding	Discuss and evaluate the feasibility of the product – internal capabilities and market potential
Cycle 4	Business Understanding, Understanding data	When clarified from the previous step, design the product
Cycle 5	Preprocess data, model	Develop the product
Cycle 6	Evaluation, Deployment	Evaluate the product from a performance and adoption criteria. Deploy the product.

Table 3-3 Action Cycles

Each of the above actions was considered a cycle (Table 3.3) and they represented the natural flow of a product launch and where the goal of the cycle was to obtain insights about that specific phase of the inquiry. Each cycle reflected one or more steps identified in the Conceptual framework (Figure 3.3. middle column). Those steps appeared on the right side of Figure 2.9 which depicted the Conceptual framework. For instance, Cycles 1, 2, 3 and 4 elaborated, expanded, and explored what was stated as a single step in the Conceptual framework and was called Business Understanding, cycle 4 also aligned with Understanding Data, cycle 5 with Data Preparation and Modeling, and Cycle 6 with Evaluation and Deployment. Each cycle was composed of four steps of Plan, Act, Observe, and Reflect – and several secondary cycles transpired within each cycle. It was understood that the cycles may not go as planned and additional cycles could be necessary. Specially, it was recognized that cycles-within-cycles will require flexibility as action, evaluation, and reflection will guide their evolution.

In accordance with Coghlan and Brannick's (2014) depiction of cycles as a clock analogy of three hands moving simultaneously, each of the six cycles was composed of many mini and micro cycles that ran simultaneously as work sessions, meetings, activities, and mini feedback sessions (Coghlan and Brannick, 2014).

3.7 The assumptions for choosing the method

The underlying assumptions for the method and data choices were as follows:

- I recognized that I am not operating in a singular reality. There are multiple, constructed, and holistic realities and I must try my best to navigate through those.
- Both knowing-in-action and reflection-in-action would contribute to data (Schon, 1983).
- The process of data acquisition will be democratic.
- Knower and Known interact and are inseparable. The researcher (I) and participants themselves are also their experiences and knowledge.
- Hypothesis are context bound and independent and can be verbally expressed as idiographic statements.
- The research takes place in a dynamic and active world where variables, actors, social and political considerations, and all other factors are simultaneously and mutually shaping their own and the holistic realities and where it is not possible to view cause and effect in their traditional logical context.
- Inquiry is value bound, implying that human judgment, driven based upon some values, also shapes part of the reality observed and analyzed.

My personal stance on research in social sciences is that social constructionism is a critical ingredient that shapes truth and hence positivism is insufficient to help reach to the truth. As an AI professional, I have observed on numerous occasions where explanatory variables are hidden in cultural, linguistic, social, and

semantic structures and not necessarily normative quantitative scientific variables (for example, GDP, income, and others). For that reason, I fully believed in the research design and approach and did not experience internal conflict related to the methodology.

3.8 The investigation process

As participants acted in their natural environment (office and workspace), they took actions, made decisions, and communicated to achieve certain goals. All of those activities – including the actions, meeting, expressions, symbols, communication, and others – became rich grounds for the systematic collection of data.

In essence, the inquiry was patterned to evaluate:

- What concepts and values do my participants use to classify their experiences? How do they prioritize and understand the AI frameworks?
- How do my participants define these concepts?
- What “Theory in-use” do my participants use to explain their experience and what is the difference between espoused vs. theory in use? (Argyris and Schön, 1974)

In each of the cycles, the 4-step analysis of plan, action, observe and reflect elucidated information and data about the three questions.

The notes from the meetings were captured and analyzed in light of the problem being investigated. As explained below in the methods, the transcribed notes were classified and organized to extract meaning. In addition to the meeting notes, documents (specifically related to the research conducted to discover, select, and design the new product) were analyzed for extracting meaning.

3.9 Data Collection Methods

An AR project provides rich and fertile grounds for data; however, methods determine what type of data is captured for analysis. I had many options for

methods. For example, I could have conducted formal interviews, captured information from observations, conducted surveys, applied ethnographic or narrative techniques, used the internet to crowdsource responses, or conducted case study-based research. In seeking a method, the following issues were considered:

- 1) I wanted to preserve the natural communication and casual interaction to gather rich information. A more formal method – such as a structured interview, closed ended questions, or surveys – would not have supported that goal.
- 2) I recognized that actions are a source of information and data, and that extracting data from the cycles of action would be instrumental to meet the needs of this research.
- 3) Finally, since the research involved obtaining deeper insights and meanings, and because I was embedded in the research and as such brought my own biases and values into the mix, I recognized that reflexivity would need to be an essential part of the research (Coghlan, 2001; Coghlan and Brannick, 2014).

Coghlan and Brannick (2014) have clarified that AR can include all types of data gathering methods – including interviews, surveys, and recorded data in organizational documents and journals (Coghlan and Brannick, 2014). As such, four methods (learning sets, observations, document inspection, and reflexivity journal) were selected for data collection. Reflective Inquiry via *Learning Sets*, a term borrowed from *Action Learning*, was chosen as the primary method to access data from participants. As Marquardt and Yeo describe (Marquardt and Yeo, 2012, p.45):

Reflective inquiry is a participatory action-learning intervention that uses both intrapersonal and interpersonal dialogue to gain an insightful perspective of the problem issues. Such a dialogue is achieved through an

interplay of assumptions, direct and indirect experiences, and the use of data by participants. An iterative feedback cycle, which encourages participants to sharpen their conceptualization and experimentation, is the core component of dynamic dialogue in group settings (Schon, 1983). This process allows groups to test and verify assumptions and mental frames through trial and error. The idea is to bring individual and tacit assumptions as well as perspectives to the fore for group critique in the form of meaningful dialogue before actual implementation.

Such a methodology is considered useful for modern day problems that are based upon complexity, opaqueness, interconnected, dynamics, and polytely (multiple goals) (Marquardt and Yeo, 2012). Participants tap into their collective experiences and explore breakthrough problem solving. Insights obtained are deployed as actions and inquiry of action provides theory. A typical reflective inquiry is based upon four stages of generative, integrative, resolution, and formulation. In generative stage questions are asked with hypothesized outcomes and scenarios and in-depth conversation (dialogue) take place that “induces a spontaneous “rupture” (enlightenment) in the way ideas are constructed and connected, leading to breakthrough thinking” (Marquardt and Yeo, 2012, p.12).

In integration stage participants collaboratively make sense of data generated in the previous stage. In resolution step, differences of opinions are discussed. In formulation stage action plans and future roadmaps are generated. Consequently, the data gathering was as follows:

Dynamic Dialogue: Dynamic dialogue, as clarified and defined by Schon (1983), was conducted in meetings designed for learning sets and as part of the action cycle and during the cycles-within-cycles (mini and micro cycles). Each of such interactions, in the form of virtual meetings was designed to invoke open and democratic discussion and dialogue about various actions. Specifically, I engaged in the following interactions with the participants: 1) As part of the

dynamic dialogue, I engaged the participants to reflect upon the knowledge gained through their actions (a requirement of Action Research) ; 2) To further clarify and generate data through actions, I asked them more specific feedback questions by presenting them a certain data – for example, I presented them with the data on ML applications and asked them to comment on it; 3) I provided them documents and also asked them to provide documents about action cycles and requested them to assess needed or implemented actions as they related to the project; 4) I asked them questions how during a cycle they performed a process or approached and solved a problem (actions) and how those actions affected outcomes; 5) I worked with them to extract the process details from applied actions and to understand what worked and what did not work and what can be improved; 6) I requested them to provide any special insights – either organization or process related during the cycles; 7) I obtained their feedback about the approach itself; and 8) I verified the findings with them. The sessions were transcribed in various physical and digital medium, and an audit trail of main topics of conversations was maintained.

Such dynamic feedback and reflexivity were not coded since coding is viewed as “a reductionist and mechanistic process that detaches the data from a broader understanding, downplaying the context” (Gjerde and Alvesson, 2020, p.130). Gjerde and Alvesson (2020) cites others and reminds that coding is not appropriate for all types of qualitative data (Bansal and Corley, 2011, p.236) or for all qualitative traditions (Gjerde and Alvesson, 2020). The reductionism can shift the context making it difficult for broad ideas to materialize (Potter and Wetherell, 1987). Since we did not have a context as what we were conducted lacked an example or an external or internal model or framework, we needed to keep our minds open to the possibilities. Coding would have done the opposite. It would have concentrated upon familiar themes that developed semantically or thematically and not new and unknown concepts. This research uses both hermeneutical principles and active participation of learning set members –

moving between parts and whole, reflecting, and relying upon some preunderstanding that I brought to the research (Alvesson and Sköldberg, 2009; Alvesson et al., 2008). My goal was to discover what was less obvious and to search for what was hidden beneath the layers of practice. Coding process – which aims to standardize and de-contextualize the data – could not have identified the complex themes. However, the discussions were recorded, analyzed, words and their contradictions were studied, meanings were developed, and practical insights were obtained via data analysis. In accordance with the definition of Action Research, such activities were conducted in the contextual setting of AIAI (Alvesson and Sköldberg, 2009; Willis, 2010).

Document Inspection: In addition to observation, document inspection was used as the third source of data. In document inspection, I, along with participants, analyzed the documents that captured research for the project and that contained information about the project status. Additionally, some documents that were analyzed contained AIAI processes and methods. Significant data was acquired from using these documents. These documents represented process insights from actual actions.

Reflexivity journal: I used reflexivity journal to journal my own experiences, record my biases, capture my feelings, identify symbols, and reflect on the research activity and related concepts on an ongoing basis.

Observation: In addition to dynamic dialogues, I also planned to use observation, albeit to a much lesser extent. In observation, researcher observes the participants in terms of their reactions and other factors such as linguistic choice, emotions, or behaviors. Since due to the coronavirus related University of Liverpool instructions, the research was undertaken virtually (phone, online), I did not have the ability to observe factors such as body language. I was able to acquire limited data on this front. With open-ended engagement and immersion, I attempted to experience reality as research participants do.

Post-research customer feedback: To understand the practical aspects, utility, and application potential of the framework beyond AIAI's use, customer feedback was obtained and reported in this research as a post-research feedback mechanism.

The combination of field notes from meeting notes (dynamic dialogues), reflexivity journal entries, inspected documents, and relatively few observations, became the primary instruments of investigation. The data was maintained in secure electronic (digital) form.

All four are well-established methods for qualitative research. The use of four methods, when combined with the literature review, maximizes the combined efficacy to extract the relevant data. Methodological coherence was established by ensuring that methods are not incompatible with each other.

In keeping with the traditions of action research the reflexivity was to capture content reflection (consider issues), process reflection (think about strategies and procedures), and premise reflection (critique underlying assumptions) (Mezirow, 1991).

3.10 Data Analysis and Synthesis

The investigation moved in natural progression of a product launch – a total of six cycles of action research. Within each cycle there were several mini and micro cycles. In each cycle, the product launch related data was captured from the four types of above-described methods and then feedback from customers. The overarching concept was that in each action and its related data, I will be able to discover important insights about NPD in ML (as shown in Figure 2.9).

As I engaged in critical reflexive analysis and dynamic dialogue with the participants, they provided detailed answers and that became a rich source of data. That data received was first understood and analyzed. Significant time and effort were spent reviewing the data and extracting concepts, meanings, and

ideas. These classification patterns were based upon first my understanding of the problem domain and then were confirmed by the participants. The ideas constituted as the context relevant accumulation of discussed concepts.

Evaluation of Documents and Participant Action Research

The crux of this research is based upon participative AR where cycles and cycles-within-cycles explored actions and reflected upon the implication for the unfolding NPD process. All cycles were based upon high levels of participation and collaborative critical reflection. Data in the form of participant inputs and inspection of documents provided the data for sensemaking related to the research question, and collaborative critical reflection guided action for the next cycle or cycle-within-cycle. General inductive thinking captured relevant themes, and through cycles of action and reflection new knowledge developed.

AR attempts to address and solve pragmatic, pertinent, context bound, complex, and real-life problems through a unique participative, collaborative, trust building, democratic and liberating process (Greenwood and Morten, 2007). Theory and action must not be separated, and the world of experience is included as a valid area to tap into (Greenwood et al., 1993). The cogeneration of knowledge happens through collaborative communication processes and democratic participation is maintained – which implies that all participants' contributions are respected and taken seriously. The workability and increased participant control over solving the problem provides the credibility-validity of AR knowledge. Meanings constructed in the inquiry sessions lead to the construction of new meanings and knowledge (Greenwood and Morten, 2007; Brydon-Miller et al., 2003). The judgment and experience of participants becomes the basis of meaning construction and knowledge creation as discussions and conversations in democratic participatory sessions are captured.

In this research, the inspection of documents and running facilitated sessions provided the locus of sensemaking. Problem context related documents provided

an avenue to extract meaning and facilitated sessions were composed of brainstorming and creative ideation driven by seeking comments, views, discussions, and intellection from problem context related stimulus, questions, and cues. The pooling of knowledge data helped narrow down the meaning as critical reflection helped drive the next set of actions. The word “facilitated” in facilitated sessions does not mean the sessions were designed to be nondemocratic, bureaucratic, or overly structured. In fact, it was the opposite. The word facilitated simply means that the group stayed focused on the problem at hand, on critical reflection, on evaluation, and on ensuring that action is followed after every cycle or cycle-within-cycle. AR knowledge creation is not context free (Greenwood and Morten, 2007) and therefore the inquiry stayed focused on the problem – but within the cycles the research was dynamic, active, vigorous, and action based. The focus remained on extracting data to produce usable knowledge that supported the goals of the participants. Due to the nature of the inquiry, which was based upon exploring a process, the local aspect of the inquiry did not necessarily prohibit the transferability of knowledge – however the limitations of generalization and universalization of new knowledge were also critically analyzed.

The discussions and reflections were recorded as meeting notes. In certain cases, meaningful themes emerged from data. A theme represents the specific pattern discovered from data (Joffe, 2012). It forms an identifiable configuration of meaning (Willig, 2013). Identifying the emergent ideas and themes involved the experience of the team and in line with the participative action research conducted over multiple cycles of action and reflection. The identified themes were not considered as rigid and inflexible – as one may find in positivist or conventional thematic analysis approaches. As a deliberate action, no formal coding (which is typically done in conventional thematic analysis) was undertaken. This was done to preserve the fast changing and exploration-based nature of the inquiry where each successive cycle (and cycle within cycle) was

shaping the meaning derived in previous cycles. Meaning attachments would have been counterproductive and even misleading. The approach in this research followed the sensemaking via rigorous inspection of documents and analytical exercises conducted to extract data about the underlying action and theoretical structures that are informative but remain hidden – however, rigidity of, and attachment to, the meaning was avoided. As such the term theme in this context does not mean as one would perform rigid coding based thematic analysis. Within the context of a business process, it can be viewed as discovering the spread between the espoused theory and theory-in-use (Argyris and Schön, 1974), where critical reflection enlightens and closes the gap. The participants questioned the very activities they engaged in and through that reflection, active discussions, actions, exploration, and comments developed a state of awareness (Argyris, 1976). The themes of the state of awareness that emerged provided the sensemaking within the cycles and become the guiding force for subsequent cycles.

At some point the existing knowledge of the researcher and participants is needed to give meaning to the emerging patterns. However, it is done via observing common themes across data. As such the patterns recognized did not represent my prior theoretical commitments – other than what was captured in the conceptual framework. Any a *priori* professional knowledge was simply used to identify (and name) the patterns that arose from the data and not to enforce limits on what could be discovered. I took special care to ensure that the existing knowledge and bias did not guide the investigation process but instead helped in articulating meaning (when needed) or crosschecking the reasonableness of the finding. Reflexivity helped tremendously in ensuring that state of awareness. Besides the professional experience, the epistemological structures of the conceptual framework (Chapter 2) helped give sensible meaning.

The conceptual framework developed in Chapter 2 was limited to identifying the six process steps and boundaries, but what constitutes within each of the process steps was approached without a theoretically informed data interpretation. What was revealed in the conceptual framework set the boundaries of the problem context and the paradigmatic conceptual framework pointed to the ethical dilemmas and social considerations when designing systems. I paid attention to the conceptual themes that emerged from the data and that were firmly grounded in the data – a process known as inductive approach (Boyatzis, 1998) – but even in the inductive approach, themes are constructed as they emerge from data (Braun and Clarke, 2006) – as it happens in the sense that it requires the interpretation in relevance to the questions being explored in the study. In my case, the participants helped with the sensemaking.

The study is related to discovering a business process that is explored and interpreted based upon the experience of the key people in the organization. In certain cases, extra layers of explanation or feedback was added to the core findings, and that was included only as an elaboration or explanation of the core findings. The business process of ML NPD was explored from the viewpoint of the boundary set by the conceptual framework explored in Section 5 of Chapter 2.

Based upon the above, the following method was adopted:

- 1) Data was obtained from the four sources of participants, documents (existing information or models), observations, and reflexivity journal.
- 2) Data was recorded in the form of meeting and exercise notes. Transcripts of meetings were used to identify patterns of meanings.
- 3) The participants actively participated and analyzed the emerging patterns and based upon the participant experience and knowledge contributed to the creation of new knowledge.
- 4) Successive cycles and cycle-within-cycles were launched in an evaluation, theory creation, and action configuration.

- 5) Sensemaking was approached as evolutionary and findings from previous cycles were analyzed for reasonableness as new knowledge was captured in succeeding cycles.

3.11 Ethical Considerations

Research ethics authorization was obtained from University of Liverpool. I recorded the data. I was cognizant that at times I had to summarize or paraphrase the data for reporting purposes. I had known the participants for years and am familiar with their linguistic and lexical preferences and contexts. Even then, in order to avoid presumption on my part, clarification questions were implemented. As such I took extra care in recording the data. The data was managed in files and saved on secure computers. The participants did not express any concerns about their safety, security, or discomfort related to the research process. I had no reason to feel any risk for myself or others. Covid related guidance was strictly followed.

As explained before, unclassified data materialized when it was determined that it was not relevant for the research related questions.

3.12 Issues of Trustworthiness

A qualitative study has its limitations. Research quality relies heavily on individual skills of the researcher and can be influenced by his or her biases. The large volume of data makes interpretation harder, and rigor can be difficult to maintain and demonstrate. The presence of the researcher can affect the responses he or she receives. To ensure that such limitations will not impact the research, the research supervisor (Dr. Paul Ellwood) actively intervened and provided guidance when research ran the risk of losing rigor. I was open to criticism and feedback and adjusted the course quickly.

Within qualitative research, the AR approach (the approach used in this research) has been criticized for lacking rigorous research design. In this research

however, the researcher used middle of the ground approach where the cycles were coordinated with the six steps of new product launch – however, the cycles were never made dependent upon the rigid structure of the six steps. The research control was independent of the steps and therefore the six steps were used more as markers rather than boundaries for cycles.

Identifying meaning can have its limitation also. Despite the flexibility, a lot is left to the judgment of the researcher. It can also suffer from a dearth of coherence as researcher develops ideas and themes that originate from the research data. Both consistency and cohesion can be improved by explicitly applying an epistemological position that can coherently reinforce the study's empirical claims. In this research, such a disciplined epistemological position was followed consistently as the interpretation of the data was being guided by the existing epistemological positions offered by the Conceptual Framework developed in the Literature Review (Chapter 2) as well as being constantly evaluated and checked with the participants.

The interpretation was guided through the framework and therefore the risk of interpretation coherence and consistency was reduced.

Care was taken to crosscheck findings. Data crosschecking was realized as data came from multiple sources (two professionals, documents, reflective journal). Theoretical crosschecking was achieved from the literature review. And the basis for inquiry were based upon multiple streams of research.

3.13 Limitations and Delimitations

This study will be limited on many levels. The limitations of qualitative research and the methodological choice limitations were addressed in the previous section. In addition to those the following limitations are important to consider: Data Science is a relatively new and emerging field. The application of data mining methodologies in the data science field (ML) is also new. My research only expresses the situation of a single entity and the viewpoint of the professionals

engaged in launching a new product within that entity. The concepts developed may not be applicable to other entities.

3.14 Summary

The method choice was to use participative AR and explore the data acquired from discussions, interactions, document inspections, observations, and reflections. The action-reflection in each cycle were expected to guide and help drive actions in the next cycles. Once the method was clarified, I was ready to move to the next step of the research. The *Findings and Cycle of Inquiry* journey lasted for several months where action and research complemented and reinforced each other. AIAI developed a product as the research guided and extracted the new product development process in ML.

Chapter 4 Findings and Cycle of Inquiry

4.1 Introduction

To recap, based upon the conceptual framework developed in Chapter 2 a baseline NPD process was assembled by stitching together two processes – the Concept Design Subprocess (what) and Product Development Subprocess (how). It was argued that the paradigmatic choices (Figure 2.9) which serve as the proxies for designer’s perceptions and beliefs about social and scientific realities influence the epistemological, ontological, and ethical properties of the system and hence influences the subsequent NPD process choices. The conceptual framework developed in Chapter 2 invokes the ML engineering process after a determination has been made about the social/scientific reality and business requirements have been determined. The engineering process is shown below. My goal was to explore and apply the process and extract those details that can help identify a more modern and relevant ML NPD framework – all with the backdrop of social/scientific paradigms. Accordingly, the first step of the ML process as identified via literature review is composed of business requirements – which, from an NPD perspective, requires a search for idea generation, clarification, and business value assessment. Only after that is completed is when design, development, and deployment (launch) can begin.

It should be pointed out that the learning set was structured in accordance with the goals of Action Research which are stated as follows (Willis, 2010, p.167):

Action research is inquiry or research in the context of focused efforts to improve the quality of an organization and its performance. It typically is designed and conducted by practitioners who analyze the data to improve their own practice (North Central Regional Educational Laboratory, www.ncrel.org).

Action researchers “see the development of theory or understanding as a by-product of the improvement of real situations, rather than application as a by-product of advances in ‘pure’ theory.” (Carr and Kremmis, 1986, p.28) (cited also in Wikiversity Action Learning article). This is a means to generate ideas (theory) that are relevant locally – to the people who are involved in the research, and to the environment in which it has taken place. (Wikiversity, en.wikiversity.org)

In the six cycles, my goal was to enrich the conceptual framework with feedback from the learning set participants. The dynamic dialogue and document inspection became the core courses of data. Each cycle was meant to represent a natural order of basic business process which captures idea-launch NPD lifecycle. The details about the learning from the cycles is presented below.

4.2 Cycle 1: How should we generate ideas about a new ML product?

Before even we start developing a product, we needed to understand how exactly do customers (users, developers) of AI develop ML concepts or generate ideas for ML products? Discover is typically the first step in new product development where companies try to identify problems that they can solve. AIAI had never launched a product before, and we were unsure about how to generate product ideas and in what areas. Would it be in supply chain, or marketing, or audit, or finance? To make that determination, two meetings were conducted. Both Randy and David participated in those meetings. They were aware that they are also participating in the research while working on the AIAI project.

Learning Set Session 1

The first action set meeting took place in early February 2020 and was attended by Randy and David. I clarified to the participants that AIAI has the capability to launch products in multiple areas, but we can only do that in one area. We must select a small set of problem domains to launch our product. In later cycles, we will develop processes to narrow it down to one area. In this cycle, we will develop a process to select a few feasible areas.

The conceptual framework developed in Chapter 2 became handy as the learning set members began the inquiry.

A typical NPD process is often started by conducting a broad survey and asking potential customers questions about their needs and wants. Deploying such an elaborate process did not meet the time and budget constraints of AIAI. I explained the dilemma to Randy and David, and they agreed that we do not have the budget or time to conduct surveys. As an alternative, I proposed that one way to study what customers seek in ML products and services is to extract rich data from customer Requests for Proposal (RFP) and Requests for Information (RFI).

We did not have direct access to customers via survey, but we did have several customer RFPs about their AI/ML needs. Customers issue RFPs to suppliers when they want to buy a product or service, and it is done as part of the sourcing process. In the RFPs customers thoroughly explain their needs and expect suppliers to explain their solution and to bid on the business. During the meeting we inspected several RFPs and the answers to the following questions I posed to the participants provided the requisite data:

- What are customers trying to achieve by bringing in the AI technology?
- What type of problems they were trying to solve with the AI technology?

To which David replied that some customers are not even giving us what they want in their RFPs, “they are asking us to tell them what they can do with machine learning.” he said.

5 SCOPE OF WORK

The Department of Homeland Security (DHS) requires services to support modernization and transformation of information technology (IT) systems and business processes used for: financial management; procurement and contract writing system management; and asset management and valuation.

The scope of this requirement includes the set of services needed to support the modernization and integration of financial management systems, procurement and contract writing systems, and asset management/valuation systems using EFiMS software. The scope includes technological

70RDAD19Q0000101


Page 31 of 71

Figure 4.1 Department of Homeland Security RFP

Some RFPs explicitly stated customer’s needs. Others did not. For example, Figure 4.2 shows one such RFP from Department of Homeland Security which has clearly specified requirements as declared by “The scope of this requirement includes the set of services needed to support the modernization...”. We observed that when it came to AI, there was a unique case of some customers who instead of giving us specific details of what they wanted, expected us (the suppliers or venders) to generate a vision for what they need and want. As we inspected those RFPs, we observed that in several cases government agencies were not providing any rigid specifications and instead were asking the suppliers to give them the plans. They called it *applied ideation or active ideation*. Details about some of these query RFP are available online (White House, 2016b, 2016a; Bur, 2019). Some of them were pure intellectual queries. Others were actual bidding for business – even though there was no indication for what the bids were for. The HHS (Health and Human Services RFP), to which our Institute responded (with other partners) and won the contract, did not ask us to submit bids on specific requirements for a product. Randy reminded us “they wanted to evaluate our abilities to create a vision for them and then build that vision.” As shown in the RFP below (Figure 4.3 the orange highlighted text) HHS was seeking *applied ideation* where they wanted “creative processes for the generating and developing

... solutions”. This implied that HHS expected the supplier to inform them what they needed to buy.

When we approached it critically, the team found this extremely odd that someone was asking for bids without knowing what they wanted. “We never saw it that way before. Strange!” Randy observed. Therefore, we concluded that when it comes to AI/ML not all customers may know with a high level of certainty about what they want. Many want the supplier to educate them on what they want and need. This, the team discussed, could be because it is a new technology and customers wanted suppliers to build a vision for them.



HHS Program Support Center (PSC)
Intelligent Automation/Artificial Intelligence (IAAI)
Solicitation Number: 19-233-SOL-00098
Data Call 1- Corporate Information and Experience

5. CAPABILITIES MATRIX
Please complete the table below with you capabilities (0= No Experience, 3=Strong Experience).

PWS Requirement Number	PWS Requirement Heading	Capability (0, 1, 2, 3)	Relevant Past Performance
C.4	C.4 FUNCTIONAL AREAS - REQUIREMENT FOR IAAI SERVICES AND PRODUCTS		
C.4.1	C.4.1 IAAI SPECIFIC REQUIREMENTS		
	The Contractor(s) shall provide IAAI solutions, support, knowledge, skills and abilities to meet overall organizational functions, mission and other activities of HHS and other Federal ordering agencies and offices. The functional areas to be under this contract are identified and described in the sections below:		
	1. IAAI Applied Ideation, Design and Process Engineering Support		
	This functional area supports the development and application of IAAI solutions. It encompasses the utilization of the creative process for generating and developing IAAI solutions to solve practice problems. This includes the scientific and analytical disciplines to conduct studies and experimentation directed towards advancing the state of the art: increasing the knowledge or understanding, concept formulation, assessment of system requirements, development, analysis and evaluation of concepts or technologies and development of operational concepts with the end goal being the application results, developing or improving existing IAAI capabilities. This functional area encompasses activities that will "enable" PSC to assist HHS and other Federal Agencies, to transition from "identifying" problems to brainstorming and designing solution. One practical example may include, but is not limited to, the conceptualization and high-level design of microservices, which involves and reengineering of underlying business processes and associated applications.		
	2. IAAI Engineering and Process Engineering Support		
	This functional area supports the development and application of IAAI disciplines to support HHS and other federal agencies, in the development of new and existing IAAI capabilities and systems. More specifically, this area supports the development of significant alterations to existing systems, integration and interfaces of existing solutions or software into different applications or platforms. Activities include, but are not limited to requirement elicitation (requirements gathering) analysis, specification development, which includes defining the architecture, components, interfaces and other characteristics. Practical examples may include, but are not limited to the technical design of microservices, the identification and engineering of interoperability requirements to include data formats and communications protocols, and the identification and engineering of the interplay between hardware, operating systems and applications (software).		

Figure 4.2 The Health and Human Services RFP (US)

ID11190043 - Advanced Analytics for the OCDO
Performance Work Statement

All tables and text in **RED** must be updated by the Contractor.

The Contractor must ensure any and all changes they make are done in **RED** text so the Government is able to review what was added or removed by the Contractor.

The Contractor must ensure the total page count does not exceed 5 pages beyond the original page count of 16 prior to the Contractor's changes, for a total of 21 pages.

1.0 Background and Purpose

1.1 Background

The mission of the Centers of Excellence (CoE), housed within the General Services Administration (GSA), is to accelerate Information Technology (IT) modernization, improve the public experience, and reduce legacy IT spending across the Government. CoE accomplish this by centralizing top government tech talent, leveraging private-sector best practices, and operating with a teaming mindset across Government departments and agencies. To better align with the dynamic effects of such a transformation, CoE provides a shared service solution for professional services to agencies adopted under the CoE initiative.

At the U.S. Department of Housing and Urban Development (HUD), the Data Analytics CoE was tasked to define the analytics outcomes and priorities while assessing existing data, tools, and systems used by HUD programs and offices. HUD also required the development of a strategy for analytics transformation that includes data governance, data management, and use cases for artificial intelligence technologies.

1.2 Purpose

HUD has a need to integrate Advanced Analytics techniques such as Artificial Intelligence, Machine Learning, Robotic Process Automation, and Geospatial Analytics with HUD's technical infrastructure and data to solve high priority issues. The team performing this work must validate certain use cases, gather and analyze data related to those use cases, develop a prototype, operationalize their prototype, and iterate and improve upon the design.

Figure 4.3 The Housing and Urban Development RFP (USA)

We then inspected Housing and Urban Development agency HUD (Figure 4.4) RFP and observed that they did provide some broad set of requirements. We observed that when narrating such requirements, these firms referred to systems by the human role the system was expected to perform. To clarify that point we discussed an example that when they were seeking a proposal for an automated invoice data entry and processing function, customer termed the solution as *Automated Accounts Payable Clerk*. As shown in Figure 4.5, the HUD RFP response also identified such roles – such as intelligent auditor, intelligent promoter, and intelligent ops manager.

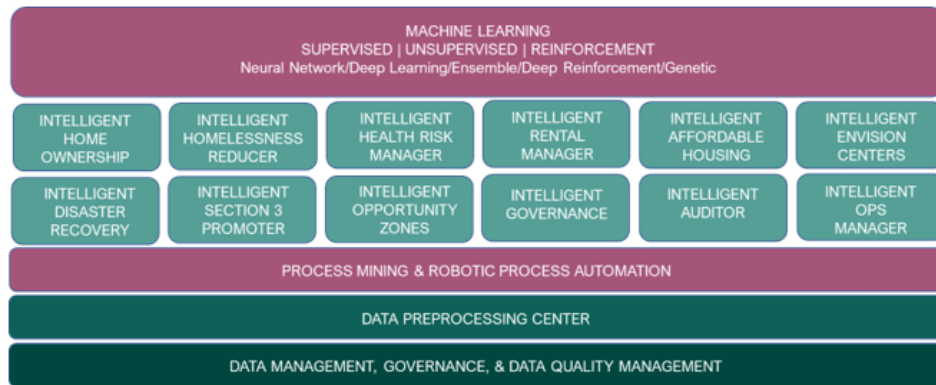


Figure 4.4 From HUD RFP Response

Related to the human-like roles, we also observed a pattern where companies were describing their needs in words such as “we are seeking digital workers”. What is a “digital worker” and what does that even mean? We asked ourselves. As our inspection of the documents continued, we also identified that customers who had given us more rigorous specifications in their RFPs were constructing their value proposition in terms of how AI/ML machines can eliminate human jobs. In one RFP related online document 5% of all jobs were mentioned as redundant (Boyd, 2019) – signifying that job cuts was appearing as a goal for automation. We recorded these observations, and a discussion took place about “why customers are viewing intelligent systems as *digital workers*?” and “why they are approaching systems implementation in terms of headcount reduction?”.

“But this is not new. Isn’t it true that all system, intelligent or unintelligent, were often justified based upon headcount reduction? What’s different with intelligent systems?” I raised the question. Randy clarified that the reduction in headcount should not be viewed as an isolated event without the accompanying use of terms like “digital workers”. “They are now looking at entire business processes and trying to automate them and replacing them with machines,” Randy said. “It is not just about making humans more productive, as it was with non-autonomous technology, it is about making humans irrelevant to perform work.” David used

the words “mass elimination of human jobs”. A digital worker signifies a standalone autonomous machine that can perform a type of human work – physical and cognitive – and can replace a human job function, the team determined.

The learning set members searched for documents that AIAI had used to answer RFPs and that are related to digital workers. One was a report authored by Deloitte (and Oxford) on automation whose press release began with the words (Deloitte, 2016):

Up to 861,000 public sector jobs – 16% of the overall workforce - could be automated by 2030 according to research by Deloitte, the business advisory firm. ... Deloitte’s previous work has shown that all sectors of the UK economy will be affected by automation in the next two decades, with 74% of jobs in transportation and storage, 59% of jobs in wholesale and retail and 56% of jobs in manufacturing having a high chance of being automated.

Looking at the percentages, the learning set participants recognized that automation is targeted to eliminate jobs and replace different forms of human work with machines (Deloitte, 2016; Rajan, 2016; Frey and Osborne, 2017). To this point, David asked that how it is different from industrial mechanization which has been going on for decades. We considered this question and the discussion pointed us to automation that is more human-like. “Industrial automation was simply automating highly repeatable tasks with limited degrees of freedom. But what we are doing now is automating tasks as humans perform tasks. Humans operate freely, move around, think, feel, and analyze. It seems that this automation is like automating a human.”, (paraphrased) David answered his own question.

However, this is not true in all cases. I reminded the team that we have received RFPs in which customers were asking for solutions that will not perform

autonomously but instead will only help humans to do their jobs better (for example, Figure 4.2). Randy pointed out that the traditional view of information technology making humans more productive will continue and AI/ML will also be deployed to achieve that. Referring to two of the RFPs we had inspected, we discovered that government customers expected 60% of workers will see a 30% improvement in productivity (Boyd, 2018, 2019). These expectations were set by the executives who issued the RFPs. This provided support for Randy's argument that AI/ML are not only sought for replacing human work but also to increase productivity. As we explored what it means for humans to become productive, we recognized that it includes the standard definition of productivity as *output increase for the same input*, but it also means having a better ability to see the future. "AI improves our ability to predict" David commented.

The first session concluded with collecting important data from RFPs and from the experience of the participants.

Learning Set Session 2: Identify New Product Launch Areas for AIAI

In mid-February 2020, we had the second meeting. The agenda for the meeting was to understand the type of problems that are only solvable by ML and then to apply that to the AIAI problem. Building upon the learning from the last meeting, the team was excited. Since we had made the decision to use our own knowledge for the discovery process and not survey customers, the RFPs and RFIs became our primary source of information about the expressions of customer requirements, needs, and wants. In this meeting we continued with analyzing the RFPs. Our strategy was to inspect the RFPs and from them try to extract the problems that would have led the customers to author those RFPs. We postulated that if we could find references to problems, terms that signified problems, or other clues about what problems customers were trying to solve, we can create a catalog of those problems and then we can group them into similar classes. The

learning set participants agreed to the research approach.

SOURCE	OBSERVATION FROM DOCUMENTS	MEANING EXTRACTION
Housing and Urban Development agency of USA (Fig 4.6). USAID (Fig 4.7)	Customers were seeking to improve predictions and reduce errors and biases in audit.	It meant they lacked predictive ability.
Department of Homeland Security RFP (Fig 4.2), USAID (Fig 4.7)	Customers wanted systems to automate not just repeatable tasks but also more cognitive (thinking) tasks.	This meant they have a problem set where more thinking is needed or human thinking to solve the problem is error prone or expensive.
Walmart (Press Release)	Customers wanted mobility solutions to be able to move products via autonomous drones, cars, and trucks. Physical inventory audit performed by robots.	This means they have the problem of not being able to move products or people around safely and cost effectively.
Merck (Press Release)	Customers required a host of solutions that included autonomous research and development.	This means they needed new ways to accelerate and improve innovation.
USAID (Fig 4.7)	Customers wanted a solution to have various capacities including social, interactive, mobility, automation, and service automation.	This means they wanted to create ways to communicate better and influence behaviors.

Table 4-1 Cycle 1 Data 1

From the RFPs in which customers had expressed their needs, we first created a set of their solution needs (See Table 4.1) – this was the piling or the inventorying

process of ideas and concepts. Then we tried to extract the solution needs and classified them into various classes. For example, both USAID and HUD wanted systems that improve prediction. So, we assumed that not being able to predict will be one class of problem. I termed it the problem of prediction. In this category, customers viewed their inability to deal with uncertainty as a problem. This includes factors such as inability to understand what is transpiring, incapacity to know or process all the variables that define their environment, and failure to understand what the future holds for them. For example, a company's failure to predict a new market opportunity will fall under this.

From our collective experience, we recalled our interaction with and research on Walmart as a potential client and remembered that Walmart was using AI to solve various mobility problems. This means any drone, robot, or autonomous vehicle that Walmart expected to use for the distribution of goods. We identified that as a unique problem area since it dealt with robotics and placed such problems of autonomous mobility and manipulation of physical environment under the class of physical mobility problems. They involve robots that move, manipulate objects, and interact with their physical environments. Cognitive processing is a necessary part to solve problems of mobility since a machine must make decisions about movement in space (for example an autonomous car).

USAID RFP (Fig 4.7) covered possible solutions such as automated doctors, automated pharmacists, and automated social workers. This was a bit complicated, and a longer discussion pursued. We decided to classify them as two type of problems – *problems of sensory enhancement* and *problems of social interaction*. Randy argued that an automated doctor must solve its knowledge problem including diagnosing and treating, and I added to it that it must also know how to socially interact with humans. We classified the former as sensory enhancement and the latter as social interaction problem. We specified that in the sensory enhancement class, customers view problems as the challenge of

machines trying to make sense of the world by using data. This can be viewed as analogous to how humans perceive their world by using visual, sounds, touch, taste, and smell to make sense of their world. The computer vision systems, facial recognition systems, voice recognition etc. fall into this category. The keywords are “enhanced senses”. We defined the problem of social interaction as involving areas in which machines can help improve social networks, give their controllers social influence, and play a role in power dynamics of a society. This goes beyond enabling interaction (which internet does) and focuses on targeting, influencing, and shaping the behavioral aspects of social interactions.

As captured in the Department of Homeland Security RFP and HUD RFP (Figure 4.6), the words “The goal is to reduce errors and biases” stood out. We realized that this was a class of its own. Randy and I did not see it as different than the problems of sensory enhancement, but David convinced us that “bias and error” are audit functions and hence is a distinct class. For a lack of a better term, we just called it *the problem of thinking*. In this class, we determined, customers view their problems to be of unclear or deficient thinking capacity. Humans lack the mental capacity to solve certain problems efficiently. Such problems can be solved by machines.

Finally, from our interaction with Merck (pharmaceutical) we recalled that the firm wanted to pursue automated research and development. I worked on that relationship and led the discussion and reminded David and Randy that they wanted to create a machine that can discover new molecules and drugs. We discussed that this type of work is neither like a doctor’s work (practice of existing knowledge) nor an auditor’s job (validating assertions). This work required discovery. We termed this type of problem as *problems of automated innovativeness*. In this class, customers view their problem as having to use the traditional methods of scientific inquiry and investigation – and consequently

lacking innovativeness or speed of innovation. ML enables faster, automated, and data centric innovation.

Now the team had determined six different areas that can be considered as problem areas that AI can solve. They included: *the problems of sensory enhancement, problems of prediction, problems of thinking, problems of physical mobility, problems social interaction, and automated innovativeness*. I was of the opinion, that there could be more, but at this time we could only identify those.

2.1 Performance objectives

2.1.1 SOO Objectives

The contractor must achieve the following objectives:

1. **Develop or enhance models** - This includes Artificial Intelligence (AI), Machine Learning (ML), and Geospatial Information Systems (GIS) models across the enterprise. The goal is to reduce scope for errors and bias in audit (or other) selection processes and minimize the necessity for humans to perform tedious tasks through process automation.
2. **Provide end-to-end analytics capabilities** - This includes providing the capabilities necessary to perform AI/ML modeling, natural language processing, text mining, cluster analysis, predictive modeling, sentiment analysis, and object-oriented programming.
3. **Advocate for the better use of data** - Establish the groundwork for enterprise adoption and integration of business intelligence capabilities to enable data-driven decision making.
4. **Implement continuous improvement** - Using Agile methodologies, the work must be iterated upon and training, operations, and maintenance support must be provided to ensure the continued success of technical solutions.

Figure 4.5 HUD RFP Front Section



Issuance Date: March 8, 2018
Open Period: March 8, 2018 to March 7, 2019

USAID Broad Agency Announcement
for
Innovations for Improving Health System Performance

Federal Agency Name: The United States Agency for International Development (USAID)

Opportunity Title: Innovations for Improving Health System Performance

Opportunity Number: BAA-GH-IHSP-2018

Authority: This BAA is issued under Federal Acquisition Regulations (FAR) Part 35.016 (c).
This is not a FAR Part 15 Procurement.

Catalog of Federal Domestic Assistance (CFDA) Number: 98.001 USAID Foreign Assistance Programs for Overseas

I. Overview

The United States Agency for International Development (USAID) is issuing this Broad Agency Announcement (BAA) to seek participants to co-create, co-develop, co-invest, and collaborate on research and development interventions to identify new approaches for improving health system performance, which move past traditional approaches and include new ways to ensure knowledge management of advances in the field.

The intent of the BAA is to allow co-creation and co-design to the maximum extent to create high quality, effective partnerships with great efficiency in time and resources. USAID will invite selected for-profit and nonprofit, public and private organizations to co-create research and development (R&D) solutions to the Problem and Challenge Statements stated in this BAA, including those organizations that have ideas, expertise, resources, and/or funding to add to potential solutions.

Figure 4.6 USAID RFP

RFPs, the documents provided by customers as expressions of their visions and needs are not just documents. They are symbols of the strategic visions of entities and their technology identities. IT and identity in the workplace have been studied extensively and IT has been seen as a direct identity referent for people and part of the extended self (Stein et al., 2013). But what we were observing was not about IT influencing (human) identity construction through work practices and role relations or societal change – it was the utter extinction of human from the workplace and the creation of a new machine centric identity of work.

With that grim recognition, the team proceeded to apply the insights to the AIAI project. We determined that in order to streamline in what areas AIAI should focus on, AIAI needed to look at: i) the business areas that automate entire processes (a job function) and not just tasks, and ii) those that lead to human headcount reduction. The learning set shortlisted such business areas as audit, marketing, and finance. The meeting concluded and toward the end I requested the team to think about what type of problems can AI solve. **Action: Apply the learning to select the focus area for AIAI.**

Evaluation:

In the absence of direct customer input, the learning set conducted the research based upon documents issued by the clients. This was proxy to a survey, however, since the RFP documents reflected the express intent of customers to buy AI systems and the descriptions of such purchases, the team deemed it to have important value for the research. The feedback from the team provided additional sources of data. The inference mechanism varied across the insights obtained in this cycle. For example, to infer that customers did not have an explicit idea of what they wanted, since they asked for advice about what they should do, was premised upon our understanding that if customers truly knew what they wanted, they would have expressed that in the RFP. In another inference, we interpreted customer's requirements to customer goals. For example, we interpreted that customers required a host of solutions that included autonomous research and development implied that they needed new ways to accelerate and improve innovation. This was our assessment based upon our experience. We used a combination of the team's experience and the written statements in the RFP's to extrapolate customer intentions and goals.

AIAI needed a process to determine how to start thinking about the type of problems that AIAI's product can solve. RFPs are the voice of the customer. They express what customers want. The first observation that the team had made was that unlike conventional systems' RFPs in which customers almost always specified their requirements, in AI/ML systems some customers expected the supplier or vender to build a vision for them. This, the team concluded, could be related to the novelty of the technology and because customers lack the knowledge about how to even conceptualize a solution.

The second observation came from the customers who did specify what they wanted in their RFPs, and their overarching goal seemed to be to replace human workers in workplace, improve productivity, and enhance predictive ability.

While conventional systems also improve productivity and enhance predictive ability, the autonomous technology specifically replaces human job function via full automation. The team observed that customers sought “digital workers” while simultaneously stating the goal of eliminating of human jobs. This, the team observed, implies that customers were aiming for human worker replacement.

The team had relied upon the RFI/RFPs to get a picture of what customers want, and a picture was emerging. A cross-sectional analysis of multiple RFPs along with the experience that the team members had in AI related business and filling out responses to the RFP, we had gained insights about the problems for which customers sought AI systems as solutions. The classification and labeling of the problems were achieved in two steps. In the first step we made a list of problems identified from various RFPs and other documents. In the second step, we classified them into broad classes. We then developed terms for the classes, identified additional features of each type of problem, and we called the **set of six problems as problem dimensions of AI**.

As an action step, we applied the problem dimensions to study how to conceptualize the AIAI product. The application helped us determine that our product will focus on problems of sensory enhancement, prediction, thinking, and social interaction. By doing that we had precluded mobility and automated innovativeness. Notice that while AIAI only elected to focus on four of the six dimensions, the team had discovered all the six areas. That is why this research was important as not only it was solving the AIAI problem, but it was also helping to create what could be a more generic methodology. By classifying and approaching in this manner, we were able to funnel and frame the AIAI concept development in terms of customer problems.

Reflection:

Our inference process could have been materially biased. We did not have direct customer input. As we extrapolate, we were basing conclusions based upon our understanding of the reality. There was an inherent risk in that inference. The linguistic and semantic information contained in an RFP could be subjectively interpreted. For example, customers not stating their specific intent does not necessarily mean they do not know what they want. It is possible that they have decided to see which supplier can give them a vision that closely resembles their own vision. While this raises some concerns, the team had experience in working with clients for decades. Additionally, the experience of the team was not limited to what was obtained from the RFP. As the team members had also interacted with customers in the past, the team possessed knowledge of client capabilities and problems and was not completely oblivious to customer conditions.

Cycle 1 Findings

The learning set had identified that AI/ML solutions can be conceptualized by exploring the need to automate entire processes and replace humans by machine, improve productivity of human work, or enhance predictability. Furthermore, we expected **the six problem dimensions** to have a generic basis to identify a problem that can be solved by AI/ML. We used these findings to solve the AIAI problem, and I added them into the methodology this research is attempting to discover (Figure 4.8). If we were to walk into a new customer who was trying to determine whether they need an autonomous solution or not, we could ask them to observe their existing business challenges and then try to conceptualize their problems in terms of the problem dimensions. For example, if a customer is experiencing decline in sales, the customer can think if it is coming from communication (social interaction problem), or inability to forecast demand (prediction problem), or inability to conceptualize a good solution (problem of sensory enhancement or thinking), or failure to innovate (automated

innovativeness), or failure to distribute and fulfill orders (mobility). This applies to any firm developing an intelligent automation project. Since the necessary common feature of potential solutions to all the six problem dimensions is “intelligence”, no such problem dimension can be addressed by conventional technology. It is not possible to architect automated social interaction or automated innovation with conventional non-intelligent technology. The contrast became clear to the team. Regular IT can only solve certain kind of problems. For the six problem dimensions, customers problems will require AI/ML technology.

Finding 1: (Figure 4.7) Customers can generate ideas about new ML products by clarifying their goals and brainstorming Six Problem Dimensions

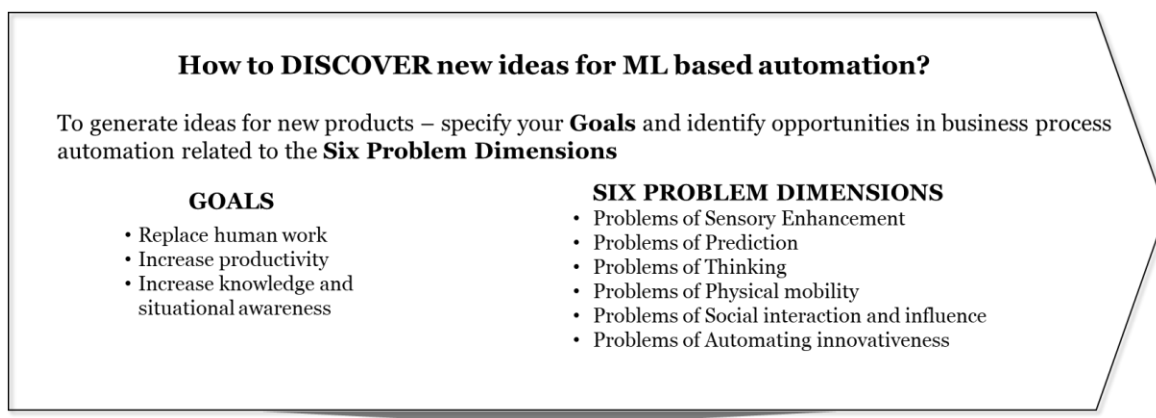


Figure 4.7 Cycle 1 Findings

4.3 Cycle 2: How to define and clarify a new ML product once its idea has been generated?

In this cycle the learning group wanted to build upon the knowledge gained from the previous cycle. The previous cycle enabled the learning set members to identify general process steps to generate product ideas.

The *Define* step identifies the type of problem that a firm will solve with ML in launching a new product. A firm can solve a set of problems, but for pragmatic and resource reasons, it must select a single focus area in which the firm will develop the product or service. I organized two meetings (late February and early March) to identify the specific area that AIAI should address. In preparation for the meeting, I collected examples of AI applications in business.

Learning Set Session 3 How to clarify and define an ML product and how is the process unique or different than other products?

Prior to the learning set meeting, I collected a sample of widely used artificial intelligence applications. For this and next two phases, I planned to use the sample to study and extract information relative for this study.

Randy and David participated in the meeting. Discussion specifically related to the study began when I asked the question: of all the areas that the institute can focus on, in which area should we launch the product? This question was not an easy question – since the institute had developed expertise in various business domains and offered courses, the product could have come from multiple areas.

I explained to the team that one way to answer that question is to study the role or function of the existing ML applications and use that knowledge to develop the criteria for determining in what area AIAI can focus best. Randy and David agreed with the approach.

I produced the list of sample AI applications and inquired the team members to state what role or function they think each application plays. Each of us took turns to go through the list of applications and narrated our view of the role and function of the application (Table 4.2). Upon analysis of the data collected and based upon my experience, three broad themes of *work automation, learning, and social*

began emerging (See Table 4.3). I asked the participants if they agreed with the classification of our feedback into the three categories and they agreed.

In the first round, the data was collected from the feedback obtained from the participants as they explored the functional elements and properties of ML systems in a non-technical descriptive fashion (Table 4.2). Three conceptual categories were extracted from the data: Learning, Work, and Social Interaction (Table 4.3).

The first conceptual category was Learning and was identified based upon the key words that related to learning. The second extracted category was work and was identified from the usage of work-related words such as *performs*, *analyzes*, and *automates*. Note that learning was different than work because a machine can automate work but not learn – for example, a non-autonomous car automates mobility, but it does not learn every time it is taken for a trip, whereas an autonomous car is expected to not only perform work but also learn from every trip.

Similarly, the class *social* was extracted based upon the word usage that neither implied work was being done nor learning was taking place, but instead some type of human to machine to machine-to-machine interaction was taking place. This class was different and distinct from learn and work because the nature of work of social interaction was different than functional performance related work or learning related work. For example, an autonomous car cab's interaction with the passenger about the destination or inquiring about customer's comfort are different than its ability to drive the passenger to the desired destination. It implied that for each technology, the application was learning something about some aspect important to a human or pertinent to a machine; that it was performing some work automation; and that it possessed some social dimension with which it played a role in human affairs and interacted with humans.

SAMPLE SET OF EXISTING ML APPLICATIONS ↓	<h2 style="text-align: center;">Summary Data from Raw Feedback</h2>
Phone personal assistant (Siri)	Humans delegate work to a machine. In the past a human would have done this work .
Facial recognition: <i>Ability to scan a crowd and identify people</i>	A machine can scan thousands of faces instantly. This work requires learning about people's faces.
Autonomous cars: No need for drivers	Automates driver's work . Constantly learns from experience of driving. Interacts with passengers and other vehicles and pedestrians.
Fraud detection	Improves machine work . Learns about malicious behavior and strategies. Reports behaviors and attacks. Interacts with auditors and risk managers.
Automated Trading	Improves machine work . Algorithms learn from markets and alternative data. Also learns how other machines trade. Algorithms function in a market and on behalf of their owners.
Text analysis and generation	Creates new knowledge. Learns about human produced knowledge. Authors . Gauges and extracts social and individual sentiments.
Health monitoring and diagnostics	Performs human work but also improves other machines. Monitors health. Performs diagnostics. Learns about human body. Reports analysis.
Discovery (R&D)	Invents. Automates new discovery. Performs scientist's work . Learns to discover. Reports findings.
Internet of things	Analyzes machines. Improves machine performance. Reports on machines. Interacts with other machines.
Social media	Enhances communications. Analyzes and profiles users. Learns about people and their behaviors . Induces interaction .
Marketing Analysis	Studies and analyzes markets and brands. Develops insights about marketing. Interacts and reports. Learns about markets.
Learning about customers	Learns about customers and customer habits. Makes buying recommendations. Interacts with customers.
Learning about weather conditions	Uses data from multiple sources to analyze weather patterns. Learns how to predict better. Reports results to humans.
Human Resources	Performs HR functions. Learns how to perform HR functions. Interacts .
Cybersecurity	Protects systems from cyberattacks. Learns to anticipate attacks. Reports back the results.
Financial Analyst	Performs financial market analysis. Learns how to value stocks. Interacts with human analysts.

Table 4-2 Meeting Notes Analysis Cycle 2

Extracted Concept Categories	Identifiers
Words/Terms related to LEARNING	Learns from experience Learns about malicious behavior Learns about faces Learns from markets Learns about human body Learns about people/behaviors Learns about markets Learns how to predict Learns how to perform Learns how to anticipate Learns how to value stocks
Words/Terms related to WORK	Performs analysis Scans Performs trades Drives (cars) Generates reports Makes recommendations Performs HR functions Protects Analyzes weather patterns
Words/Terms related to SOCIAL INTERACTION	Interacts with people Interacts with customers Trades [with other people or machines] Interacts with passengers Induces interaction. Interacts [with employees] Reports back results [to humans]

Table 4-3 Discovering Emerging Concepts

After the meeting, I used the feedback obtained from the team by reverse engineering the insights into their question forms (Table 4.4). Based upon the feedback, this could be the criteria using which AIAI could select its product area. I shared the criteria with the team in a follow-up call, and they agreed that this could provide the criteria that AIAI needs.

	HUMAN	MACHINES
LEARN	In a given problem domain, what can we learn about humans that we do not know?	In a given problem domain, what can we learn about machines what we do not know?
WORK	In a given problem domain, what human work can machines do or improve?	In a given problem domain, what machine work can machines do or improve?
SOCIAL	In a given problem domain, what human social interactions can be enabled by AI?	In a given problem domain, what machine social interactions can be enabled by AI?

Table 4-4 Model to generate ideas about Products/Services

Before the meeting concluded, the team participants discussed the broader meaning of having a system with working and learning system with social manifestation. “It is like a human worker” Randy expressed. **An action step was outlined to apply the model to the AIAI problem.**

Learning Set Meeting 4 Applying the Learning to Define the AIAI Product [Action]

This meeting was action-oriented meeting where the criteria developed from the previous step was used to develop a product/service area for AIAI. For each of the subject areas we taught at AIAI, we answered the above questions about the some of the problems we were familiar with. Randy and David participated in the meeting. As the team generated ideas, the research findings were put to test via action to identify the area in which AIAI will build its product (Table 4.5). Since all the details of that session are not central to this thesis, I will provide a summary representation of the ideas generated across multiple fields, which shows the application of the criteria developed in Meeting 3.

	Examples of Products	Work	Learn	Social
Marketing	Brand Management	A bot that automates	Learns to identify customer needs	Socially interact with customers
Finance	Investor relations Manager	A bot that automates the work performed by IR departments	Learns about investor expectations.	Interacts with customers.
Audit	An audit support product	A bot that provides preaudit information	Learns about potential audit clients	Interacts with auditors.

Table 4-5 AIAI Identifying its Product

During the exercise, the AIAI team determined that audit automation offered a better opportunity for AIAI. Before accepting a client engagement, auditors typically develop a thorough understanding of the client. This assessment can only be done from publicly available information. However, it is difficult for auditors to perform that review manually. When the engagement begins, auditors need to keep an eye on firm's activities and economics. Furthermore, auditors must have a thorough analysis of firm when pursuing it as a client. Finally, public information can also be a source of verification procedure during audit. For instance, the public information about a company's products, as acquired from customer comments, can provide a sense of company's sales projections or sales information. It is not efficient for each audit firm to develop this capability on its own. AIAI can provide that capability in a software as a service (SaaS) model. A decision was reached to pursue the audit product.

Prior to the conclusion of the meeting, I inquired if we had missed anything. Randy and David made some important suggestions to expound the *work*, *learn*, and *social* model. To further expand and improve the criteria it was suggested by the team members to incorporate three important aspects: 1) the automation part in

work automation has grades – as some parts of work can, or should, be fully automated while others would require human machine combined work; 2) The learning part is not a given – as depending upon the data and the algorithm performance, some tasks can be learned better than others. This means there is no guarantee that a system could be developed. The system is shaped by trial and error; and 3) The ethical and governance challenges must not be ignored. I captured this feedback and recognized that these elements were also reflected in my literature review. A sample of actual words (the raw input) and their inferred implications as enhancements to the model are shown in Table 4.6.

	Meeting 4 Recommendations’ Direct Quotations	Implication for the Core Model
Work	<ul style="list-style-type: none"> • “Not everyone wants their systems to be fully autonomous.” • “How much control you want to give to a system?” 	Automation is a relative process where some parts of the work are more automatable than others. ML Systems have grades of autonomy.
Learn	<ul style="list-style-type: none"> • “Unlike IT, you don’t know if an ML system would work.” • “Unless you train it, you don’t know if it is trainable.” • “The end result is always so uncertain.” • Customers are (or should be) concerned about the uncertainty 	A machine learning system is not guaranteed to work. Unlike a programmable or deterministic IT system, it could be impossible to build a solution given the current state of the technology.
Social	<ul style="list-style-type: none"> • “What about governance and ethics?” • “I recall two factors – accountability and transparency – from your model.” 	Ethics and governance are critical parts of the solution. The social aspects imply ethics and governance must be fully integrated into the solution scope.

Table 4-6 Identifying the Governance Needs

The research in audit revealed the key insight: specifically understanding the audit quality, the audit demand drivers, the audit supply drivers, and the measurement of success of an audit are important considerations (DeFond and Zhang, 2014). Audited financial reports serve as the core information upon which financial markets, regulators, investors and others rely. Despite being the foundation upon which the global economic activity is conducted, audit has failed the stakeholders on numerous occasions. In a dynamic business environment, audit risk is changing and automated audit is becoming far more prevalent (Meservy et al., 1992; Hoffman, 2017; Gepp et al., 2018; Merkl-Davies et al., 2011; Giroux and Cassell, 2011). There is even talk about constant audit and that it has the potential to become a new standard for governance (Chou et al., 2007; Jans and Hosseinpour, 2019; Dmitrenko and Matsegora, 2017). The AIAI system will observe numbers and records and evaluate vulnerabilities, discover inaccuracies, identify discrepancies, and provide intelligence. The system will see through the data, listen to the earnings calls, and develop instincts by learning how to classify problems.

Evaluation:

The criteria identified was extracted from studying the existing applications of machine learning to determine the common roles and functions across various applications. The convergence across the three themes of work, learn, and social was not only an indication of how to apply ML to develop products and services, but also represented a departure from the traditional approaches used in conventional systems. It manifested an acknowledgment that ML systems displayed anthropomorphism aspects in terms of that had what one would expect from a human – work, learn, and social.

There were three levels of inferential steps undertaken in this cycle. In the first step, the repeating theme of existing applications being used for work, learning, and social were identified. In the second step, the three functions were converted into questions. In the third step, the three themes were expanded and elaborated

to include autonomy, ethics, governance, and uncertainty. The first inference was based upon recognizing the linguistic and semantic usage of the terms. The second was not as much an inference as it was assuming that the themes of work, learn, and social when presented in a question form can be used to extract important system details from a user. In the third, the additions were experiential insights from the two practitioners.

Even if customers were able to express perfect requirements and it was possible to create a perfectly autonomous solution which will be functionally trustworthy for customers, we recognized that customers will still be concerned about “explainability”. In general terms, explainable AI is being able to explain to humans how machine made a specific decision (Kevin Casey, 2019). Customers are concerned about being able to explain machine decisions. In many algorithms when machines make decisions, such decisions are not explainable in terms of how the machine reached to that specific decision. In ML, machines make decisions based upon patterns and not necessarily reasoning. “If a judge wants to know why a machine recommended denial of loan for a given party, it is not possible to explain that in a deep learning system” remarked Randy. This means that in decision-making systems, customers will not be comfortable with the systems in which there is a likelihood that the owners or makers of those systems may not be able to explain the thought process of the system. This would be driven by regulatory, legal, and liability reasons.

Building upon the customer concerns topic, research data showed that customers are cautious about the uncertainty inherent in the AI systems. Unlike conventional systems whose plans are deterministic and outside of project management or technical failures, it is not possible to have engineering design failures, ML systems have significant uncertainty about if they will work or not. The uncertainty comes from the fact that at the beginning of the ML systems development, developers do not know if the algorithms will optimize given the data they have. This uncertainty

is concerning for customers. The larger the project scope, the greater the uncertainty. In other words, unlike conventional systems where proper planning can reveal a clear path to project success, ML systems are unpredictable in terms of whether they will work or not and that their development process is iterative. Unlike conventional systems where customers seek automation in areas based upon deterministic mathematical processes (addition, subtraction etc.), in machine learning systems they anticipated automation in areas of pattern recognition. Thus, while one type of system can only perform specific programable tasks, the ML system can learn and adapt.

Critical Reflection

The three concepts of *work*, *learn*, and *social* were extracted from the linguistic and semantic usage of those words to describe the functions and roles of ML systems. Such a pattern assumes that the linguistic and semantic consistency – implying that the learning set members usage of the terms was indeed meant in the sense they were stated. This implies that if the inquiry was undertaken from a different method, it may have resulted in different results. Secondly, the nature of the inquiry does not ensure that an exhaustible set of factors have been identified.

Turning them into questions is another leap of faith. If there was a bias (either due to linguistic or semantic concern) it now has been interjected into the mutated form of a question. This means that when the question will be posed to an audience, the effect of the bias could be magnified as it now has two degrees of separation from the original intent.

Keeping those constraints and risks in mind, we did proceed to apply the new learning to solve our problem. This was the nature of the inquiry we had undertaken. We had to live with those risks.

Additionally, our decision to select a focus area was largely based upon our internal AIAI factors. We did not explore the market need or assess the market. We

assumed that we are familiar with the market since we teach courses in them. We assumed that our courses provide us with enough interaction to learn about market needs. Our assumption could have been biased.

Was using human attributes to define machines an ostentatious undertaking which despite its obvious efficacy lacked substance? This bothered me. After all humans are complex, uniquely gifted, and self-conscious beings. How can something as complex and unique be copied into a machine?

Human values can be biased. The data on which systems get trained can be biased – since data is captured from human practices. When human practices are biased, they get captured in data. Using such data to give perception to a machine implies that what machine is learning could be biased. The paradigmatic inquiry is designed to address these questions. The Define phase ended and we were ready to launch the Design phase.

Finding 2: Customers can answer the Learning, Work, Social questions to define the product and can clarify product scope by addressing the MPD questions. (See Figure 4.8)

How to DEFINE new products for ML based automation?

To define and clarify new products – use LWS and MPD questions below to generate scope clarifications

Learn, Work, and Social (LWS)	Machine Personality and Decisions (MPD)
<ul style="list-style-type: none">• What can we learn about humans?• What human work can machines do or improve?• What can we learn about machines?• What machine work can new machines do or improve?• What social traits you want your system to have for human or machine interaction?	<ul style="list-style-type: none">• What type of thinking you want your systems to do?• What do you want your system to sense, hear, see, feel?• What do you expect your system’s personality be?• What kind of values you want your system to have?• How do you want your system to make decisions (for example, more vs. less conservative)

Figure 4.8 Cycle 2 Findings

4.4 Cycle 3: Learning how to design the ML system?

At this stage, we had enriched the conceptual framework on the first step of the ML methodology (Business Requirements) by extracting the steps involved in Discover and Define processes. This was helping us bridge the gap and move from a pure engineering methodology to an NPD framework. On the action side, based upon the insights from the previous two cycles, AIAI was able to pinpoint the area in which it will develop its product. The next step was to engage in the design activity. The goal of this step is to apply and extract design methodologies that are specific to ML and then compare and contrast them with the conventional methods.

Learning Set Meeting 5 Identifying Design basics in ML and how are they different or unique

In this meeting, the team focused on design considerations. I opened the meeting by asking the question: how should AIAI approach designing the problem of audit we identified in the previous cycle?

The current applications sample set had worked for us in the previous cycle. When I raised the question about the research approach for this cycle, the team agreed that we can continue that approach to identify the design issues. Since design can be broadly described as a depiction of inputs, processing, and outputs, we decided to study the inputs, processing, and outputs on the existing ML systems. Our goal was not to simply identify the inputs or the outputs – but to extract common themes across applications at a higher level of abstraction. However, to do that, we first proceeded to describe the inputs, processing, and outputs. The work product from the transcripts from the meeting are shown in Table 4.7.

SAMPLE SET OF EXISTING ML APPLICATIONS ↓	Inputs	Processing	Outputs
Phone personal assistant (Siri)	Received human voice used as command.	Understands the human command. Performs functions. Learns.	Completed task of what was assigned to it
Facial recognition: <i>Ability to scan a crowd and identify people</i>	Pictures or video data.	Learning about images. Classifying images.	Image match.
Autonomous cars	Visual input. Machine input. GPS data.	Applies driving skills. Learns how to drive.	Drives safely.
Fraud detection	Textual data files.	Detects fraud. Learns how to detect fraud	Stops intrusion.
Automated Trading	Textual data. Trading numerical data.	Learning about trading and how other machines are trading. Deciphers goals. Performs trades.	Makes a trade.
Text analysis and generation	Textual data.	Learning about human language. Performs specified functions.	Analysis report.
Health monitoring and diagnostics	Many types of digital data – pictures, numerical, text (clinical).	Performs analysis on clinical info. Learns.	Issues reports. Takes clinical actions.
Discovery (R&D)	Industry specific data from different fields – could be text, chemicals, mathematical relationships.	Invents and learns to invent more.	Outputs discovery.
Internet of things	Machine meta data.	Monitors machine health. Learns to monitor machine health better.	Machine health report.
Social media	Textual, image, audio, video	Facilitates communications. Establishes bonds. Learns to connect.	Analysis about people.
Marketing Analysis	Textual, image, audio, video	Performs analysis. Learns.	Reports. Recommendations.
Learning about customers	Textual, image, audio, video	Makes recommendations. Learns about customers.	Reports.
Learning about weather conditions	Textual, image, audio, video	Predicts. Learns to predict.	Reports.
Human Resources	Textual, image, audio, video	Performs HR work. Learns to perform work.	Various reports and actions (for example hiring decision)
Cybersecurity	Digital data, machine code, metadata, others	Safeguards assets. Learns to identify and tackle threats.	Intervenes and stops threats.
Financial Analyst	Financial data. Numerical and text.	Predicts. Learns to predict.	Issues valuation reports.

Table 4-7 Input, Output, Processing Data

The approach developed was to use the examples of existing systems and determine their inputs, outputs, and internal processes.

As we looked at the data from *input* entries (Column 2, Table 4.7), we realized that an intelligent system consumes and ingests data. It was not very different than how humans receive data from ear, nose, touch, eyes, or taste. Randy commented that all the input data implies that something is receiving it, storing it, and making it ready for the processing. Randy said that in the AI field, that is known as sensing and the device used to do that is known as a sensor. Thus, every ML system, we inferred, would need sensing.

When analyzing the processing data (Column 3 of Table 4.7), the team recognized the duality of what machine knows and its ongoing learning. We observed from the data that each of the processing example was composed of some unique function of performing an act that was based upon preexisting learning (for example, matching an image, reading a text, writing text, performing a diagnostic, etc.) and some type of learning from the specific experience of performing that act. Based upon some preexisting training of the algorithm, it is expected that the computer brain will identify some known pattern and respond to a given problem as an output. To codify it, the team called that process as “applying”. David commented that applying is not learning, it is using what has already been learned to issue a response. Just as someone who knows English would identify the difference between two letters in English alphabets based upon preexisting knowledge of the language, the computer offers a response that it already knows.

However, in ML, processing also involves some learning. Learning would be analogous to the person in the previous example trying to learn a foreign language and she will have to learn new alphabets that she did not know before. We codified that as “learning”. We recognized that an ML system will have two parts: 1) *apply* what it already knows; and 2) *learn* from the experience.

The team recognized that the data for *output* indicated (Column 4 in Table 4.7), that the system interacts back with the environment in the form of some action. That action, for example in the case of an autonomous car, could be to make a right turn. David mentioned that the term used in AI literature is “actuator” where the function of an actuator is act back upon the environment in which the system functions.

To develop a higher order design sequence, synthesizing the above four areas, I formulated the insights shown below:

Sensing: The receiving of data in the form on visual, sound, textual, machine data, or other types of input that comes into the system.

Applying: Using a trained algorithm to determine the course of action based upon the data. For example, if the input is a picture and the algorithm is trained to match a picture, the application of its current training will be Applying.

Learning: This is where the system accumulates experience by learning from each interaction of Sensing and Applying.

Acting: This is where the system acts back to the environment.

As the four areas became transparent, I termed that as the SALA model (acronym for sensing, applying, learning, and acting), and restructured them into the following four questions: *1) What will the system sense and what sensors will be needed? 2) What learning will system apply? 3) What will the system learn from each interaction and how will it accumulate experience? And 4) What specific action will the system produce back to its environment?* The answers to these four questions were expected to illuminate the design elements of the product. Before the meeting ended, we assigned ourselves the exercise to fill in the SALA model for the audit product we defined in the previous cycle.

Learning Set 6 Putting SALA into Action to design AIAI system

As the team met, the discussion began with synthesizing what each one of us had filled in as responses to the SALA model. There was some overlap in our thinking. Table 4.8 shows how we applied the SALA model to determine the design elements of the AIAI product. The AIAI product will have sensors that will accumulate the qualitative and quantitative data on firms; it will develop a profile of potential audit clients; it will learn from every inquiry; and it will issue reports back to the audit firms.

SALA	The AIAI Product
Sense	Qualitative and quantitative information of the activities and characteristics of the firm - Data from social media, financial statements, annual filings, patents, regulatory, and other such data
Apply	Develop a profile of a firm based upon several financial and nonfinancial factors. Produce an assessment of audit client's profile.
Learn	Learn from each interaction.
Act	Inform users about the profile of the audit client.

Table 4-8 The SALA Model

The remainder of the meeting shifted back to the issues of *ethics and governance* that constitute as one of the consistent themes in this thesis. Randy and David emphasized social dynamics of systems, pointing to the finding that people are viewing systems as humans in an organization. Despite having the basic design elements, participants recommended that ML artifacts need governance. The participants agreed and insisted that without governance ML artifacts can be risky and disruptive if not governed properly. Statements such as Randy's emphasis "AI must be governed." and David's reminder "you can't unleash AI without governance" were recorded as key data items for this critical area of design. The governance entails making sure that the artifact is safe and performs within the strict guidelines of expectations about the functionality of the artifact.

Using terms such as **bonding** and **deciphering goals** of other entities (see Table 4.2) is not something we do for inanimate entities. The discussion data from the team meetings showed that unlike conventional machines, ML based machines are active participants in social dynamics. As covered in the literature review (Section 2.4 of Chapter 2), technology, even when under human control, leads to social power, and with having its own autonomy can introduce unexpected power patterns. This implies that ML data has embedded social power in it and that power can be unleashed by using the data in different ways. For example, the report on audit client will greatly determine how the audit firm will perceive that client; for automated parole proceedings different combinations can create different outcomes for parolees; and for people seeking credit from automated credit authorization service different combinations can award or deny credit. Thus, different combinations of data can lead to different patterns of influence and social power. This implies that the designer's values and goals about how to design the system from data are important. Ethics matter at all stages of NPD of AI systems.

Since conventional products do not exhibit intelligence, the traditional NPD process does not contain any reference to understanding the social dynamics of a product. Furthermore, ML systems automation capacity varies with limited autonomy to full autonomy.

Evaluation:

The inferences made in this cycle were based upon the data collected to explain input, processing, and output. The inference for inputs was not a direct inference in terms of linguistic features. It was a logical inference based upon the argument that if a system receives data to solve a cognitive problem (a problem requiring thinking) it must receive some external inquiry in the form of data and that implies the system has sensors (akin to human eyes, ears etc.) that can capture the relevant data and pass that data onto the system's brain. This logical inference can be viewed as deductive in the sense that for the data to enter the

brain, it must pass through some physical medium (sensor), however it required the logical leap for interpretation.

The inference about the brain as having the two functions of using the existing learning to solve a problem while simultaneously learning from that experience was derived directly from the data as in each example from the sample, the data showed the presence of an action based upon an existing capability and a learning goal. Observing these common themes across multiple applications reinforced evidence about the dual nature of the artifacts. The common themes were that the system was using its existing learning to do something but also learning simultaneously. For example, Siri is responding to a command by its master (a human) but is also learning about his or her accentual variation. The similarity of the conjecture to human brain was incidental and did not go into the data as a variable.

Finally, the acting part of the application was also logical inference from the observation that in all forms the system was trying to affect its environment. This was also enabled by a level of abstraction where reason or rationale for that conclusion flowed from observing that a system incapable of acting upon the world will not be able to generate change.

Reflection

The overall approach to reduce the system into input, processing, and output necessarily required to view the machine in terms of predetermined reductionism. This epistemological view of a machine necessarily constrains the system to the mechanical viewpoint of cybernetic systems (Stacey, 2011). My bias toward viewing the world from a cybernetic viewpoint influenced me to breakdown the inquiry into those reductionist components. An alternative viewpoint – for example viewing a system as a complex adaptive system – may have led to a different insight.

At a first glance, the logical inference that data must flow through a physical medium was neither revolutionary nor in need for formal research to elucidate. However, in the context of this study, it was profound because it pointed to the variety of data (images, voice, temperature, videos, etc.) that can enter the system as sensory data and that the type of sensors needed should be considered with respect to not only the nature of the problem but also its future manifestations. For example, in tornado forecasting, we can receive data from satellites, but also from on the ground sensors (cameras) that may visually detect a tornado forming.

The contrast between learning and applying the existing learning seems a bit constraining as it does not consider those systems that may not have any preexisting learning at all. For example, within the domain of ML methods, unsupervised learning methods may not have any previous learning about the problem domain. They may simply (and madly) churn out millions of patterns based upon data fed to them, without knowing what to do with those patterns. The opposite is also true. A learning system may be designed in a manner where its learning has become so mature that it has no capacity to learn further. It will perform – but will not learn more from each interaction.

When built upon the learning from the previous cycles, a general theme was appearing: *intelligent machines are a manifestation of human as they exhibit human-like characteristics*. This was a remarkable departure from conventional systems.

Finding 3: Customers can efficiently design a system by identifying the sensing, applying, learning, and acting features of a product. Additional aspects of social power dynamics, ethics, and governance must be identified. (Figure 4.9)

How to DESIGN new products for ML based automation?

To efficiently design new products – describe the intelligent behavior and system functionality (SALA)

System Behavior

- What type of intelligence you want to embed in your system?
- What are social dynamics and behaviors of the system?
- What are the governance standards?

SALA Model

- What data would the software use “Sense”?
- What algorithm and model will be used “Apply”?
- What learning will materialize, and will it be ongoing “Learn”?
- What actions would the agent take to perform its function “Act”?

Figure 4.9 Cycle 3 Findings

4.5 Cycle 4: Learning how to understand the market potential or value creation of the system?

Discuss is a pause step after Design which precedes the development. The value creation potential of a new product is determined before investing in the development process. It provides insights about how to communicate the benefits and also how to position and sell the product. For the team, the goal of this step was to understand in what ways does the ML product creates value for clients and then contrast and compare that to conventional systems. The value creation is measured in financial terms.

Learning Sets 7 and 8: Learning how ML products create value

The team met to discuss how to determine the value creation potential of a technology. We decided that our focus of exploration will be pragmatic and financial. We recognized that while technology creates long term value for the economy, we wanted to focus on the benefits of the technology from a user’s perspective.

The research method applied to for this phase was to ask participants to answer the question “How is value created by this system?” for the sample set of existing applications. The meeting notes and feedback about value creation was then classified into the three Value Elements obtained from Cycle 1: Work Automation, Productivity Increase, and Enhanced Situational Awareness.

Each of the Value Elements was then linked with the *Financial Effect* to highlight on what specific financial dimension (for example, revenues, costs, profits, cost of capital, etc.). The link was established based upon the existing financial theory.

As we took turns to give our ideas about what value is being created by an ML system, the data in Table 4.9 was recorded. While commenting on value creation, the participants, in some cases also described how some of the Value Elements are measured. I am showing those key phrases in Column 3 of Table 4.8 in quotes as they enhance the data captured in Column 2. Keeping them in actual quotes form helps clarify the intent.

The determination that the raw data seemed to converge on the three aspects of the inquiry in cycle 1 came as a surprise. For example, the value proposition that less people will be needed to do a job was related to the automation; the observation that predictive ability will increase is related to the prediction ability; and the ability to serve more customers to the increase in productivity. Once classified, we needed to link them to a financial measure. Due to my extensive financial experience, this part of the exercise was primarily led by me. I helped establish the relationship between the automation and its impacts on profitability, cost of capital, growth, and other such financial measures.

Productivity improvement implies the output to input ratio becomes larger, and that means profits increase. A suboptimal state of awareness can be viewed as an alternative measure of higher risk. The team agreed with my assessment.

SAMPLE SET OF EXISTING ML APPLICATIONS ↓	How is value created?	Some key phrases from participants
Phone personal assistant (Siri)	Human assistant not needed. Personal efficiency increase.	“Headcount reduction is the best measure”
Facial recognition: <i>Ability to scan a crowd and identify people</i>	Humans not needed.	“Output increase” “More output for the same input” “Productivity increase”
Autonomous cars: No need for drivers	Human drivers not needed. Traffic efficiency increase.	
Fraud detection	Less fraud means cost reduction. Risk reduction.	“Risk reduction” “Risk of not automating” “Risk mitigation”
Automated Trading	Greater profits. Human replacement.	“Knowledge value” “Value of knowing” “Value of better prediction”
Text analysis and generation	New content generation. New knowledge. Human replacement.	
Health monitoring and diagnostics	Performs human work but also improves other machines.	
Discovery (R&D)	Increase in the quality, quantity, and efficiency of innovation	“Knowledge value” “innovation acceleration”
Internet of things	Machine performance improvement	
Social media	Communication and knowledge sharing	
Marketing Analysis	Learning about markets; ability to target new markets	“Cost reduction calculated by personnel cuts”
Learning about customers	Recommendations engine improves sales	“increasing revenues from better prediction”
Learning about weather conditions	Advance warning	
Human Resources	Better hiring. Less cost. Talent retention.	“Time saved in FTEs”
Cybersecurity	Cost reduction. Risk reduction.	“Measure cost of risk”
Financial Analyst	Faster, better analysis with less human resources.	“Cost reduction calculated by personnel cuts”

Table 4-9 Value Creation Model

The insights were framed into questions that constituted a model which can be used to determine the value creation from ML. This model, Value Creation from ML (VCML) is shown in Table 4.10.

QUESTIONS → VALUE ELEMENTS ↓	By how much does the product increase revenues?	How much cost reduction happens?	How does your technology improve capital efficiency?	How does your technology reduce risk?
Automation				
Productivity				
State of Awareness				
Estimated Dollars by Years				

Table 4-10 Value Creation from ML (VCML)

Filling in responses to these questions can provide a financial measure for implementing ML technology. The working session ended with the team deciding to update the VCML model for AIAI.

In a follow-up call (Work Session 8), using the VCML model, we established how an audit firm can calculate the return on investment from using our technology.

AIAI measured value by: a) calculating cost savings from number of FTE (full time employees) replaced by automation, b) calculating the productivity improvement (output vs. input); and c) calculating the value obtained from knowledge enhancement (for example, the financial value of having greater insights into a supply chain or financial market) and d) risk reduction.

The first three categories are related to the profits of a business. They impact costs (such as headcounts) and productivity (which can include both reduction in costs and increase in revenues). The knowledge enhancement is also connected to profits. For example, the knowledge about market opportunities and customers

translates into greater sales and profits. Similarly, the knowledge about market anomalies in trading can lead to investment profits.

Risk, the fourth category, can also be translated into financial value. It impacts a firm's cost of capital. This means that firms with higher risks will have to deliver higher economic profits to investors to attract capital. Firms with higher risk have a higher cost to attract capital.

The productivity increase could result from acceleration in innovation.

Participants noticed that ML could lead to acceleration in innovation. Innovation is typically an indicator of a firm's growth potential. Automating innovation implies that the firm can accelerate its growth prospects, increase revenues, and stay competitive. The meeting concluded as the participants approved the AIAI's application of the VCML model.

Evaluation

The research approach used in this cycle leaned heavily on prior findings. The inference of classifying the raw data into the three thematic elements of automation, productivity, and state of awareness was based upon the understanding of the team from the three concepts developed in Cycle 1. For example, the identification of human replacement or human labor reduction did not require greater interpretation than classifying that under automation.

However, there seems to be an overlap between productivity improvement and automation. Doesn't automation increase productivity by definition? To resolve that, productivity increase classification was used only when such an increase did not involve headcount reduction as a value proposition. For example, when a company uses recommendation engine (an ML product used by retailers like Amazon) it increases its sales but there is no associated headcount reduction since it was not the case that the company had employed thousands of people who were doing online recommendations and that the implementation of the

recommendation engine will lead to layoffs. It was a new process that improved the productivity.

The linking of the three factors of automation, productivity, and state of awareness was based upon my experience as former CFO. In that regard, the change in profits or risks can trigger any number of financial value creation. For example, supply chain risk reduction can lead to revenue increase from product availability, cost reduction from less spoilage, capital efficiency improvement from inventory reduction, and cost of capital reduction from less operating volatility. These relationships are widely available in the financial literature. The novelty of the findings was not as much as identifying these relationships but instead the model of value creation which provides a mechanism of determining the financial benefits of ML in NPD.

Reflection:

The linkage between the raw data and our cycle 1 findings may have been because we had gone through the cycle 1 and had that information. This could have produced the anchoring bias which means that unconsciously we were classifying the raw observations into classes that we were familiar with. We would not have known if we missed a specific class that was not part of the cycle 1. For example, how will we classify or place some financial value on an ML artifact that matches people for dating – our model cannot address that easily. It can be argued that matching dates reduces risk and saves time and cost of trying various dates before settling with a romantic partner, but that would be a stretch in our model.

The novelty of the technology sometimes makes it impossible or extremely hard to see how the technology will evolve. Our model assumes that the designer and offerors of technology will have precise information about the value creation potential. This is an exuberant and ambitious assumption. Technology development can take many paths and the benefits it provides to the customers

or users could turn out to be widely different than initially anticipated. For example, how can we estimate the value from e-mail.

Our findings and approach are intuitive to understand, but it was not the first time someone addressed this issue. I researched and discovered that in the early stages of CRISP-DM scholars did attempt to link the methodology with value and risk. For instance, a study by an industry consortium did reach to the similar conclusion (Chapman et al., 2000) as AIAI. Figure 4.10 shows their findings. However, those findings were limited to data mining and did not transcend to machine learning level (Pomykalski and Buzydlowski, 2017) which AIAI has now accomplished. Despite that, it served as a crosscheck to our finding.

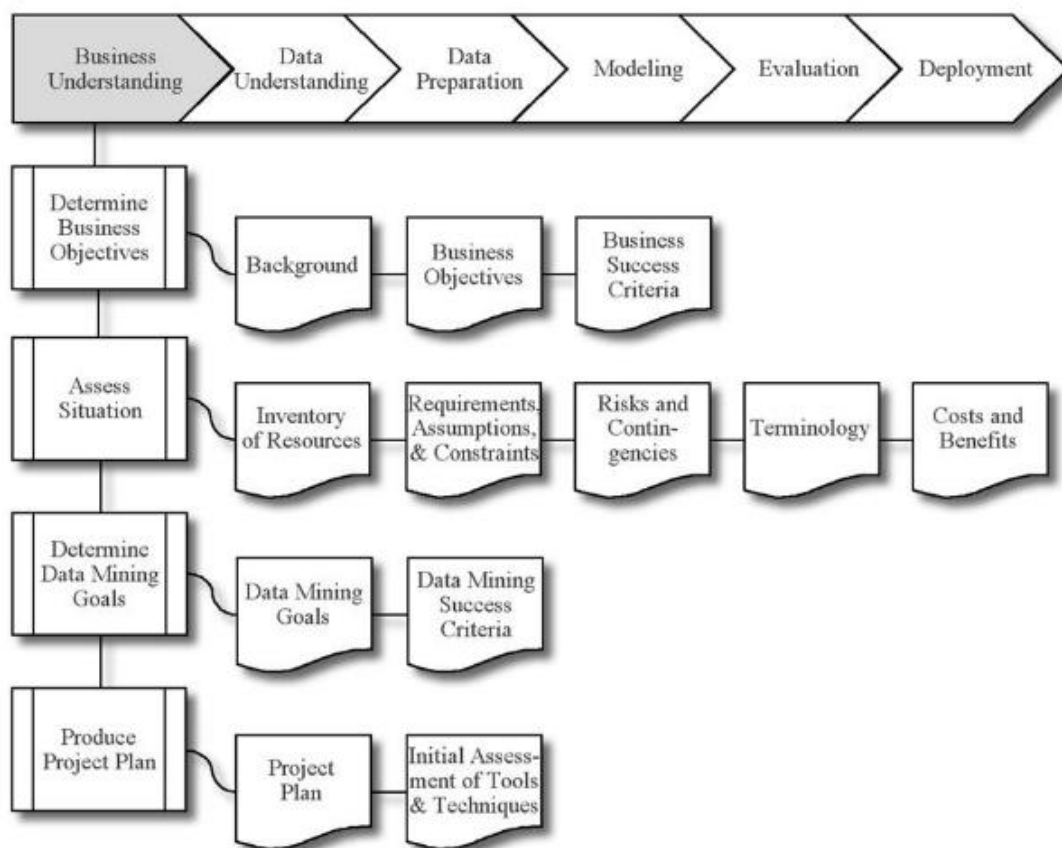


Figure 4.10 Adopted from Chapman et al. 2000

The data from this cycle linked the solution engineering and design with customer value creation in the business world. This step is necessary as customers need to build a business case to make investment in a product or a service. It is also critical so AIAI sales department can help customers understand the value of the proposed solutions.

Finding 4: Customers can estimate business value and develop a business case for AI/ML by calculating impact on profits, risk, and growth. While the general categories identified are the same as for conventional systems (increase profits, reduce risks), accelerated innovation is unique to automation. The concept of *innovation acceleration* is novel. (See Figure 4.11)

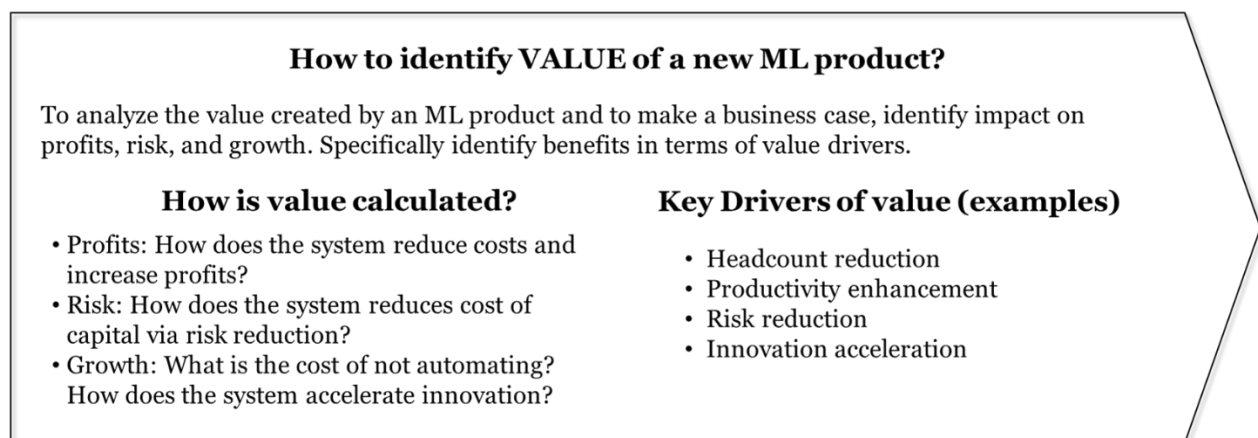


Figure 4.11 Cycle 4 Findings

This concluded our study of how ML creates value for customers. We were now ready to develop the product.

4.7 Cycle 5: Learning to Develop the ML Product

Moving from design to development stage was rapid. The goal of this cycle was to study the development process in ML systems and to observe any departures in comparison to the conventional systems. The team had to learn in action while conducting the development of the new product.

Once Randy and I had the design elements outlined, we both went in high gear to get a prototype ready. It was a highly collaborative and intense experience. Launching Cycle 5 meant that AIAI had completed a significant part of the journey. The firm's and the team's success or failure was now sealed. Once product development begins, there is no going back. There is little that can be changed in terms of business strategy.

Learning Set 9 Identifying the ML development process

Even though I am calling this Learning Set Session 9, several activities, meetings, and calls took place in this phase. Specifically, this phase started with the product development and continued over the course of the entire product development and testing. In the first meeting of the cycle, Randy and David had informed me that the AIAI process will not be exactly similar to the established CRISP-DM process (Section 2.2, Chapter 2). To study the actual process employed, at each stage of the development, I posed the following question to the team:

What step are we applying to develop the product and how does this step contribute to a developed product?

Randy, David, and I took turns to describe our experience. The process began when we agreed that the recording of the steps undertaken for product development will represent our collective learning. Table 4.11 shows the steps undertaken by AIAI. I named the steps after the meeting.

Steps	Name	Description
Step 1	Modularize	We divided the problem into lowest common denominator “thinking agents” that will be needed.
Step 2	Data Inventory	We performed data review what data we have, how much, its quality, and what would we need to purchase.
Step 3	Data Preprocessing	We prepared the data for ingestion into the algorithm,
Step 4	Feature Engineering	We studied the various variables we had included in the data.
Step 5	Develop Model	We developed the learning model and selected various algorithms to try learning.
Step 6	Train Model	We trained the algorithms on the data.
Step 7	Test Model	We tested the training results to see which model was performing the best.
Step 8	Deploy metamodel and meta cognition	We deployed the final model for each module and developed inter-agent protocols so multiple agents can work collaboratively.
Step 9	Implementation	We implemented the software, post development.

Table 4-11 The Development Process

David, Randy, and I reported undertaking a total of 9 steps. Step 1 meant that the intelligent artifact was not viewed as a single algorithm with a single function but instead a composite of various agents working collaboratively and each having its

own algorithm. This ensemble of agents required breaking down the product into modules or parts. In step 2, for each of the agent, we reviewed our current data and evaluated if we had enough quality data to undertake training the algorithm. In step 3, for each module we prepared a dataset for ingestion into the algorithm. In step 4, for each dataset we studied the mathematical properties of the data and tried to assess if we had enough variables. In step 5 we built the model and selected algorithms on which we will train our learning system. In step 6, we trained each of the algorithm. In step 7, we evaluated the performance of the algorithm. In step 8, we evaluated how the various modules will work with each in a collaborative manner. In step 9 we implemented the system.

These were the steps taken by AIAI and for the purpose of this study, I compared it with CRISP-DM and conventional systems development process and obtained the Table 4.12. This assessment was deductive as the certain actions are not undertaken in conventional systems.

Preprocessing (Step 3) is used to organize ML data so it can be fed into an algorithm. In AIAI's case, the development team conducted the following tasks: 1) Downloading ML data from public sources; 2) Organizing the ML data; 3) Studying the predictive power of ML data even before it is fed into the algorithm. The team discussed that preprocessing can be further segmented into the following sub-steps: **ML data governance, ML data quality management, ML data labeling, and ML feature engineering**. We reviewed and agreed with the list and commented that the sub-steps can also be automated.

Step	Name	CRISP-DM	Conventional
Step 1	Modularize	✗	✗
Step 2	Data Inventory	✗	✗
Step 3	Data Preprocessing	✓	✗
Step 4	Feature Engineering	✓	✗
Step 5	Develop Model	✓	✗
Step 6	Train Model	✓	✗
Step 7	Test Model	✓	✗
Step 8	Deploy metamodel and meta cognition	✗	✗
Step 9	Implementation	✓	✓

Table 4-12 Comparison with Conventional

The other set of activities (Step 5) identified is model building. In this step data scientists decide what type of learning (supervised, unsupervised, reinforcement, other) takes place and which algorithm will be used to train the learning artifact. Each learning method contains several algorithms, and the developer has to select the best algorithm. In AIAI's case, the firm was trying to build multiple predictive and analytical products, so the firm dealt with multiple models. Each model required a different algorithm. The participants reported that the developers began by looking at various models and testing them with sample data to determine their performance. We recorded that while they knew what type of learning will be undertaken (for example supervised vs. unsupervised), they did not know beforehand that which algorithm will work and what parts of data will be more

informative for learning in a specific model. It turned out to be an iterative process where the development team continued testing various algorithms for performance. I recorded that the process can be segmented into: a) Identify the learning method; b) Select possible algorithms as candidates (this selection is done based upon the characteristics of data and the scope/goal of the solution); c) Test the algorithm's initial performance; d) Iterate till results are obtained. The participants reviewed and acknowledged the process steps; and suggested that model building is unique to machine learning. "Conventional systems are not built with algorithms and models" Randy commented.

Another set of activities was performance evaluation (Step 7). The team reported that after an algorithm was selected, it was further scrutinized for performance. Once AIAI ran the algorithm, model's performance was evaluated to determine false positives or false negatives. This was somewhat equivalent to testing in conventional systems but while testing in conventional systems tests for errors generated by a deterministic system, ML tests for learning ability and performance of a multidimensional learning system.

Finally Step 8 was identified as the meta model development. A meta model is where the designers and developers of the model understand how various algorithms work collectively to create a shared result. In this case, it is understood that a single model may not be enough to solve a problem and a composite model will be necessary. The model of models (ensemble) is known as the meta model.

Evaluation

Unlike other cycles, the research method in this cycle was not based upon first discovering and postulating a process method, framing it, and applying it for AIAI purposes. It was the opposite. In this cycle we first undertook the development steps and then reported what we had done. Since this was our first time developing a product, it was deemed best that we start with a generic methodology and then

tread through the process by filling in the gaps. The gaps were filled as the team reported back that many steps undertaken were not part of any of the existing methodologies (See Table 4.12). The insights were extracted from experience and then a name was assigned to them. Care was taken to ensure that their stepwise order was maintained. The data shows what Randy and David did at AIAI and provides mutually exclusive boundaries between various steps. In other words, each process step can be viewed as a standalone action step and contingent upon the completion of the previous step. The successful completion of the AIAI product can be viewed as the single point of validation which provides evidence that at least for AIAI, the pursuance of these process steps led to a successful product development.

Reflection

The fact that this cycle's research reflects reporting on a process that was discovered after it was undertaken implies it has limited forward looking value. In other words, we know it worked at AIAI, but we do not know that if we apply it again, will it work for us or not. We also do not know if the process is transportable to other companies. Like any other process, we made some mistakes along the way and adjusted our path. Those mistakes are not reflected in the 9-Step process map identified here (Table 4.11). It shows what to do, but not what to avoid. The negative feedback value is not embedded in the model.

The two main findings that made this cycle valuable were: a) approaching ML development from an enterprise perspective requires building data management capabilities; and b) approaching ML systems from a multiagent interaction perspective which requires metacognition to manage the interactions of agents. These findings were based upon the example of the product that AIAI was designing. Do they contain the AIAI bias? I reflected upon this question and realized that even when companies may develop multiple ML products, they may

not necessarily approach them as interconnected agents. They may simply approach them as a collection of separate and individual entities with no intelligent interaction between them. However, I argued, that may not be the most optimized solution for a firm and hence firms may derive greater value from an integrated solution.

Unlike what we discovered; the baseline methodology tends to view the ML artifact development as a singular endeavor. In other words, it views development as if a firm will develop only one product. But in today's world companies may develop hundreds of ML solutions simultaneously. This implies that the baseline methodology should be expanded to include broader enterprise processes that form the foundation for providing and supporting data requirements for the multiple ML projects. The baseline methodology does not provide any means for establishing enterprise data management processes that support multiple ML projects. The critical insight was that data management and organization related activities should be part of the ML development process. If you are building a single ML product, you may not need an enterprise-wide formal data management program. However, as firms begin to launch company-wide ML programs, building standalone and formal data management and organization programs may become necessary.

Secondly, the insights that a ML artifact can be composed of various interlinked models was an important consideration. It can be viewed as a network of agents that work collaboratively to achieve goals. This network of agents is a necessary element of conducting and organizing work in modern times. For example, a recommendation algorithm for an online retailer may not work in isolation. It can be accompanied by others such as intelligent pricing agents, fulfillment agents, supply chain agent etc. to complete work. This implies that a metacognitive agent of some type will be needed to manage the interaction of various agents. No such

measure is included in the baseline methodology. I termed this step as metamodel planning.

Findings: To become avid developers of ML, firms may need to deploy some enterprise capabilities. While some of the development steps are common with conventional methods, two critical steps unique to ML systems are: 1) data management at enterprise level; and 2) metamodel planning. (Figure 4.12)

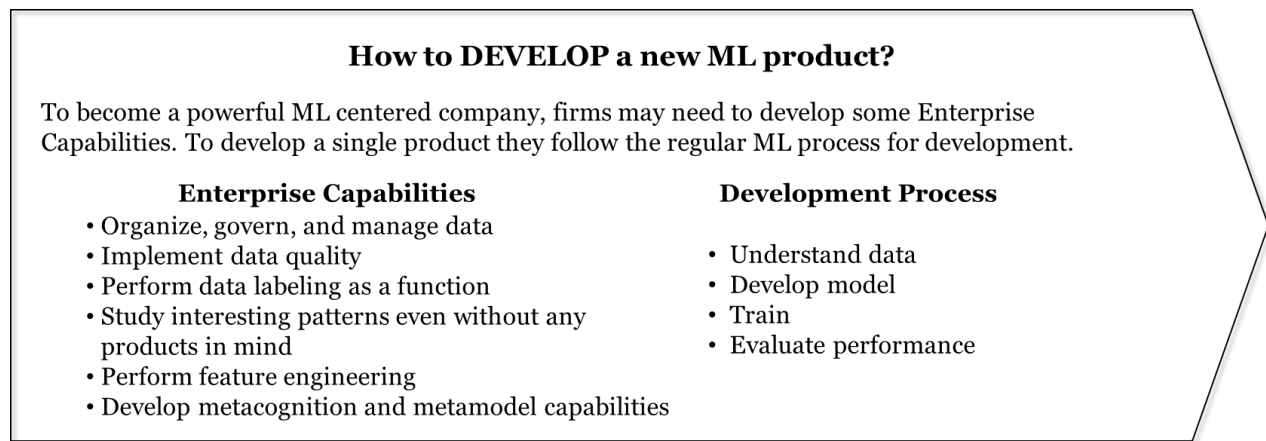


Figure 4.12 Cycle 5 Findings

We were now ready to deploy the product. It involved launching the next cycle of investigation. It is covered in the next section.

4.8 Cycle 6: Learning how to deploy the ML product

Once software is developed and tested, it is ready to be deployed in service. That activity is known as deployment.

Learning Set 10 Learning how to deploy ML product

Implementing an ML solution is not like implementing a conventional system, this much we knew. But to understand the specific requirements that led to the

deployment success at AIAI, this part of the research was one of the most active parts of the project. The decision-making related to deployment was rapid and the team met in short sessions spread across three workdays.

Data was gathered from two activities. First, a sample of existing ML applications was provided to the participants to gauge their feedback about what is the deployment method used for those applications.

Second, I asked the following specific question to the participants: What specific steps you took to deploy the ML product developed by AIAI?

Randy was primarily responsible for the technical deployment. He reported that his biggest concern was to select the delivery mechanism for the software. The intelligent artifact can be made available to clients via cloud (online) or from AIAI's own servers. When asked about how he will make the decision between those options, he replied that "it came down to cost and security." AIAI did not have the resources to manage its own data center or servers. AIAI needed to depend upon existing cloud providers like Amazon.

Randy and David reported that other deployment also includes other IT related activities such as loading the software on the cloud, securing it (cybersecurity), testing it in live production environment, integrating with other software (if needed), and optimizing are also part of the deployment. It may also include other aspects such as how will users interact with the system, user interface, etc.

While technical deployment is self-explanatory, *Organizational Deployment* refers to the aspects of training people on how to use, manage, and interact with intelligent systems. I raised that as an issue and participants agreed that this was an important consideration. This implied *learning about a learning system*. In some ways, this reminded me of the single-loop and double-loop learning concepts (Argyris, 1976). We discussed this concept and realized that we must factor in the state of change that will arise as we implement new learning in the

organization – but we must also factor in that a learning system is also constantly changing and evolving. Therefore, the change it is producing in the environment is not a one-time change. In reality, it is forming a loop where changes in the environment affects both learning systems and the learning organization (human managed) in which the learning systems are deployed and the interaction of the two affect each other. This means that organizational change management will be an ongoing factor. For example, as we introduce autonomous job interviewing robots, we will have to train the interviewees, the human interviewers, the management team, the HR group, the business group, the legal group, the compliance group, and more. As the interviewing system becomes more intelligent, it will gain greater social power and the change management process will be ongoing. This part of the thinking was primarily the concepts that I thought about, however when I discussed them with the team, there was wide acceptance that organizations should be aware of such changes.

The learning set participants reflected back on the set of existing ML applications used in Cycle 2, 3 and 4 and realized that the systems can be deployed in multiple platforms.

SAMPLE SET OF EXISTING ML APPLICATIONS ↓	How is this system deployed?	Additional Explanatory Comments
Phone personal assistant (Siri)	iPhone	
Facial recognition: <i>Ability to scan a crowd and identify people</i>	Servers/Cloud	
Autonomous cars	In a car (automobile)	“multiple intelligent systems have to work together” “monitoring this system means to have some type of meta-cognition”
Fraud detection	Servers/Cloud	
Automated Trading	Servers/Cloud	“could be multiple systems”
Text analysis and generation	Servers/Cloud You own computer	
Health monitoring and diagnostics	Servers/Cloud	“multiple clinical systems work in collaboration”
Discovery (R&D)	Servers/Cloud	
Internet of things	Servers/Cloud	
Social media	Servers/Cloud	
Marketing Analysis	Servers/Cloud	
Learning about customers	Servers/Cloud	
Learning about weather conditions	Servers/Cloud	
Human Resources	Servers/Cloud	
Cybersecurity	Servers/Cloud	“multiple intelligent systems”
Financial Analyst	Servers/Cloud	

Table 4-13 Exploring Deployment Paths

The discussion developed to an important topic. An intelligent system constantly learns. Arguably, the team members noticed, it also collaborates with other intelligent and unintelligent systems. This means if a cluster of ML artifacts are

working in collaboration with each other, the deployment considerations must address how the systems will communicate with each other. Furthermore, monitoring those systems, as reported in column 3 of Table 4.13, required metacognition of the states and interactions of various systems.

I incorporated the feedback into a model that was used to determine the deployment path for AIAI and termed it as Deployment Evaluation model. It is shown in Table 4.14:

TECHNICAL	
How will your customers access the solution? [Cloud, Internal or Device, or On-your-Servers]	
What is the cost of each deployment option?	
Does your system contain multiple intelligent agents?	
How do your intelligent agents communicate with each other?	
How do intelligent agents work together in a collaborative manner?	
How do you monitor the state of the system when multiple agents are involved?	
What can fail?	
ORGANIZATIONAL	
What change management and organizational issues are you concerned about?	
What type of training will be needed to learn how to use the system?	
How role will the system play in an organization? How will it shift social power in an organization?	
What other groups in the firm will need to be informed and trained about the system?	

Table 4-14 The Deployment Process for AIAI

Evaluation

The insights obtained in this cycle were based upon the experience of the participants and the specific requirements of AIAI. The process of raw data collection to extrapolate special insights was based upon two parts. In the first part a technical deployment evaluation is made. This was accomplished by recording the specific process steps taken by the AIAI team. Since the team was implementing an ML system, it is inferred that the team would have taken steps to make the ML deployment successful. Thus, recording the specific steps undertaken can provide some insights to the functioning of ML or Data Science departments.

Reflection

The AIAI team developed the deployment steps to support the specific requirements of the AIAI product launch. The data related to process steps, therefore, will be highly biased to the AIAI circumstances. However, when I further reflected upon that, it became clear to me that the technical considerations would not be that different across other implementations. Questions such as should the system be deployed on Cloud or company's own servers are standard concerns. What made the team insights extremely different were the aspects that were specific to the ML systems. ML solutions are learning systems and they may need to constantly learn from new experiences.

Second, the participants recognized that an intelligent system, specially the one that is part of a work chain or interacts with other system, is a social system. Being part of a social network implies that the system needs to be adaptive and responsive to changes in the environment. Third, the participants also recognized that the system needs to have some type of a metacognition. This means it needed to be aware of its own states.

Deployment processes can be broken down into two main categories: Technical Deployment and Organizational Deployment. The recognition that our learning system needs to be adaptive and social had a special meaning for us. It meant we would constantly need to update the learning capacity of the system and that the system needed to have metacognition and knowledge about its own states. This implies that in addition to the traditional deployment issues (such as integration, cybersecurity etc.), we needed to consider the mechanism by which changes in the environment will be tracked to constantly evaluate if the learning system is adapting to the changes in the environment or not.

Findings: In addition to the technical aspects, *deployment* is a social process manifested in educating and communicating with stakeholders. It also requires viewing learning technology as a growing and adaptive system that interacts with its environment and hence changes in the environment and the associated adaptability needs to be monitored. (Figure 4.13)

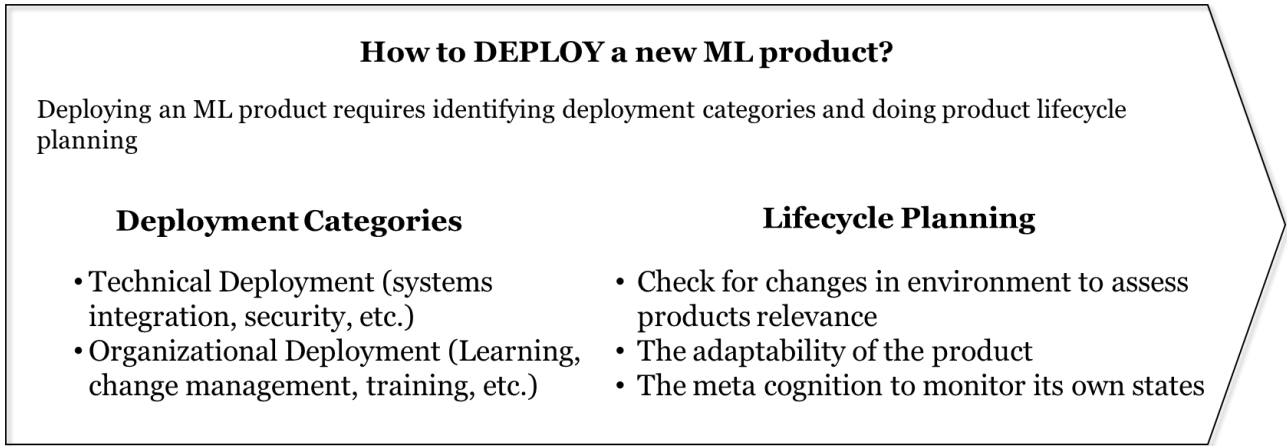


Figure 4.13 Cycle 6 Findings

4.9 Customer Input and Verification

A previous study performed to launch new data driven products vs. conventional products (Lee et al., 2019) identified four iterative phases of Discover, Define,

Develop, and Deliver. During the product development process, we were unable to receive direct feedback from customers. While the business problem was solved – as AIAI was able to develop and deploy a product – the customer feedback was necessary to determine the model’s applicability and generalizability. This would ensure that both relevance and rigor requirements of the action research methodology were met. I received direct customer feedback on the AIAI NPD framework. This was done selectively and confidentially as AIAI did not want to share the model outside the firm. A NY based data company, a NY based hedge fund, an airline, and a tech firm were contacted to receive feedback on the AIAI NPD framework. In general, all the practitioners agreed they do not have NPD framework or methodology for developing ML products. They also acknowledged that they are developing ML products and consider such products as their competitive advantage. Despite such an important role these products play in their business, these customers did not have an NPD framework to guide them in the idea-to-launch process for ML products. Overall, the customers were very interested in adopting the AIAI NPD framework. Specifically, it gave them many ways to generate ideas about products and also to explore the various manifestations of the product based upon social-scientific paradigms. In terms of limitations, two practitioners did point out that some parts of the framework required making major changes in the firm and that implementing a change of that scale will require a board and CEO level action and initiative.

4.10 Competitive Advantage for AIAI

AIAI’s goal was to develop a product and while developing the product discover the framework. With the positive reviews of the customers, it was clear that in the process AIAI had developed another offering: the AIAI NPD Framework. The framework discovered is a marketable product and AIAI intends to offer this to its clients. This has already led to the creation of a competitive advantage for AIAI as no other firm, to my knowledge, offers a comprehensive ML NPD framework.

We conducted this research in an AR setting and obtained tremendous benefit from the research (Greenwood et al., 1993; Torbert, 2001). This was no ordinary change. We were redefining the future of the company. We were transitioning the firm from an academic institute to a Silicon Valley type company. We were cognizant that authenticity is exemplified by four factors: be attentive to the data; be intelligent in inquiry; be reasonable in making judgements; and be responsible in making decisions and taking action (Coghlan and Brannick, 2014; Coghlan, 2001). AIAI had a product that AIAI can demo to clients and had discovered an NPD framework in the process.

Through the AR, AIAI had successfully launched its product, all while discovering and extracting new knowledge. In summary, the findings of the six cycles of action research showed that the baseline engineering methodology was enriched to a significantly broader and rich collection of capabilities to conceptualize, design, develop, and implement ML products. The AIAI NPD Framework and its application are discussed in the next chapter.

Chapter 5 Discussion

5.1 Introduction

In this chapter I will discuss the findings, contribution to knowledge, and then discuss the applications of the research in practice.

AI and ML have become pervasive and powerful forces for business transformation (Makridakis, 2017; Wright and Schultz, 2018). As companies are recognizing that ML is the primary driver of creating competitive advantage in the contemporary times, they are embracing the technology. Embracing the technology implies that the technology is properly integrated into the strategic frameworks of a firm, that it is properly conceptualized as a driver of value and competitive advantage, that it is adopted with an understanding of the state of the technology, and that it is implemented with success (Naqvi and Munoz, 2020). Unfortunately, the recent data reported about the technology adoption of artificial intelligence shows a failure rate of 75 to 85% (Nimdzi Insights Pactera EDGE, 2019). With over \$7 billion being invested into the AI field each year (Shoham et al., 2018), and a failure rate of 75 to 85% implies unless remedial improvements are made, over \$5 billion annualized value will be at risk. The action research in this thesis focused on developing a framework that can improve the practice outcomes for new ML products.

However, developing an ML system is not just an engineering. It has two other extremely important dimensions. The business dimension captures the business logic and case, business requirements, concepts, business needs, and other such factors that are relevant to both conceptualize and financially justify the system. The engineering methodology cannot be initiated without first architecting the business dimension. The third is the social dimension. Systems conceptualization, design, and development are social processes with underlying ethical and governance consideration, and hence social and ethical aspects must

be considered. Collectively, the three dimensions of business, engineering, and social constitute as critical elements of an effective NPD methodology.

5.2 The AIAI Model – A Review

The research began with first establishing a baseline NPD process for launching intelligent products. The conceptual framework, developed in Chapter 2, served as the baseline and was composed of a paradigm plus engineering process (see Figure 2.9). In this action research I field tested the baseline conceptual framework in an active work setting of launching an actual new product launch project for AIAI. The inquiry lasted six action cycles.

The conceptual framework developed in Chapter 2 (Figure 2.9) assumed that the product construction or building process will begin after reaching a “Business Understanding”. The term implied that business requirements for the product or service being developed will be known. However, in practical applications, such an understanding is usually unclear. Specially in new product development areas, this clarity is greatly absent. As we went through the action research to build our product, we gained knowledge in action and with that knowledge filled in the gaps in the conceptual framework. This gap-filling not only added several new steps in the process but also expanded upon the existing steps by making them more intuitive and detailed as well as adding *ethics and governance* into the framework.

The findings from the previous chapter showed that once findings and lessons were classified, they represented an NPD process. The findings led to the creation of a new framework, which was developed using the skeleton of the baseline methodology, enriched via fieldwork, and turned out to have meaningful and substantial differences than conventional models. The new framework is composed of 6 categories, 12 distinct action sequences, and over 30 tasks (see Figure 5.1). To incorporate the social dimension, the model derived from the

research is also composed of paradigms of ML which serves as the gateway to the NPD process (STEP 2). Collectively, the model is known as the AIAI NPD Framework.

The left side of the AIAI NPD Framework (Figure 5.1) is composed of four paradigms of machine learning. The four paradigms represent the theoretical assumptions that the designers make about the world in which their artifacts (ML products) dwell. These were covered in Chapter 2.

The purpose of exploring the scientific-social realities of intelligent systems is to be able to design, develop, and govern these systems in accordance with their underlying social and scientific assumptions. All the subsequent steps in the NPD process – for example the discovery, definition, design, development, and deployment parts – will be approached differently based upon the STEP 1 selection.

The ML NPD Process in Figure 5.1 is composed of two primary steps. In STEP 1, customers are expected to understand their assumptions about the social and scientific realities. Based upon this model, after making the paradigmatic determination, the product or service development teams sequentially operationalize the 12- action sequences to get to the final product or service. It is expected that as NPD teams conceptualize and plan new products, they will constantly reflect back on their paradigm assumptions so they can readjust the scope of the products. In accordance with the AIAI experience, and what this research suggests, following the sequence of steps implies that the product development team performs a set of actions that, if done right, facilitate the development of the product and service and increase the likelihood of success.

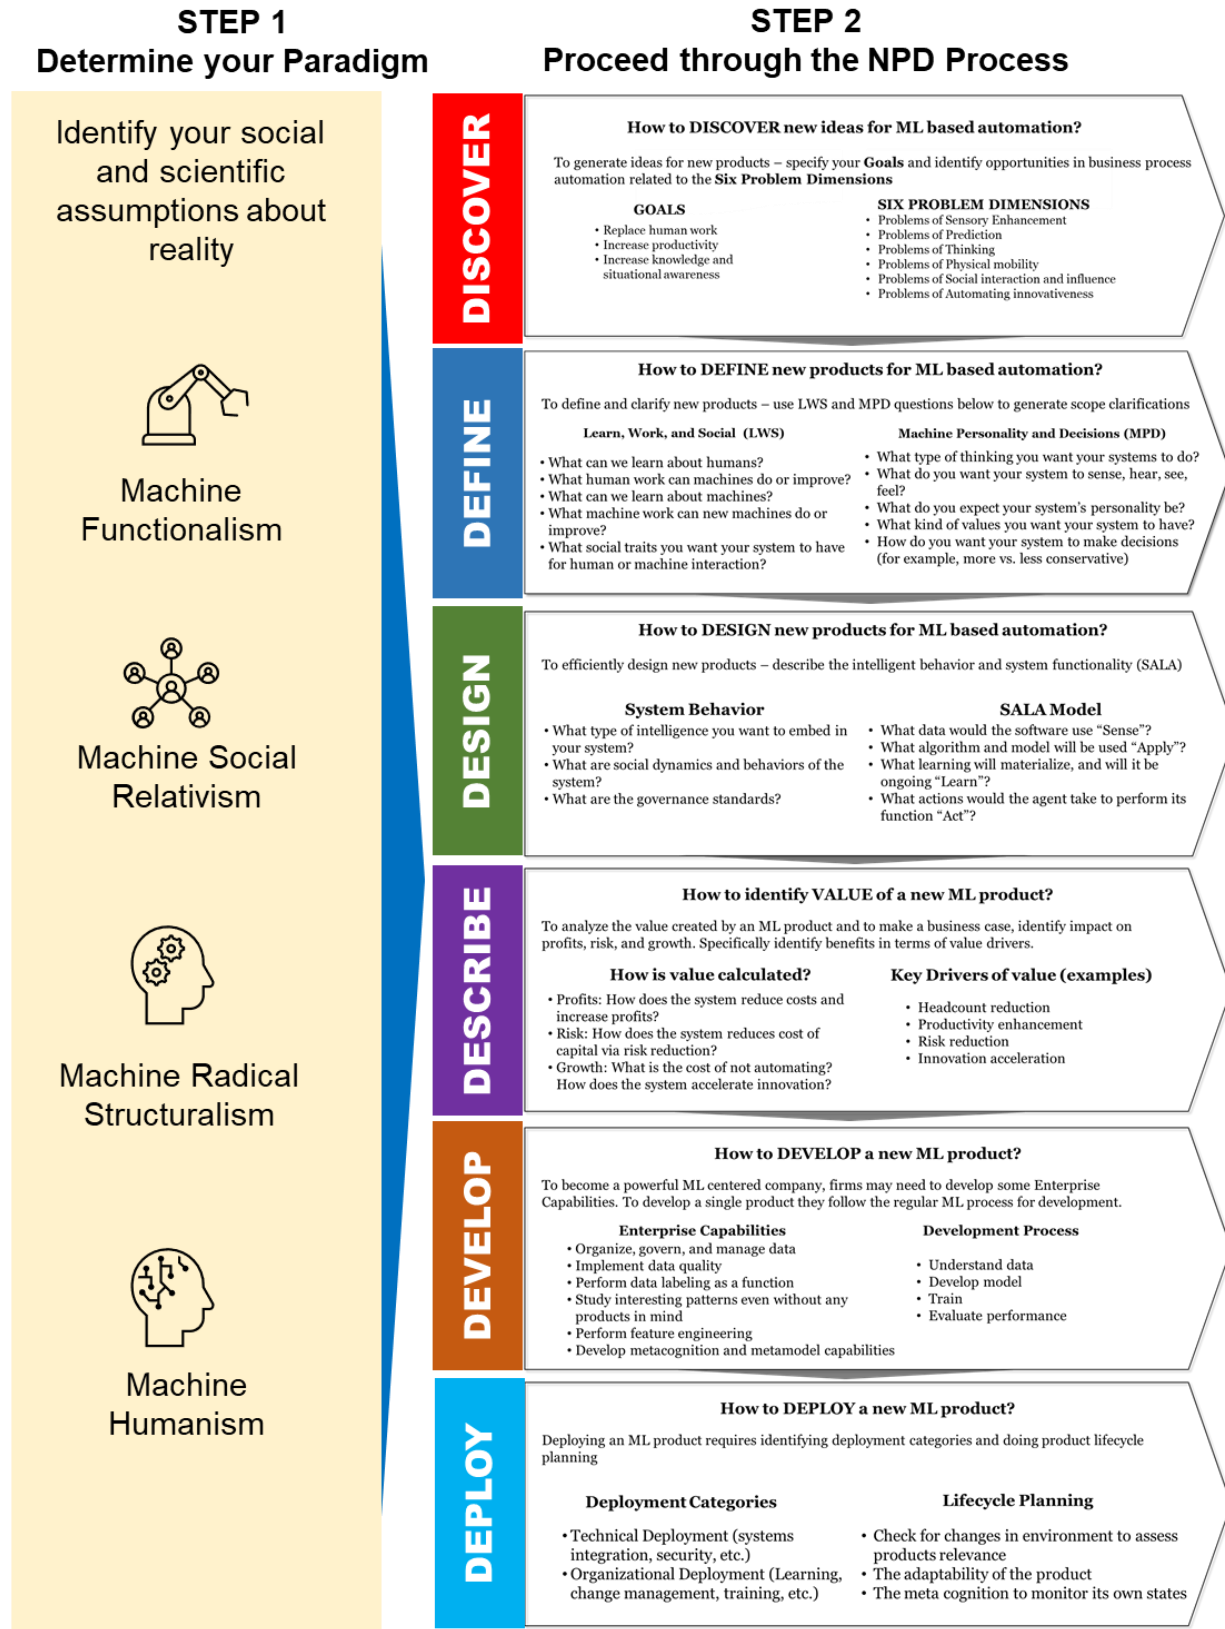


Figure 5.1 The AIAI NPD Framework Action Sequences

5.3 The AIAI Experience

When the NPD project for AIAI was launched, we had no NPD framework to work with. The dearth of methodologies was documented in Literature Review Chapter 2. Our risk for launching a new product without a methodology was profound. A conceptual framework was developed in Chapter 2, and it was field tested in this research.

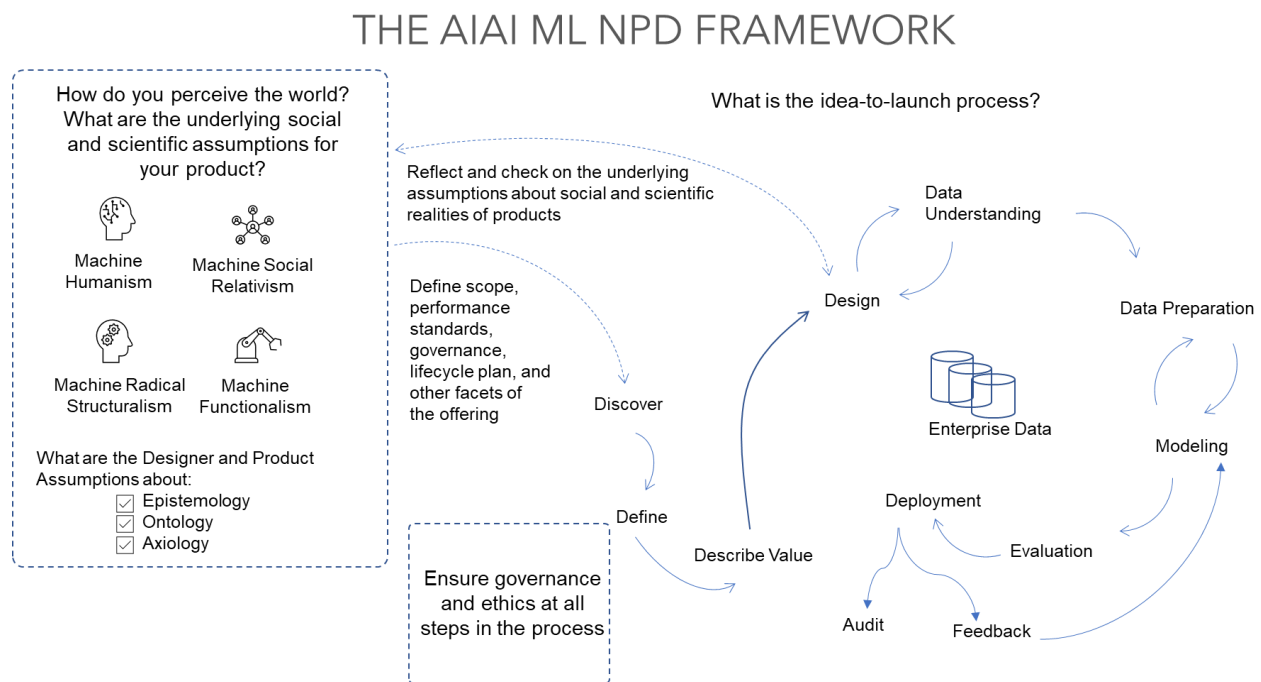


Figure 5.2 The AIAI ML NPD Framework

In the field test we discovered that ML is a different process than traditional or conventional product development in information technology (Figure 5.2). Before even one can conceptualize a specific solution, it helps to first understand what type of roles ML solutions can assume and what kind of problems are addressed by ML. In my research I attempted to explore these questions under a cycle which was later termed as *Discover*. I discovered that to clarify roles, it is helpful to begin by first asking why a firm wants an ML application. Specifically,

AIAI ML NPD identified three possible explanations: a) automate human work, b) improve human productivity, and c) increase human prediction capacity. To address the type of problems ML can solve and to generate new product ideas, we identified six problem areas including the problems of prediction, of thinking, of physical mobility and related cognitive processing, of social interaction, and of automated innovativeness. Notice that conventional IT does not address such problems in the manner that ML does. The key difference being that a learning system constantly learns, adapts, and improves its performance whereas a conventional system is static in terms of performance. Hence, identifying the problems of prediction, of thinking, of physical mobility and related cognitive processing, of social interaction, and of automate innovativeness, when approached from the backdrop of learning systems provides a powerful mechanism to not only solve the problem's current manifestation, but also its future manifestations. In other words, the solution evolves as the problem evolves. Every day that an autonomous car is out on the road, it learns new things. Every day that an algorithm conducts trades in financial markets, it becomes smarter. Articulating and knowing this helps to ideate product concepts. At AIAI, we used these to brainstorm examples of products that we can launch. Unlike conventional systems or the baseline NPD process, this research had identified a process to discover broad ideas for automation. For example, inquiring what human work can be replaced by machines enables one to brainstorm ideas for automation. Similarly, questioning what social interaction or mobility problems a firm can solve also helps to produce ideas for automation. Discovering problems that can be tackled by ML helped us narrow our options for launching a new product at AIAI. We discovered that we had the capacity to launch products in different areas, including in marketing, finance, supply chain and audit.

As we moved to the next step of defining the product, our goal was to narrow down the above options into a single product area that will be pursued by AIAI.

The research to uncover that showed what we termed as the *Define* cycle and which is facilitated by generating scope specific clarifying ideas for automation and by humanizing the systems. Generating specific ideas can be done by answering the questions shown in Table 5.1. The table was developed by creating a 3X2 from the two dimensions of Learn/Work/Social and Human/Machines. In this context *learn* simply means its regular usage in common language – such as one can learn about a human. For example, one can learn about the positive and negative habits of a person. Similarly, learning about a machine could be learning how to improve the performance of an HVAC system. Work/Human includes identifying what human work can be automated or improved by machines. And Work/Machine relationship looks at existing machines and determines if the work of an existing machine can be improved. An example of that would be taking an autonomous car and improving its features and functions. Social means how machine interacts with humans or other machines. Using these questions, we were able to identify that AIAI can learn new and interesting things about the management of firms under audit and can automate the work of an auditor who performs preaudit functions. Preadit functions include reviewing a firm's financial statements, looking at previous audits, and acquiring general knowledge about the auditee.

While answering these six questions can generate specific ideas for automation, they do not help to articulate the requirements of the system. To further clarify the user needs, we found that extracting requirements becomes much easier if we encourage people to talk in terms of humans rather than IT systems. This means to encourage them to answer humanizing questions – such as what type of work you would expect from a perfect human or what do you want your system to see and feel or what values you want to embed in your system – greatly helped us to define a product. No such mechanism exists in conventional systems because conventional systems are not designed to automate cognitive work of humans. Using this methodology, we zeroed in on the audit product. As AIAI clarified the

type of intelligent product the firm wanted to pursue, the next step was to start what turned out to be the *Design* cycle.

	HUMAN	MACHINES
LEARN	In a given problem domain, what can we learn about humans that we do not know?	In a given problem domain, what can we learn about machines what we do not know?
WORK	In a given problem domain, what human work can machines do or improve?	In a given problem domain, what machine work can machines do or improve?
SOCIAL	In a given problem domain, what human social interactions can be enabled by AI?	In a given problem domain, what machine social interactions can be enabled by AI?

Table 5-1 Generating Ideas for Automation

Designing an intelligent product required first understanding what *intelligence*, in terms of an intelligence product, is. We recognized that intelligence implies the ability of a product to resolve uncertainty in accordance with a goal. This implies that there will be grades of intelligence as simpler situations would require less intelligence to resolve whereas more complex situations will require greater intelligence. While intuitive, this distinction leads to different design choices. For example, a low uncertainty situation can be resolved by a simple physical or digital robot. More complex situations, such as making a reservation in a restaurant would require greater intelligence by the artifact. We recognized that knowing how intelligent we want our product to be is an important consideration. Knowing that helps to direct proper resources to the product and avoids over or under engineering. Connected to that, the second important consideration is how social we want our product to be. For example, do we want the product to

communicate with other intelligent machines or humans or we expect the product to be isolated. The third and important consideration is what type of governance is needed for the product. A high-risk product requires governance. Also, a product that is based upon the paradigmatic assumptions of social conflict requires governance. When combined, these three determinants provide a good design starting point for the product. To take it to the functional level, our research identified a methodology known as SALA. This methodology was originally developed at AIAI as a conceptual framework and our finding was based upon its actual application and testing in a real NPD process. SALA stands for Sense, Apply, Learn, and Act. An intelligent system performs these four functions. The design part involves clarifying answers to questions such as: what the product senses (data) from the environment in which it is expected to operate; what it would analyze, what type of decisions it will make; what type of actions are expected from the product; and what type of learning will be undertaken by the product. Answering these questions helps gauge functional requirements of the ML system. AIAI applied both methods discussed in this paragraph and was able to design the AIAI Audit product based upon that. Notice that the steps pointed out in the *Design* process are profoundly different than how the conventional systems are designed. The conventional systems design is not architected for learning and its interaction with the environment is limited and unchanging.

Once the design is completed, the next step was to clarify the business case and value of the product. Products that fail to create value for customers do not sell. Hence making a business case for customers was a critical part of our journey. We termed this cycle as *Describe Value*. The word implied determining the benefits of and the value creation by the product. The AIAI team objected to sharing any information outside the firm due to confidentiality and product secrecy reasons. To do that we identified various areas for value calculation. The first area identified the relationship between the artifact and its impact on

corporate earnings (revenues, costs, and resulting profits), risks, and growth. Using value drivers, in the second step, one can calculate and measure the profits, cost of capital, and growth factor. These measurements can help evaluate the return on investment. In 1956, Gordon and Shapiro, developed a financial model that uses profits (cash flows), cost of capital, and the growth rate (Gordon and Shapiro, 1956) as inputs and shows the valuation of an investment. In the research, the three elements of profits, cost of capital, and growth rate are strategically identified to be able to understand the value of investing in ML. They fit into the Gordon and Shapiro (1956) model. The research findings in this cycle did not materially depart from identifying value creation in conventional systems but there were some unique aspects. The two unique attributes are: a) risk assessment of not automating; and b) innovation acceleration. Innovation acceleration happens from using the automation technology as a scientist. In that role, the intelligent automation works to make new discoveries. These solutions are typically employed by research and development departments. For example, such ML systems played a major role in the discovery of vaccine for coronavirus (Kaushik and Raj, 2020). Conventional systems cannot act in this role or capacity and therefore no value of this kind can be associated with them. AIAI deployed the research learning from this cycle to develop a value creation measurement for the AIAI product. After measuring the returns, all the business steps are complete. The next step is when the actual development of the product is carried out and the engineering steps begin.

This next cycle was termed as *Development* cycle and here research findings suggest that the methods of developing conventional systems do not apply to ML systems and the data mining methods (such as CRISP-DM) are too limited. ML systems are often composed of multiple systems or agents, and they work in an integrated manner. In simple words, modern day ML systems are not standalone artifacts, they form a society. There are two implications of such a setup. First, companies need to build enterprise-based data capabilities. Second, designers

must have a mechanism to manage the interaction, individual and collective performance, optimization, and performance of the society of systems. No such requirements are deemed necessary for conventional systems or individual data mining projects. Thereby, the AIAI NPD Framework not only expanded the traditional engineering methodology, but also elevated it to become an enterprise based vs. just a standalone project based.

Post development, we explored the deployment part and realized that *Deployment* in ML is far more than the technical deployment. It requires metacognition management of the system and necessities monitoring changes in the environment that induce changes in the system. The system also provides audit information and feedback to improve its performance. Notice the arrow going from Feedback to Modeling in Figure 5.2. Since the system is also interacting with the organization, both organization and the intelligent system also impact each other. The learning system learns, adapts, and evolves. As such, unlike in conventional systems where once deployed and tested, the system is expected to be stable, ML systems require constant post deployment monitoring. Finally both ethics and governance were viewed as critical for machine learning products.

5.4 Reflections on the Social Paradigms of Machine Learning

With deployment, our findings for the process side of NPD were complete. However, this research was not just about process steps. My scholarly research also pointed out something that as a practitioner I was least expecting: Paradigms influence design considerations. This finding was about the epistemological, ethical, and ontological assumptions about the world. Our underlying perceptions about the world, our values, and our mental frameworks influence the choices we make in conceptualizing systems. Using the same data

and model, we can design a system that solves a customer's problem or that can exploit or manipulate the customer. Both are possible but it is up to us how we conceptualize a solution. This research could have been only about the process steps – but that will not be consistent with the values of our organization. The paradigmatic inquiry makes us aware of what would otherwise be unobservable choices.

To further formalize the paradigmatic inquiry, I have outlined the various considerations that may help determine what is your (a human's) mental model before you engage in conceptualizing a product. Your values and mental models are reflected in your designs. One way to determine what your mental models are is to ask yourself if you expect your system to function in a frictionless world and be fully autonomous or you expect conflict and tension. If you expect disorder (friction, tension, conflict) due to regulatory, ethical, or other concerns, you must embed governance into the product. If you expect a frictionless world, then you will only focus on the functionality of your product and ignore the governance part. For example, if your autonomous vehicle or a stock trading algorithm operates in a world with no regulations, concern for liability, ethical or legal constraints, then you may deem functional performance as sufficient. However, if you consider there is social conflict involved, then your considerations must take governance and ethics into consideration.

What about the existing ML products that a firm has? How can it determine the dominant paradigm under which it was developed? To test which paradigms our current or past choices of ML systems represent, one can use Table 5.2. The columns headings in Table 5.2. show the four quadrants of the paradigm model developed in Section 5 of Chapter 2. The rows show the actual system requirements as captured by designers. Once the designer inserts what type of autonomy, social interaction, data, change, and governance aspects are part of the solution, one can point to the appropriate paradigm. For example, a high

autonomy, high interaction, inflexible data, and well governed system indicates that it will fall under Machine Social Relativism. The reason this system will have inflexibility of data and low governance is because the system design assumes no conflict.

	Machine Functionalism	Machine Radical Structuralism	Machine Social Relativism	Machine Humanism
Autonomy in a System	High autonomy	Low autonomy	High autonomy	Low autonomy
Social Interaction	Can work with other human and machines; Limited interaction	Can work with other human and machines; Significant interaction	Works with other human and machines; Significant interaction	Works with other human and machines; Significant interaction
Data of the System	Datasets are fixed	Datasets may not be fixed	Datasets are fixed	Datasets may not be fixed
Change in System	Inflexible	Flexible	Inflexible	Flexible
ML Governance	Low or no governance	Moderate governance	Moderate governance	High governance

Table 5-2 Identifying Paradigms in Existing Systems

Applying this part in AIAI, we recognized that our motivation for launching a new product was based upon Machine Radical Structuralism. We were opting for an autonomous system that assumes that there will be disorder in the environment. Disorder could imply regulatory issues, ethical issues, or organizational conflict. When one uses the Machine Radical Structuralism approach, one recognizes that the solution must be backed by a strong governance framework and recognize social and ethical issues. Note that developing the governance framework is not part of the paradigmatic inquiry. Paradigmatic inquiry only points out the need

for such a step. The governance criteria are developed in the NPD framework (bottom part of Figure 5.1).

5.5 Contribution to Knowledge

As discussed in the Literature Review, leading voices in the industry were concerned about the dearth of methodologies that deal with product development in ML. When scholars pointed out recently that ML systems have all of the problems of non-ML software systems plus an additional set of their own specific issues (Wan et al., 2019) and that AI domain is not like prior software development (Amershi et al., 2019), an urgent need for a methodology became obvious. The existing dominant methodology, from which other all other methodologies were derived, was deemed insufficient to embrace the new challenges of modern AI and data science (Martinez-Plumed et al., 2019). The limitations of the existing methodologies were pointed out in recent literature (Rollins, 2015; Piatetsky, 2014; Marbán et al., 2009; Studer et al., 2020). On one hand we observe a dearth of methodologies, on the other hand we have a surge in the launch of ML based data driven products (Li et al., 2019; Hesenius et al., 2019) – creating a perfect storm leading to a high failure rate.

Notice that all the above-mentioned researchers were only referring to the engineering side of the NPD. Add to that the business and social dimensions and we observe an extremely troubling absence of guidance on how to develop new ML products – and therefore a need for a framework.

The AIAI NPD Framework is one such model to fill the vacuum. To my knowledge, this is the first model that not only establishes an enterprise centric methodology for the ML engineering side but also includes the two dimensions of business and social considerations. This makes the AIAI NPD Framework the first in the industry for ML that is inclusive and in line with the depth and comprehensiveness of models previously developed in conventional information technology (they were extensively covered in Literature Review Chapter 2). It is

expected that the combination of ML engineering, social (managerial), and business dimensions makes the model unique and better equipped to handle risks in NPD.

A generic NPD process is composed of two areas of Concept Design Subprocess (*what should a firm do – social and business aspects*) and a Product Development Subprocess (*how should a firm do it – engineering aspects*) (Wang, 2016; Shepherd and Ahmed, 2000). The AIAI NPD Framework addresses both areas. For example, the first four cycles (Discover, Define, Design, and Discuss) covered the social and business areas inherent in an NPD methodology. The last two cycles (Develop, Deploy) covered the engineering part. When stitched together the methodology provides a comprehensive and exhaustive view of undertaking NPD in ML.

On the engineering side, the AIAI NPD Framework elevates the engineering problem from a firm trying to engineer a single ML product to building an enterprise-wide capability. That is why when the *Development* findings call for establishing a data management, data quality, data governance, and data preprocessing capabilities at the company level. The findings propose an enterprise vision where ML will become pervasive, and more than one product will be simultaneously developed. Existing frameworks and methodologies only focus on a single product engineering. The enterprise aspect of the AIAI NPD Framework is further magnified when it is observed that the model proposes a society of machines (agents) working together and achieving goals collaboratively. These machines need to be managed and hence the metacognition and metamodel elements form the critical process components in the *Develop* and *Deploy* cycles. To my knowledge, the existing methodologies do not address either building enterprise capabilities as part of the ML development or envision a system of multiple agents trying to achieve collaborative work goals.

On the business side of NPD, we know that NPD requires planning (Salomo et al., 2007), best practices NPD (Kahn et al., 2006) and managerial sensemaking (Christiansen and Varnes, 2009). Along with a broader framework for the engineering methodology, the AIAI NPD Framework was developed to provide planning, best practices, and managerial sensemaking. The AIAI NPD Framework maintains the process dimensions of a typical NPD (Shepherd and Ahmed, 2000; Grönlund et al., 2010). However, the process is built upon social and scientific reality assessment – known as paradigms.

The *Discover* process highlights that knowing the various roles of ML and the types of problems ML can solve can help strategists and managers drive ideas for automation. The *Define* cycle findings showed that funneling and filtering the concepts identified in the *Discover* cycle can be accomplished by humanizing (anthropomorphizing) the solution requirements. No existing methodology, in my knowledge, approaches envisioning a solution from that angle. In the conventional software development, agile methodology introduced new approaches to requirements analysis that included greater human interaction (Darwish and Megahed, 2016). A recent literature review of agile requirements engineering in conventional systems calls for greater accommodation for Human Centered Design (HCD) and User Centered Design (UCD) (Schön et al., 2017). Some recommend using ethnographic analysis (Meligy et al., 2018; Surendra, 2008). In practice, goal-oriented requirements analysis (GORA) is also applied (ElSayed et al., 2017; Kinoshita et al., 2017). In conventional systems, feelings, emotions, values, and motivations were integrated into the requirements analysis framework and was termed as value-based analysis (Thew and Sutcliffe, 2018). But the difference between all of the above conventional methods and the AIAI model is that we seek attributes such as feelings, values, motivations, and goals not just for the human user but also for the machine itself. Machine, therefore, becomes the goal-oriented entity with feelings, emotions, cognitive structures, likes and dislikes, and interactions with other intelligent entities.

The *Design* stage calls for using the SALA method and that itself is a unique concept. In my knowledge, no existing methodology outside of AIAI has outlined a product planning framework using the SALA model. The *Discuss* stage is similar to the existing approaches – that is to calculate the financial benefits of a solution – however, it does highlight that ML can not only automate work, it can also accelerate innovation by creating automated scientists. Even those methodologies that do include business case building via financial analysis do not consider the profoundly powerful dynamics of the ML economy where firms can automate and accelerate the pace of innovation itself. For instance, the use of ML was pervasive in discovering the vaccine for Covid19 (Kaushik and Raj, 2020). Measuring acceleration in innovation as a financial measure is unique to our model.

Central to the AIAI NPD Framework is the concept of governance. This concept is reflected both on the business side and the engineering side in the model. The governance and social consciousness of the AIAI NPD Framework was patterned after the widely used and revolutionary model of social paradigms – originally envisioned by Burrell and Morgan (Burrell and Morgan, 1979) and later applied by Hirschheim and Klein (Hirschheim and Klein, 1989) in information systems.

Below I discuss the contribution of the AIAI NPD Framework in terms of the social paradigms. I divide the coverage into the two areas of the ontological insights and the epistemological insights.

5.6 The Applied Examples of the Framework

This is an applied example for the AIAI NPD Framework for ML. In this example, suppose we are trying to conceptualize an autonomous car – such as Tesla (Figure 5.3). As per the model, we will first distinguish between the four types of product paradigms that will determine not only the various manifestations of the product but also will influence changes in the subsequent steps in the NPD process.

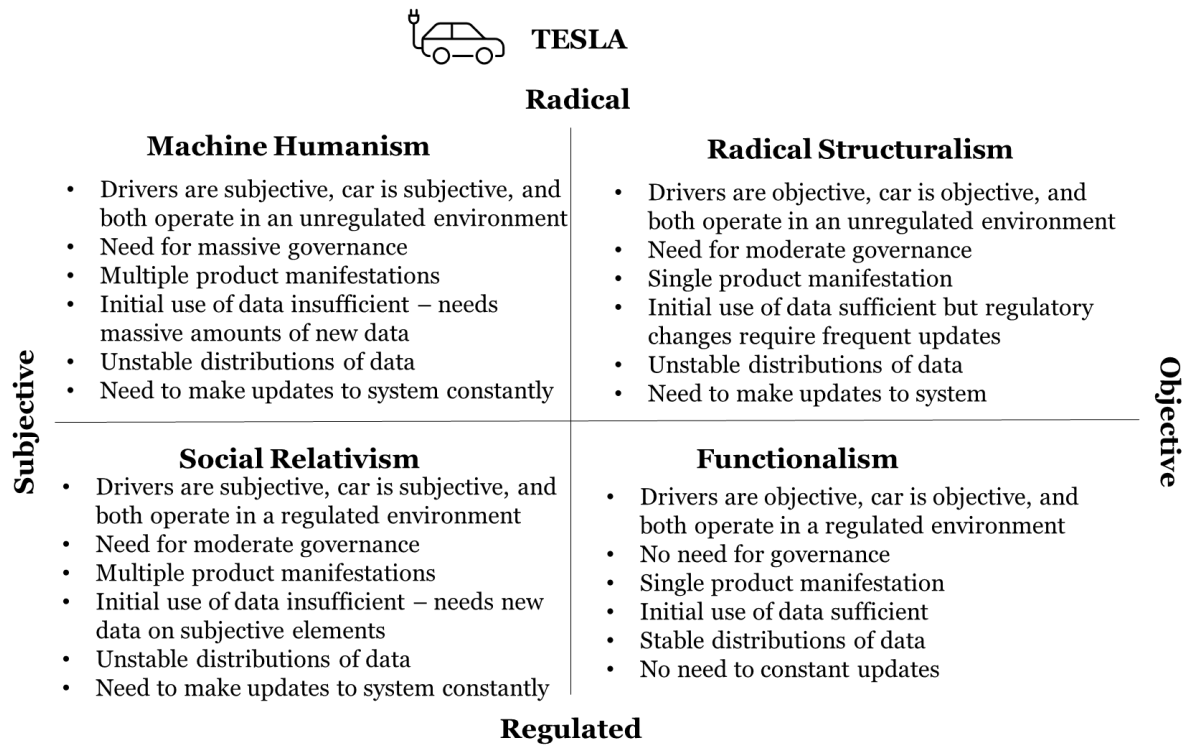


Figure 5.3 STEP 1 Analysis for Tesla

Machine Functionalism: An autonomous car conceptualized under the functionalism paradigm will be based upon objective reality and regulated assumptions. It means that the product will be envisioned based upon assumptions that both the human designer and user approach reality as objective – and the machine itself approaches reality objectively.

Furthermore, a regulated assumption implies that the product must be made compliant with laws and regulations. Since compliance with laws and regulations are embedded within the designs, the governance will be less desirable. Product has standard features and follows established legal and regulatory norms. Here are the six NPD steps of the product will be based upon the functionalism paradigm. For example, at the concept stage, the car will be expected to represent objectively determined needs of customers and as an autonomous operator will abide by all regulations. A customer

may ask the car to speed over the speed limit, but the car will not accept that directive. At the define stage, the car will be defined in terms of standard functions and features. At design stage, stable and fixed datasets will be used to train the algorithms, and the models and algorithms will stay stable. The value creation can be objectively assessed, and customers will be able to produce concrete evidence of value creation. The development process will require training the vehicle on stable datasets. Once trained the models will be deemed sufficient to represent the entire problem domain. A car operating in New York will be the same as one operating in California. New learning capacity shall not be deemed necessary during the product lifecycle. Deploying the car will require less governance since the design will embed regulatory compliance and legal constraints. Software updates to the autonomous driving features will be rare or nonexistent. If Tesla was designing the car under this paradigm, it will not anticipate major changes in legal/regulatory environment and will expect a certain level of objectivity from humans and will embed machine learning in a manner where the car will expect the human to operate the vehicle objectively. There will be no need to place automated controls to stop a car if a human driver would want the car to jump off a cliff.

Machine Radical Structuralism: In this paradigm humans and machines function objectively in a radical or unregulated environment. The scientific reality in this paradigm is stable and objectively determinable – hence product features, functions, and customer expectations will be stable and standard. The social reality is unstable, implying conflict or lack of regulation. This could mean that dominant influences will be active in the society. Keeping track of those influences, social conflict, and instability will be critical. The product design will change as law and regulation develops.

This car will require moderate governance because while drivers are expected to be objective; the instability is introduced from the system. This could also mean that the designers will expect that the driving rules will be different, rapidly changing, or evolving. In this case, standard learning will not be sufficient. The designers would be expected to continue to improve learning of the car as the environment changes. This means that the NPD process will be dynamic. The product will continue to change and evolve. On the design side, the car will represent standard functions and features with less need for customization. The underlying assumption that people will have needs and wants that can be objectively determined will ensure a steadier design about factors such as seating, comfort, and other owner preferences. However, the unstable regulatory environment will imply that the car's autonomous driving software will need to be updated constantly with new data. If Tesla was designing its autonomous car under this assumption, the firm will recognize that constant updates will be made to the car as regulation related to autonomous cars will develop. Unless required by law, Tesla will be unlikely to place controls in the car to stop the car from jumping off the cliff if a human operator wants it to jump.

Machine Social Relativism: In this paradigm, both humans and machines are expected to operate subjectively, however they exist in a regulated society. Controls, comfort setting, decision-making on road, may mimic the owner's preferences and states. However, all such factors follow established regulatory norms. The scientific reality of entities is modeled by multiple factors and is dependent upon the vantage point of the observer. Vantage point can be based upon a multitude of shifting factors that are internal to the entity (moods, safety concerns, political preferences, etc.). The external environment of the entity is deemed stable and regulated. This

car will require significantly larger datasets to model multiple factors associated with human behavior, moods, and preferences. This will need strong governance. For example, a person may want the car to drive extremely fast, but the regulatory assumption will override human preference. Unlike functionalism, in radical structuralism, the car will view human riders as having subjective realities. Similarly, designers will attribute humans as subjective. Without objective criteria, during the discovery and defining process, the features of the car will not be based upon objective reality calculated and measured by closed ended surveys. Instead, it will be developed by using ethnographic methods to develop a broad profile of customer needs and moods. The car itself will make decisions based upon human moods – for example, depending upon a person's mood may change the route on which it drives a passenger and take a more scenic but longer route if it senses the passenger in a romantic mood. If Tesla was designing a car under this assumption, it will make sure that human passengers' and drivers' habits and behaviors – including driving behaviors are modeled and tracked. The car will study human behavior. In this case, Tesla will expect that a human will not drive the car off the cliff and will likely expect rules and regulations prohibiting a human from doing that. In cases where such rules and regulations do not exist, it may simply use human behavior to discourage a person from doing that.

Machine Humanism: In this paradigm machines and humans operate under the subjectivity assumption in an unregulated society. This would be the hardest to model situation in designing an autonomous car. This assumes that human reality is subjective – and hence not entirely determinable and subject to change due to moods and emotions – all while the environment in which the car operates has major instability related to

governance and regulations. This means extremely large datasets will be needed to train the autonomous driving features. Human operator and his or her environment will be deemed unstable. At the discovery stage, it will be hard to determine the needs and wants and ethnographic studies will be used. The design will remain variable and constantly evolving. The product design will stay in a perpetual flux. The regulatory structures may change rapidly, or the new legal structures can develop without notice. In this case the car will require immense governance. If Tesla was designing a car under this assumption, the firm will ensure it can model human behavior and also all the environmental (regulatory, legal and other) factors. The autonomous features of the car will require constant updating. If a customer wants to jump the car off the cliff, the car will stop functioning and will not allow that to happen. Hence, the intensity of governance will be viewed as inversely proportional to the combined values of objectivity and regulation.

The above discussion was meant to introduce the NPD Framework's application in a real situation. The framework can be applied in nearly all scenarios. For example, a stock trading autonomous software could have four different manifestations based upon the underlying paradigmatic assumptions – and such a determination will necessarily invoke different assumptions in the subsequent NPD process. For those readers who want to explore the underlying philosophical assumptions of the framework, the discussion below will be helpful.

5.7 The Ontological Insights

The Ivori (1991) model included ontological assumptions about data, information system, organization, and humans (Ivori, 1991). In the AIAI ML social paradigm, these ontological assumptions are important considerations.

Data in Machine Learning: The AIAI model also challenges the existing distinction of data in terms of realism and nominalism. Ivori (1991) explained that the data and information are composed of both descriptive facts and constitutive meanings (Ivori, 1991). It is important to recognize that Ivori was using the term “data” to connote data as representing a fact or a socially constructed meaning. Ivori acknowledged that data and information classifications under *ontology* exposed the risk of overlapping them with epistemological inquiry. The factual and constitutive distinction implies that there are two types of data, one that conforms with the facts (for example as specified in the natural sciences), for instance the temperature at a certain location at a certain time, and the second that is imaginary or human perception based. Hirschheim and Klein (1989) four paradigms assumptions were derived from their former article in which they clarified that a data model is the development of a schema that formally represents a “Universe of Discourse” (UoD) (Klein and Hirschheim, 1987b). A paradigm consists of at least two sets of assumptions: ontological and epistemological – and the paradigm of data modeling was based upon two basic ontological positions of realism and nominalism. Realism, they argued, postulates that UoD comprises of objective immutable objects and structures that exist as empirical entities independent of the observers’ application of them. Nominalism, in contrast, is where reality is a subjective construction of mind, and socially determined names and symbols create the perception and structure of reality. It is only experienced in observers’ appreciation. The AIAI NPD Framework challenges that assumption of data and

shows that such a contrast for machine learning is not necessary. Data is understood to be observations of real-world phenomena. As covered in the *Develop* cycle, data provides a little window into the reality. It gives a small dot which is used to connect the dots. If an element of data is provided as a single value, it may not mean much for a machine. The machine does not care what kind of reality is expressed by the data. Since the machine is not reasoning with the data, instead, it is using it to learn via a mathematical model, the elements of human subjectivity or objectivity are not relevant. In fact, data, from a machine learning viewpoint can be viewed as consisting of simultaneous existence in multiple forms of realities – for instance, objective, subjective, virtual reality, augmented reality. The value of temperature at a certain time and place can be used to learn about the weather conditions of a geography, or about the variation of moods of a human with weather, or to predict the performance of an engine for a truck, or to assess the romantic appeal of one human for another. From a machine perspective, the data is neutral and simultaneously belonging to various realities. This insight is unique to the AIAI NPD Framework.

Information System: Iivari (1991) claims that the classification of information systems is based upon the Goldkuhl and Lyytinen's (Goldkuhl and Lyytinen, 1982) concept that IS can be viewed as “technical systems with social implications” or “social systems only technically implemented” (Iivari, 1991). In that perspective, Iivari points out, the mechanistic view of information system calls for it being viewed as a tool or an artifact, while the institutional view demands it to be seen as a social system. ML introduces the presence of digital worker or machine intelligence in a typical human centric organization (Brynjolfsson and McAfee, 2015; McAfee and Brynjolfsson, 2016). As shown in my research in *Define*, *Design*, and *Develop* cycles, when machines make decisions – including decisions about promotions, hiring, firing, financial, legal, and others – they are no longer merely the providers of information. They have turned into decision-makers and are part of the cognitive and social structure of a

firm (Wright et al., 2017). This is a major change and in accordance with the reasoning provided in this research, its implication is that every ML system is a social system belonging to a society composed of other machines and humans. This recognition is not prevalent in conventional NPD models. As long as it is recognized that intelligent machines are now contributing parts of the social structure, and not just artifacts, the ontological distinction between *technical* and *social systems* can be maintained in ML. The difference between the conventional and the AIAI model being that in the AIAI NPD Framework the social system itself is enhanced to include intelligent machines as part of that structure.

Human: The AIAI model alters the debate about human determinism and voluntarism. The distinction between human deterministic and voluntarism ontological assumption is maintained both in Burrell and Morgan (Burrell and Morgan, 1979) as well as in Hirschheim and Klein (Hirschheim and Klein, 1989). Prior to analyzing this distinction, it is important to observe that this distinction itself points to the claim made in the literature review that the Burrell and Morgan's and Hirschheim and Klein's four paradigm models were indeed based upon human centered realities – since both *deterministic* and *voluntarism* apply to human. The deterministic view implies that humans activities are determined by their circumstances and voluntarism implies that humans shape their circumstances (Iivari, 1991). As the findings from the *Define* and *Design* cycles suggest, the inclusion of intelligent machines in human societies implies that human actions and their outcomes can be affected and shaped by the machines that are part of their environment. For example, machines can develop precise cognitive profiles of people and can understand their preference structures to a point that they can influence people's choices (Kaya and Salah, 2018; Mehta et al., 2019). Technology determinism, expressed by Ellul (1965), argued that technology requires its own techniques and processes and those techniques may not be shaped by human choices as humans would have to follow those

techniques and best practices to run the technology (Ellul, 1965). Ellul was referring to conventional systems and therefore his idea of technology determinism was limited. As our findings show (in the first four cycles) that autonomous machines perform functions independently and autonomously, the concept of technology determinism must be evolved to incorporate the reality of ML centric machines. Accordingly, technology determinism no longer only means that human has become a prisoner of technology user manual, but also that technology can function and make intelligent decisions independent of humans.

Organization: The AIAI model recognizes organizational dimension but introduces the concept that decisions made by machines affect organizations and organizations (humans) affect machines. Human decisions are a source of data and machines absorb that data to learn. But autonomous machines are making decisions independently – and their decisions affect organizations. Therefore, the new definition of organization must include intelligent machines as participants in the constructs of organizations. The organizational dimension is the ontological split between realism and nominalism (Iivari, 1991). Realism implies a world exists outside the perception of the observer and that world is tangible and has immutable structures. Social world, like natural world, is real, concrete, and hard. Nominalism implies that the social world results from the appreciation and perception of the observer and it is a social construction. ML plays a major role in integrating both concepts – since a machine does not have to make that discrimination between realism and nominalism. Machine is viewing the world from the nomothetic angle and its perception is shaped by data (Mitchell, 1997; Russell and Norvig, 1994). As the research findings show in the *Design* cycle, the data can be about human emotions and feelings, or it can be about what natural sciences define as facts – for the machine it does not make a difference. Machine, when looking at an organization, assesses the organization not only in terms of processes and business or organizational structures, but can also recognize the deep social, political, and psychological structures that exist in the organizations

(Straub et al., 2016; Fire and Puzis, 2016; Boxwala et al., 2011). Hence, the AIAI NPD Framework argues that a machine's perception can be far broader and deeper about an organization than the human perception.

5.7 The Epistemological Insights

The epistemological assumptions of both Burrell & Morgan (1979) and Hirschheim and Klein (1989) were consistent with the human centered scientific exploration in an age when data was sought as a result of the inquiry. In ML however, inquiry can result from the inevitable link between ML data and ML algorithms. This reversal of the scientific process was pointed out in the *Discover* cycle. As we move from functionalism to social relativism, the application of the scientific process is not changed, even though subjectivity of meaning, sensemaking, metaphors, and symbols is acknowledged. The shift from singular objective truth to accepting the possibility of a wide spectrum of reality is embraced. Despite the change in the ontological assumptions, the fact that reality – perceived as actual or constructed – is assumed as knowledge stays intact in the traditional models. Data gathering follows research design and model selection. The objectivity assumption in radical structuralism also dictates the application of the scientific process – whether approached as positivist or anti-positivist – the data gathering, and analysis are deemed essential to understand the nature of conflict, class interests, and other economic and social considerations. The neohumanism also requires data related to human emancipation, liberty, and freedom from natural and social constraints. Burrell and Morgan (1979), as well as Hirschheim and Klein (1989), never consider a world where machines can identify problems, deploy models, select variables, and find optimum solutions – all from data. Data was no longer an afterthought; it became the only reality. Positivist approaches and antipositivist approaches are both human centered cognitive structures. However, the AIAI model challenges that long held assumption. Gregory Wheeler of Munich Center for Mathematical Philosophy explains the nature of the problem in his article Machine Epistemology and Big

Data (Wheeler, 2016). He clarifies that statistics addresses two questions of: 1) what can be inferred from data, given the modeling assumptions of the researcher; and 2) reliability of those inference. He explains that since data is acquired through deliberation and action, the assessment of reliability follows the choice of deliberation and action. Prior knowledge, he argues, becomes a tool for new inquires rather than a normative standard against which new conclusions are evaluated. However, the long tradition in epistemology from the philosophy side acknowledges the objects of knowledge as prior to, and unchanged by, the cognitive activity to know them. As such, statistics, he illuminates, becomes a puzzle for traditional epistemology as it fails to answer on what grounds you select a model and what justifies certain data being used to model an uncertain event. Wheeler says that this is the wrong way to think about things. Uncertainty needs to be explored and not ignored, and knowledge must become a means of control and not a state of mind, and inquiry must not determine the epistemic notions but instead they should be derived from the role they play in the inquiry. He cited Hilary Putman (Putnam, 2004) who demanded pragmatic enlightenment in epistemology and Dewey (Dewey, 1929) who observed that uncertainty is a practical matter. Wheeler argued that ML, like statistics, is also interested in finding out what can be inferred from data – but unlike statistics, ML tries to circumvent setting of explicit modeling assumptions before drawing inferences and can also deploy algorithms that can learn on their own which modeling assumptions to select (Wheeler, 2016). Yet if modeling a problem, experimenter's choice of which features he or she selects or deems important will necessarily require background knowledge and assumptions, and this absence of general principles has been a source of criticism by philosophers (Whitehead, 1925). However, as discovered in my research from the *Discover, Define, and Design* cycles, the relevant and informative features in ML can be determined without any prior knowledge and this can unleash new discovery and new knowledge, above and beyond the experimenter's experience or assumptions of

inquiry. The AIAI NPD Framework encapsulates this unique aspect which leads to automated discovery where a machine can learn hidden structures as substitutes for background knowledge.

5.8 Application in Practice

In this section I will only discuss some of the potential applications of the findings in practice. In the next chapter I show the limitations of the applications in practice. ML has now become a primary source of competitive advantage in business (Naqvi and Munoz, 2020). It holds transformative power (Makridakis, 2017; Wright and Schultz, 2018). Despite such potential, data reported about the technology adoption of artificial intelligence shows a failure rate of 75 to 85% (Nimdzi Insights Pactera EDGE, 2019). With billions of dollars being invested in AI/ML (Shoham et al., 2018), the investment risk remains high. I showed in the literature findings that the failure rate can be attributed to the dearth of an appropriate methodologies to develop new ML products. Accordingly, a company seeking to build AI and ML products is left with no good options. If it does not develop new products, it will face competitive wipeout. If it does, it faces a failure rate where 3 out of 4 products may not achieve their goals.

It is expected that the AIAI NPD Framework may help overcome some of the problems identified in the previous paragraph. As illustrated in the US government example in Chapter 4, many firms are unclear about what to automate. Notice this question is different than “how” to automate. These firms are seeking guidance on ideation, conceptualization, and envisioning of ML solutions. That is why the learning from the Discover cycle is useful for the industry. The initial ideation begins when a practitioner applies the AIAI NPD Framework to conceptualize opportunities by analyzing the three goals of automation: automate human work, increase productivity, and increase predictive ability. Adding to that the six problems of prediction – of thinking, of physical mobility and related cognitive

processing, of social interaction, and of automated innovativeness – practitioners can brainstorm and ideate thousands of ideas across a firm.

But not having ideas for automation is not the only limitation for firms. Inability to shift to the ML world semantically and cognitively from conventional non-intelligent technology is also a major barrier. Working downwards in a funnel like manner in the AIAI NPD Framework, practitioners can synthesize ideas generated into specific concepts for automation. They do that by humanizing the system and asking questions that apply to human workers. This makes it easier to ideate at a lower level. The humanlike qualities of intelligent systems enable us to anthropomorphize the system. This exercise invigorates the system's conceptualization. Once that problem is solved, practitioners can jump into the *Design* part.

Without a relevant methodology, there are two major roadblocks to designing ML products. First, practitioners are approaching ML on a project-by-project basis and not on an enterprise level. As discussed above, not having an enterprise perspective implies that companies will have to reinvent the wheel every time they develop a new product. Many processes can be formalized as capability areas and departments in companies. The AIAI NPD Framework addresses and points out those capabilities such as data management, data governance, and data quality. This is also the model's greatest limitation (discussed in the next chapter).

Second, practitioners need a way to develop lower-level expectations or requirements for the system. The conventional methods do not work since they were not developed for learning systems. The AIAI NPD Framework can help construct the requirements of an AI system. It is known as the SALA process. Filling out the SALA requirements can enable practitioners, regardless of their background, to intuitively think about how to design an intelligent system. Identifying what the system will sense, apply, learn, and act (SALA) provides an intuitive way to design intelligent products.

Since to create a business competitive advantage, it is likely that practitioners will be tasked with deploying not one but many intelligent systems, they would be expected to view these systems as a network or society of intelligent agents. The AIAI NPD encourages this line of thinking and through a series of questions, enables practitioners to develop that perspective. This perspective – once translated into vision, strategy, and execution – creates interconnected set of capabilities. An example of that will be Amazon, a company in which multiple AI systems function as an interdependent and interconnected network of capabilities. Through collaboration of multiple intelligent systems, Amazon can simultaneously make recommendations, support orders and fulfillment with automated robots in their warehouses, and manage shipping and distribution (Levy, 2018). AIAI NPD Framework is designed to support such collaboration.

The *Discuss* cycle, which focused on value estimation provides practitioners with guidance on how to measure value creation from ML. In addition to reminding practitioners that automation can create profits, reduce risks, and create new growth opportunities, the AIAI NPD Framework also shows that innovation acceleration can be a key attribute of ML development. Innovation acceleration happens when ML is applied in Research & Development functions of companies. This also points to the fact that with AIAI NPD Framework ing practitioners will be inclined to think about automation in terms of the entire firm and across all departments and not in the limited sense of a single solution or a use case. The AIAI NPD Framework elevates all aspects of NPD to an enterprise level.

The AIAI model can also provide guidance for ML development and deployment. The development and deployment capabilities in the AIAI NPD Framework are not specific to a single project or use case. They are driven by having an enterprise viewpoint. Enterprise factors require setting up formal corporate structures to support data related activities and they demand incorporating metamodel and metacognition in deploying a society of intelligent artifacts. Capabilities will be

needed to govern, regulate, and manage the behaviors of the participants in the society of agents. AIAI NPD Framework informs practitioners that such an enterprise thinking will be critical for enterprise-based NPD. The AIAI model does not stop at deployment as it captures the concept that ML systems are learning systems and they may possess evolutionary or adaptive capabilities. Tracking the environment and the associated changes in the society of artifacts will be critical.

The model goes beyond the NPD part and incorporates a way for practitioners to decipher the social paradigms under which they are developing systems. This implies including ethics and governance as critical parts of the model and not afterthoughts or extras.

Despite such advantages, the AIAI NPD Framework is just that, a model. It has limited testing. In the Conclusion chapter I discuss the limitations of the model. In addition, I will provide a summary of findings and give suggestions for future research.

5.9 Sensemaking

Alluding to intelligent machines as social systems is a big claim that requires rethinking machines as part of the organizations. My research had not covered all the rich dynamics that are associated when an intelligent entity acquires a social role. Focusing on that would have been out of scope for this work, however, some critical insights are necessary. One can approach this issue as what would constitute as machine's sensemaking? Also, how would such a change impact human sensemaking? I referred to Weick's seven core ideas (Weick, 1995) and reconstructed them to think about machines. After all, humans will be making sense about machines and machines will be making sense about humans and other machines. For each of the seven core ideas of Weick (1995), I can think of the following questions as they relate to ML systems. For *identity* in Weick's model, what would machines' role be with respect to their environment? How would the

introduction of intelligent machines lead to identity change in firms? For *retrospective*, what meta patterns would develop in machine and human coexistence and co-work situations? How would those patterns be codified into organizational memory? For *Social*: How will machines redefine socialization? Would they force humans to develop a singular personality? Would it allow humans to develop uniqueness? When you author an article on blogs, machines provide feedback to authors. Does it mean authors will develop singular voice? The feedback is based upon the popularity of previous likes and dislikes. Thus, humans will experience massive uniformity. For *Ongoing*: Machines that learn and adapt will have redevelop understanding of the world as the world around them changes. This will happen from data, data distributions, and features. For *Extracted Cues*: Cues extracted from sense and perception lead to cognition and are articulated in writing and speaking. Machines are very good in analyzing speech patterns and are now beginning to understand contexts. How would that change the human expression? Would it make people more careful about what they write or say? Would it lead to social despotism? In this regard humans are making sense of being analyzed by machines. For *Plausibility and Sufficiency*: The problem with plausibility and sufficiency is that machines can make sense out of things that humans ignore or cannot process. On the other hand, humans are good at making sense even with little information – but machines are catching up. This means that at some point humans will have to rely upon machines for sensemaking far more than they rely upon their own senses.

These questions are not in the scope of this research, but they are critical issues of our times that deal with emancipation, justice, and human rights. These are the type of issues for which AR is perfect. AR exposes such problems.

5.10 Summary

In summary, the AIAI NPD Framework is proposed to be a comprehensive methodology for launching new machine learning products. The framework was

discovered through an action research conducted in one firm. Since the model has not been applied in other firms, no claims can be made for its generalization. However, several areas of the model are intuitive, and practitioners can apply them to improve their ML NPD processes. In the next chapter I will elaborate the limitations of the model.

Chapter 6 Conclusion

6.1 Introduction

The fundamental question addressed by the study was: what is a new product development (NPD) framework for designing and developing machine learning products? NPD is not a new field. The above question has been addressed for conventional technologies. What made this study unique was its focus on Machine Learning (ML). The resulting methodology from the research suggests that the NPD process for ML is materially different than for other technologies and product/services. The difference comes from the fact that ML builds intelligent machines and prior to the advent of AI humankind has never built synthetic intelligence. As this goal manifests in business processes needed to launch successful intelligent products, it requires looking at both business and engineering domains. A positivist centric study could have approached that by hypothesizing each sub-element of an NPD model and testing them for relevance, efficacy, and adaptability. However, this research approached the problem as an action research project. As AIAI developed a new product, an associated NPD model took shape. Some parts of the model were researched and preceded their application in the AIAI product development. Other parts were abstracted and extracted from the decisions and actions made during the product development. Theory guided action and action guided theory such that both AIAI product and the AIAI NPD Framework fomented and solidified into existence almost simultaneously. Both tacit and explicit knowledge and learning was embedded in a framework that metamorphosed the generic baseline methodology and reoriented it for a ML centric economy, business processes, and innovation. As the AIAI NPD Framework emerged as a model that can be applied in practice, it exhibited some unique features.

The AIAI NPD Framework follows the same sequence of activities as any NPD model – the six phases of discovery, definition, designing, discussing, developing, and deploying. But what happens within each of these categories – the process steps, the approaches, the concepts, the articulation of requirements, and the tasks – changes substantially. What also changes are the corporate capabilities. The AIAI Model proposes enterprise level changes. While not uncommon in NPD models seeking enterprise level coherence to drive innovation, what is different are the specific modifications the AIAI model calls for. AIAI NPD Framework suggests building separate focus areas and departments at an enterprise level (such as data governance, data preprocessing). The AIAI model also departs from the conventional models by introducing capability building that accelerates innovation via automating innovation itself and altering the scientific process (such as creating autonomous machines as scientists). More importantly, the AIAI model establishes a balance between the forces of engineering, business, and social/ethical by configuring them as the necessary interdependent processes.

While the business and engineering sides of the NPD exhibit explicit changes – the social dimensions of information systems development are approached from the unique angle of social paradigms. Patterned after Burrell & Morgan (Burrell and Morgan, 1979) model, the AIAI model explores the rationale behind systems development. Systems can be developed for functional purposes only. Systems can also be created to give power to the weak. ML can also be viewed as enforcing order. And it can be a source of social change and emancipation. The social dimension of NPD is a unique addition because it does not start the vision-setting exercise of a system at the discovery point where firms reach out to customers to understand problems. It creates a preliminary prerequisite step that forces management to think about the social role of the system they are creating. Would the system derive power from its functionalism – i.e. it works and that is all that is needed? Or would it empower the weak and uphold justice? These debates are

necessary in the world where ML has been used to manipulate elections and referendums, exploit customers, steal identities, perform cyberattacks, conduct espionage, create social conflicts, and concentrate power in the hands of few. This necessary prerequisite step in the NPD process introduces ethics and governance as integral parts of NPD. Given the explosive power of AI and ML technologies when combined with the possibilities of what manufactured synthetic intelligence can do, the inclusion of such considerations becomes necessary.

The AIAI NPD Framework gives practitioners a way to approach both capability building and product development in a single model. It guides them to incorporate best practices specific to ML but also challenges them to instill enterprise-wide capabilities that can make NPD as part of their operational DNA. AIAI was successful in launching its product and hence one can assume that the methodology worked for AIAI.

6.2 Limitations

While the novelty and uniqueness of this study, as discussed above, make it relevant for modern day practice environments trying to develop ML products, there are many limitations.

The obvious limitation is that the AIAI model was developed from the experience of a single firm. While the model did include some prior work of AIAI that was used to train other firms, and utilized literature review to develop a conceptual framework, the application in AIAI offers extremely limited evidence of its application. The feedback provided by customers was based upon the structure and concepts of the methodology and not in them applying the model.

Secondly, the model calls for top-down change that requires building many corporate capabilities. As pointed out by two customers, even if the efficacy and effectiveness of the model are deemed favorably, are managers and executive ready to embrace that level of change in their firms? Would the CEO and board be ready to jolt the existing structures of a firm to introduce new departments

(e.g. data preprocessing) and restructure and reorganize the firm to respond to the new competitive challenges? The model does not address that.

Thirdly, the model assumes that managers and executives will be well-versed and educated about their own competitive realities; that they will rise to the occasion to confront the competition. Internet has been around for over two decades and many companies have not even embraced the power of the internet. Thus, expecting business to change because competitive dynamics have shifted may not be grounded in reality. The model ignores that part.

The entire methodology assumes that the organizations pursuing this will have extensive data and their other systems (conventional) will be in perfect state. In an ideal world, managers will be bringing AI as a new layer of competency built upon a perfectly functioning existing IT infrastructure. This is far from true. In many companies the conventional IT is in a disconcerting state and to expect those firms to consider an ivory tower model that worked for a firm which did not have these challenges is too much to ask for.

In practical considerations, social paradigms are not considered. To assume that companies first think about ethical and governance before engaging in projects reflects a rather naïve understanding of the world. In the real world, ethics and governance are often considered as the last things, that is if they are considered at all. Hence, as shown in the movie Brexit (formerly named Brexit, an uncivil war) (Haynes, 2019), the Brexit campaign designers did not consider the ethics of using ML to instigate sentiment and conflict in the UK which may have led to the assassination of Jo Cox (a British MP). This shows that strategic goals of entities become more important than their ethical considerations.

On a more functional level, while the first two steps of the model capture discovering and defining enterprise-wide projects, the model does not provide any basis of how companies should prioritize the discovered projects. Perhaps the discuss step, which focuses on financial value creation, can be expanded to

include a project prioritization mechanism. Return on investment alone cannot be used as the sole dimension to prioritize which projects will be implemented first. Project risk, a firm's capabilities, and other factors must also be considered.

As amply pointed out in conventional systems literature, even by incorporating the human in the agile design framework, using lightweight and flexible approach, and performing ethnographic analysis, one cannot ignore the biases and cognitive limitations and skills of the analyst (Morales-Ramirez and Alva-Martinez, 2018). Such limitations of analyst who gathers information for requirements can impact the quality of the design (Pitts and Browne, 2004). The AIAI NPD Framework offers no protection against people injecting their own biases in the design stage.

Furthermore, once requirements are obtained, validating them requires checks such as consistency, completeness, and realism (Bilal et al., 2016; Bendík, 2017). The AIAI model does not specifically calls for such validation features – even though they can be easily inserted into the model.

As software moved to cloud and microservices models developed, attempts were made to standardize requirements process by projects such as undertaken by Unicorn (Trihinas, 2017). Significant research also focused on requirements prioritization (Hudaib et al., 2018; Qaddoura et al., 2017). Other, more targeted, areas of requirements such as mobile, privacy, and other requirements were also analyzed (Zimmeck et al., 2016). Software development flexibility was deemed important and adaptive software requirements capability was developed (Ali and Hong, 2018). The AIAI model is a general-purpose ML model and does not consider specifics of infrastructure (cloud, telecom, mobile, etc.) and other details.

Given the above limitations, it is likely that managers may use the AIAI NPD Framework as a guideline and pick and choose some aspects of it while ignore others.

6.3 Ideas for Future Research

The AIAI NPD Framework lays the groundwork for establishing several interesting and relevant areas of exploration. From that perspective, the most obvious area will be to improve the AIAI model by eliminating some of the flaws pointed out in the previous section. The AIAI NPD Framework has established not only the need of ML to have its own dedicated NPD model but has also shown how to do that. Even those who might consider the model to be barebones skeleton, the task of adding flesh on the model can be an undertaking for decades to come. The AI economy is here to stay and the research on the technological side of AI is happening exponentially. In that scenario, there will never be a shortage of methodologies.

One of the unique attributes of the AIAI NPD Framework is the use of social paradigms to determine the ethos of a product or a service by understanding the perception of the management team about the social and ethics values. Today, we are quite aware of the ethical problems related to AI. Some AI problems arise from using biased data. At other times they result from malicious and sinister intentions. The AIAI Model does not go far enough to develop a comprehensive framework for governance and ethics. Much research is needed to understand how to ensure that learning machines are trained to be unbiased not only at the inception but also throughout their lifetime.

The AIAI model introduces the concept of the society of interdependent intelligent agents working collaboratively where agents interact with each other and with the environment in which they operate. These goal-oriented agents may also interact with humans. Since agents are learning systems with adaptive features, changes in their environment can affect them to change, evolve and adapt. This configuration and clustering can be viewed as a complex adaptive system. The AIAI model views

them as forming a complex adaptive system. Cilliers (1998) identified the ten properties of complex systems as follows (Cilliers, 1998):

- 1) Complex systems are composed of a large number of elements (Elements, Agents).
- 2) Elements interact with each other (Interaction).
- 3) Interaction among elements is rich. However, the behavior of the system is not determined by the exact number of interactions (Nature of Interaction).
- 4) Interactions are nonlinear (Math of Interaction).
- 5) Interactions have short range (Range of Interaction).
- 6) There are loops in the interaction that provide feedback; such that elements can feed information back to themselves (Recurrency).
- 7) Complex systems are open systems in the manner that they interact with the environment. As such scope of the system is determined by description of the system and the position of the observer (Framing).
- 8) Systems are far from equilibrium and depend upon constant flow to energy to stay dynamic and vibrant (Equilibrium).
- 9) History is important for complex systems as they have a past and past plays an important role in present behavior (History).
- 10) Each element of the system does not contain the consciousness of the entire system and lacks knowledge about the behavior of the system as a whole (Behavior Knowledge Paucity).

Each of the above ten points can offer an avenue of further research to develop the theoretical foundations of clusters of intelligent agents (humans and machines) functioning in a CAS. Analyzing and developing theory for clusters of intelligent machines and humans as a complex adaptive system and studying the behavior of that system will be a worthwhile exercise in years to come. Managing the states, health, and status clusters of systems will require metacognition.

I expect that this research has opened doors to significant follow-up research in the area of machine-human interaction to accomplish work. Specifically, the area of adaptive clusters of machines that can adapt, learn, and detect changes in the environment, that can change their collaboration and coordination preferences and work goals with other systems, and that form interdependent relationships with other systems and function as a complex system – can be the focus areas of study for decades to come.

The AI revolution has just started. This research was an attempt to formalize NPD in ML. It is understood, with a certain degree of human humbleness, that even that process itself will soon be automated. Automated NPD implies machines will discover, define, discuss, design, develop, and deploy new products – all by themselves. If that brings to mind a certain dystopian future, it also explains the obsession in this research with social paradigms, ethics, and governance. With every ML system born without those considerations, humankind is inching closer to the future we may come to regret – with or without an NPD methodology.

6.4 Future of NPD in ML

It is my expectation that companies will embrace autonomous automation as a strategic goal and will try to build companies around AI. Blake Morgan reported in Forbes magazine that Amazon has reorganized itself around artificial intelligence (Morgan, 2018). From all internal processes to all external processes, the firm is trying to develop a master plan for total and integrated automation. For such companies, the framework such as AIAI NPD Framework will become indispensable. These firms will try to develop a both bottom-up and top-down approach to develop total automation plans.

Eventually, the NPD process itself will likely be automated. That is where autonomous systems will design product concepts, generate launch plans, design and develop products, and launch the products without human intervention. In some cases, this is already happening. AIAI has presented a use case to a financial firm where credit cards can be customized for each customer and an intelligent software will automatically issue and manage credit cards, payments, supplier contracts, and offerings autonomously. There will be great need to NPD frameworks in the ML space.

Chapter 7 Reflections of a Scholar-Practitioner

Implementing the action research process in this study was one of the most powerful and moving experiences of my life. The collaboration, insights, and participation of the team members was inspirational. This action research was not just about discovery, it was also about understanding my own thoughts and paying attention to my experiences. Throughout the process I learned to critically analyze my own thinking and assumptions and recognize my biases. During each cycle, I had to pay attention to reflect upon and uncover the flaws in my research assumptions. It taught me to critically reflect and think about my own thoughts.

The first lesson – and perhaps the most important one – was a strange recognition about my own ignorance. Never in my life I felt so intellectually challenged as I felt during the study. The process of discovering, researching, and implementing – all at the same time – was profoundly exciting. But behind the veil of a confident executive, as the study progressed, an overwhelming sense of fear dawned on me. I recognized that without the scholarly side, how unprepared I was as a practitioner. More importantly, how little self-knowledge I possessed to know what I did not know. This wave of intellectual humility and intellectual integrity materialized as I learned to question my own plans, thinking, and research approaches.

I recognized that I had carried a false sense of security for most of my life and that was that *I know all the answers*. It came from being the person who gets the job done and that getting the job done implied to strategize, plan, and execute. That is what practitioners do. But this challenge was different. In this I had to work with a team to discover new knowledge. The plan was not derived from some consultant presentation, an HBR article, a business book, or a textbook. It was not created from my own thinking, imagination, and creativity. For those were the things I had done as a practitioner, but in this case, it required an intermediate step to approach matters from a scientific perspective, while at all times knowing that the experience of the team members is also a rich source of data. If this was a bridge to becoming a scholar-practitioner, being on that bridge brought a deep sense of humbleness. The substantial

knowledge I gained from this study is also a stark reminder of how little I knew about business and myself. The double helix nature of scholar-practitioner became obvious.

It is true that I had to explore multiple fields for this research. I researched NPD but also information systems development methodologies, machine learning frameworks, philosophy of knowledge, and other business areas that needed coverage to develop an entire product development process. This project was not just about making the product development work at AIAI, it was also about extracting usable information about the process that makes ML NPD successful. It was about de-risking the process of NPD in intelligent systems. A lot resided on the success of this project. But despite all the research that was done, and all the books and articles analyzed and read, what created powerful dynamics was the interaction with the team members. “They know so much” was the first feeling of appreciation that came to me naturally – and was followed by questioning “Why had we not tapped their knowledge and experience before?”. The fact that people are so much more than their resumes, annual reviews, and titles became obvious. The team members shared important insights that I did not expect. Their passion to make the project successful, their devotion, and their intellectual contribution was immense. I recognized that as a leader, we are often boxed into viewing people in an extremely limited functional sense. When given an opportunity to develop and think of solutions, their creative sides flourish. Their imagination flows and they become leaders. A leader’s primary function is, I learned, to bring out the best in his or her people and the only way to do that is to have trust in their potential.

My learning also transpired in deploying the scientific process from a non-positivist angle. This learning not only manifested from the application of action research, but also from the fact that in literature review I analyzed and reconfigured the Burrell and Morgan model (Burrell and Morgan, 1979). To structure the paradigm centric model for AIAI NPD, I needed to acquire significant expertise in the four scientific paradigms. It involved an in-depth study of the underlying epistemological, ontological, and ethical attributes of the four paradigms. As I learned about them, I recognized that while paradigms have been used to analyze scientific ideas and for scientific discovery, they can also be used in product and services development – specially to develop intelligent products. Intelligence, like science, is not only positivist. For example, natural language

processing shows us how to extract human sentiment. This is an extension of ethnographic studies. It implies that intelligence can be built from at least four different angles – functionalism (positivism), social relativism, radical structuralism, and neohumanism.

Despite the knowledge of the above four paradigms, there has been a dearth of approaching systems development from a neohumanism paradigm, I realized. We are building systems as a response to business needs but not for social good. We assume that social good will manifest if enough systems are in place or that an invisible hand will somehow create social good. We place all faith in our capitalistic endeavors and in the power of technology. But designs that are not baptized for social good at their inception do not result in upholding justice and good at their maturity. For example, a building constructed to house people that does not consider environmental damage, carbon emissions, health considerations of the dwellers, and effects on the neighborhood, can be a functional or an aesthetic masterpiece, but its design was grounded in social disservice and during all its later functional stages will continue to hurt the human civilization. This was my greatest learning from this study: design must manifest social good at the inception stage. Whether in physical products or intelligent products and services, the social good must start at the inception stage.

During this process I questioned my fundamental purpose of being, of doing anything and everything I am engaged in. I recognized that it is that purpose that defines what questions I will ask and what problems I will solve. At a very fundamental level the structure of thought, I learned, is tied to the sense of purpose a person possesses. I discovered that social justice was the driving force behind everything I pursue in life. That is why the link between NPD and ethics/social justice became apparent to me and my subsequent thinking after the literature review carried that theme across the six cycles. I recognized that purpose often defines what type of information we collect and how we analyze information or form assumptions. AIAI did not have the corporate bone. It was not a large company with deep structural fiefdoms and cutthroat politics. It was due to the academic roots and the social consciousness embedded in our values that our approach for product envisioning was not purely commercial and during the research I did explore the merits of ethics and social justice as drivers of NPD. AIAI's

NPD Model begins with the paradigmatic clarity and calls for ethics and social justice to be included in designs. As the interdependence between thoughts, feelings, and desires (purpose) became more obvious, I recognized that as a leader of AIAI I did not just seek financial returns. I wanted social justice reflected in our pursuits and designs and felt strongly about it – and as a result my thoughts guided me to assign greater value to that part in the design.

I also learned that clarity of thinking comes when one escapes the two traps of the dominating ego and the submissive ego. I had fallen victim to both of these traps in my previous practitioner-only centric life. I had strong ideas and feelings about how things should be done, and I pushed them as part of the agenda. In other cases, I strictly abided by the guidance and direction provided by investors or board members – refusing to question. Both of these states of ego are counter to what the DBA program and the research has taught me. I have developed, and I continue to develop, an open mind that is neither controlling nor submitting without thinking and questioning the basis of thinking, reason, assumptions, and data. I must acknowledge that it is not something that comes easy to me and that I still struggle with this on a daily basis. But I am now aware and have created reminders for myself to apply the right type of thinking mechanisms and to question. Learning about biases was fundamental to critically reflect on issues and decisions. Biases come from many sources and having the ability to analyze your own thought patterns requires a meta state of awareness. This, I believe, is an acquired skill and I am trying to develop that.

The three fundamentals of work, mutual understanding, and emancipation around which the society is organized require the development of symbols and shared meaning. But despite the calls for the development of shared meaning humans remained unprotected against bias and injustice. The lack of truth and deterioration in quality of human condition are evidenced by the rise of nationalism, environmental degradation, and fake news. The civilization is under the threat of self-enforced annihilation. And in the midst of all that we have introduced intelligent machines that lacks any feelings and empathy. It grinds and thinks and decides and acts but lacks any moral sense. If there was any reason to regulate all human affairs with a desire for truth and to establish justice as the norm, it is now. This research was about learning systems and the human

who conducted this research has learned that learning machines must be designed and governed with the human principles of social justice.

References

- Agrawal, Ajay, Gans, Joshua & Goldfarb, Avi (2019) *The economics of artificial intelligence : an agenda*. Chicago: The University of Chicago Press.
- Allen, Peter M. & Varga, Liz (2006) A co-evolutionary complex systems perspective on information systems. *Journal of Information Technology*. [Online] 21 (4), 229–238.
- Alter, Steven (2008) Defining Information Systems as Work Systems: Implications for the IS Field. *Business Analytics and Information Systems*. [Online] Paper 22. [online]. Available from: <http://repository.usfca.edu/at/22%0A>.
- Alvesson, M. & Sköldbberg, K. (2009) *Reflexive Methodology*. London: Sage.
- Alvesson, Mats, Hardy, Cynthia & Harley, Bill (2008) Reflecting on Reflexive Practices in Organization and Reflecting on Reflexive Practices in Organization and Management Theory. *Journal of Management Studies*. 45 (3), 480–501.
- Amershi, Saleema, Begel, Andrew, Bird, Christian, DeLine, Robert, Gall, Harald, Kamar, Ece, Nagappan, Nachiappan, Nushi, Besmira & Zimmermann, Thomas (2019) Software Engineering for Machine Learning: A Case Study. *Proceedings - 2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice, ICSE-SEIP 2019*. [Online] 291–300.
- Araujo, Theo (2018) Living up to the chatbot hype: The influence of anthropomorphic design cues and communicative agency framing on conversational agent and company perceptions. *Computers in Human Behavior*. [Online] 85183–189. [online]. Available from: <https://doi.org/10.1016/j.chb.2018.03.051>.
- Argyris, C. & Schön, D. (1974) *Theory in practice: Increasing professional effectiveness*. San Francisco, CA: Jossey-Bass.
- Argyris, Chris (1976) Single-Loop and Double-Loop Models in Research on Decision Making. *Administrative Science Quarterly*. 21363–377.
- Avison, D. E. & Fitzgerald, G. (1995) *Information Systems Development: Methodologies, Techniques and Tools*. 2nd edition. New York: McGraw-Hill.
- Aydin, Mehmet Emin & Fellows, Ryan (2017) *A reinforcement learning algorithm for building collaboration in multi-agent systems*. [online]. Available from: <http://arxiv.org/abs/1711.10574>.
- Aydin, Mehmet Emin & Fellows, Ryan (2018) Building collaboration in multi-agent systems using reinforcement learning. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. [Online] 11056 LNAI201–212.
- Bansal, P. & Corley, K. (2011) The coming of age for qualitative research: Embracing the diversity of qualitative methods. *Academy of Management Journal*. 54 (2), 233–237.
- Becker, B. (2006) Rethinking the Stage-Gate process—A reply to the critics.

Management Roundtable July 12.

- Bishop, C. M. (2006) *Pattern recognition and machine learning*. Springer.
- Boxwala, Aziz A., Kim, Jihoon, Grillo, Janice M. & Ohno-Machado, Lucila (2011) Using statistical and machine learning to help institutions detect suspicious access to electronic health records. *Journal of the American Medical Informatics Association*. [Online] 18 (4), 498–505.
- Boyatzis, R. E. (1998) *Transforming Qualitative Information: Thematic Analysis and Code Development*. Thousand Oaks, CA: Sage.
- Boyd, Aaron (2018) Management Chief: Workforce Must Change to Reach IT Modernization Goals. Nextgov [online]. Available from: <https://www.nextgov.com/it-modernization/2018/04/management-chief-workforce-must-change-reach-it-modernization-goals/147643/>.
- Boyd, Aaron (2019) *The new Intelligent Automation/Artificial Intelligence contract will be housed in Health and Human Services but available to any agency looking for AI solutions* [online]. Available from: <https://www.nextgov.com/emerging-tech/2019/01/hhs-contract-will-offer-ai-tech-support-all-government/154078/>.
- Braun, Virginia & Clarke, Victoria (2006) Using thematic analysis in psychology. *Qualitative Research in Psychology*. [Online] 3 (2), 77–101.
- Brockhoff, Klaus (2003) Customers' perspectives of involvement in new product development. *International Journal of Technology Management*. [Online] 26 (5–6), 464–481.
- Brydon-Miller, Mary, Greenwood, Davydd, Maguire, Patricia & Bradbury, Hilary (2003) Why action research? *Action Research*. 11 (11), 9–28.
- Brynjolfsson, Eric & McAfee, Andrew (2011) *Race against the machine : how the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*. Lexington: Digital Frontier Press.
- Brynjolfsson, Erik & McAfee, Andrew (2015) Will humans go the way of horses? Labor in the second machine age. *Foreign Affairs*. 94 (4), .
- Brynjolfsson, Erik, Rock, Daniel & Syverson, Chad (2017) *Artificial Intelligence and the Modern Productivity Paradox :*
- Bubenko Jr, J. A. Gustafsson, M. R. & Karlsson, T. (1982) 'A declarative approach to conceptual information modeling. In', in T W Olle, H G Sol, & A A Verrijn-Stuart (eds.) *Information Systems Design Methodologies: a Comparative Review*. The Netherlands: North-Holland. pp. 93–142.
- Bur, Jessie (2019) Government wants to know how to make data better for AI. Federal Times [online]. Available from: <https://www.federaltimes.com/it-networks/2019/07/10/government-wants-to-know-how-to-make-data-better-for-ai/>.
- Burrell, Gibson & Morgan, Gareth (1979) *Sociological Paradigms and Organisational*

- Analysis: Elements of the Sociology of Corporate Life*. New York, NY: Routledge.
- Carr, W. & Kemmis, S. (1986) (1986) *Becoming Critical: Education, Knowledge and Action Research*. London: Falmer.
- Carr, W. & Krennmis, S. (1986) *Becoming Critical*. Lewes: Falmer Press.
- Chandler, Dawn & Torbert, Bill (2003) Transforming Inquiry and Action: Interweaving 27 Flavors of Action Research. *Action Research*. [Online] 1 (2), 133–152.
- Chapman, Pete, Clinton, Julian, Kerber, Randy, Khabaza, Thomas, Reinartz, Thomas, Shearer, Colin & Wirth, Rüdiger (2000) *Crisp-dm 1.0*. [online]. Available from: file:///C:/Users/alnaq/Downloads/crisp-dm.pdf.
- Chevalier, J. M. & Buckles, D. J. (2013) *Research: Theory and Methods for Engaged Inquiry*. UK: Routledge.
- Chou, Charles Ling yu, Du, Timon & Lai, Vincent S. (2007) Continuous auditing with a multi-agent system. *Decision Support Systems*. [Online] 42 (4), 2274–2292.
- Christiansen, John K. & Varnes, Claus J. (2009) Formal rules in product development: Sensemaking of structured approaches. *Journal of Product Innovation Management*. [Online] 26 (5), 502–519.
- Cilliers, P. (1998) *Complexity and post-modernism: Understanding complex systems*. London: Routledge.
- Coghlan, David (2001) Insider Action Research Projects. *Management Learning* 32 (1) p.49–60.
- Coghlan, David & Brannick, Teresa (2014) *Doing Action Research in Your Own Organization*. 4th edition. UK: Sage Publications Ltd.
- Cooper, Robert G. (2014) What's next? After stage-gate. *Research Technology Management*. [Online] 57 (1), 20–31. [online]. Available from: <https://doi.org/10.5437/08956308X5606963>.
- Cooper, Robert G. & Sommer, Anita Friis (2018) Agile–Stage-Gate for Manufacturers: Changing the Way New Products Are Developed Integrating Agile project management methods into a Stage-Gate system offers both opportunities and challenges. *Research Technology Management*. [Online] 61 (2), 17–26.
- Coughlan, Paul & Coghlan, David (2002) Action research for operations management. *International Journal of Operations and Production Management*. [Online] 22 (2), 220–240.
- Dåderman, Antonia & Rosander, Sara (2018) Evaluating Frameworks for Implementing Machine Learning in Signal Processing A Comparative Study of CRISP-DM, SEMMA and KDD. *EXAMENSARBETE INOM TEKNIK, GRUNDNIVÅ, 15 HP STOCKHOLM, SVERIGE* 2. 1–43. [online]. Available from: <http://www.diva-portal.se/smash/get/diva2:1250897/FULLTEXT01.pdf>.
- Darwish, Nagy Ramadan & Megahed, Salwa (2016) Requirements Engineering in Scrum

- Framework. *International Journal of Computer Applications*. [Online] 149 (8), 24–29.
- Deetz, Stanley (1996) Crossroads — Describing Differences in Approaches to Organization Science : Rethinking Burrell and Morgan and Their Legacy. *Organization Science*. 7 (2), 191–720.
- DeFond, Mark & Zhang, Jieying (2014) A review of archival auditing research. *Journal of Accounting and Economics*. [Online] 58 (2–3), 275–326. [online]. Available from: <http://dx.doi.org/10.1016/j.jacceco.2014.09.002>.
- Deloitte (2016) *Deloitte: Automation set to transform public services* [online]. Available from: <https://www2.deloitte.com/uk/en/pages/press-releases/articles/automation-set-to-transform-public-services.html>.
- Dewey, J. (1929) *The Quest for Certainty*. 1960th, 3rd edition. New York: Capricorn Books.
- DiMasi, J. A., Feldman, L., Seckler, A. & Wilson, A. (2009) Trends in Risks Associated With New Drug Development: Success Rates for Investigational Drugs. *Clinical Pharmacology & Therapeutics*. [Online] 87 (3), 272–277. [online]. Available from: <http://dx.doi.org/10.1038/clpt.2009.295>5Cnpapers3://publication/doi/10.1038/clpt.2009.295.
- Dirican, Cüneyt (2015) The Impacts of Robotics, Artificial Intelligence On Business and Economics. *Procedia - Social and Behavioral Sciences*. [Online] 195564–573. [online]. Available from: <http://www.sciencedirect.com/science/article/pii/S1877042815036137>.
- Dmitrenko, Maksim & Matsegora, Andrey (2017) Continuous Auditing As an Element of Effective Corporate Governance. *Scientific Journal of Polonia University*. [Online] 24 (5), 91.
- Domingos, P. (2015) *The master algorithm: How the quest for the ultimate learning machine will remake our world*. New York, NY: Basic Books.
- Donoho, David (2017) 50 Years of Data Science. *Journal of Computational and Graphical Statistics*. [Online] 26 (4), 745–766.
- Eason, Oriana Karina (2016) Information Systems Development Methodologies Transitions: An Analysis of Waterfall to Agile Methodology. *University of New Hampshire*. 1–23. [online]. Available from: <http://scholars.unh.edu/honors%0Ahttp://scholars.unh.edu/honors>.
- Eden, Colin & Huxham, Chris (1996) Action Research for Management Research. *British Journal of Management*. [Online] 7 (1), 75–86. [online]. Available from: <http://onlinelibrary.wiley.com/doi/10.1111/j.1467-8551.1996.tb00107.x/pdf>.
- Ellul, J. (1965) *The technological society*. London: Cape.
- ElSayed, Iman Ahmed, Ezz, Zeinab & Nasr, Eman (2017) Goal modeling techniques in requirements engineering: A systematic literature review. *Journal of Computer*

- Science*. [Online] 13 (9), 430–439.
- Farhi, Edward, Goldstone, Jeffrey & Gutmann, Sam (2014) *A Quantum Approximate Optimization Algorithm*. 1–16. [online]. Available from: <http://arxiv.org/abs/1411.4028>.
- Fayyad, Usama, Smyth, Padhraic & Piatetsky-Shapiro, Gregory (1996) From Data Mining to Knowledge Discovery in Databases. *AI Magazine*. [Online] 17 (3), 37–54.
- Fire, Michael & Puzis, Rami (2016) Organization Mining Using Online Social Networks. *Networks and Spatial Economics*. [Online] 16 (2), 545–578.
- Frey, Carl Benedikt & Osborne, Michael A. (2017) The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change*. 114254–280.
- Friedrich, Thomas, Schlauderer, Sebastian, Weidinger, Julian & Raab, Maximilian (2017) On the Research Paradigms and Research Methods Employed in the BISE Journal - A Ten-Year Update. *Wirtschaftsinformatik 2017: Proceedings*. (February 2017), 1111–1125.
- Gepp, Adrian, Linnenluecke, Martina K., O'Neill, Terrence J. & Smith, Tom (2018) Big data techniques in auditing research and practice: Current trends and future opportunities. *Journal of Accounting Literature*. [Online] 40 (February), 102–115. [online]. Available from: <https://doi.org/10.1016/j.acclit.2017.05.003>.
- Giroux, Gary & Cassell, Cory (2011) Changing audit risk characteristics in the public client market. *Research in Accounting Regulation*. [Online] 23 (2), 177–183. [online]. Available from: <http://dx.doi.org/10.1016/j.racreg.2011.06.009>.
- Gjerde, Susann & Alvesson, Mats (2020) Sandwiched: Exploring role and identity of middle managers in the genuine middle. *Human Relations*. [Online] 73 (1), 124–151.
- Goldkuhl, G. & Lyytinen, K. (1982) 'A language action view of information systems', in C Ross & M. Ginzberg (eds.) *Proceedings of the Third Conference on Information Systems*. Ann Arbor, Michigan: . pp. 13–30.
- Gordon, Myron J. & Shapiro, E. L. I. (1956) CAPITAL EQUIPMENT ANALYSIS : THE REQUIRED RATE OF PROFIT. *Management Science*. 3 (1), 102–110.
- Greenwood, Davydd J. & Morten, Levin (2007) *Introduction to Action Research: Social Research for Social Change*. 2nd edition. Thousand Oaks CA: SAGE Publications, Inc.,
- Greenwood, Davydd J., Whyte, Willian F. & Harkavy, Ira (1993) Participatory Action Research as a Process and as a Goal. *Human Relations*. 46 (2), 175–192.
- Gregor, Shirley (2006) The nature of theory in Information Systems. *MIS Quarterly: Management Information Systems*. [Online] 30 (3), 611–642.
- Grönlund, Johan, Sjödin, David Rönnberg & Frishammar, Johan (2010) Open innovation and the stage-gate process: A revised model for new product

- development. *California Management Review*. [Online] 52 (3), 106–131.
- Gruner, Stefan (2011) Problems for a philosophy of software engineering. *Minds and Machines*. [Online] 21 (2), 275–299.
- Gruner, Stefan (2010) Software engineering between technics and science: Recent discussions about the foundations and the scientificness of a rising discipline. *Journal for General Philosophy of Science*. [Online] 41 (1), 237–260.
- Gruner, Stefan & Kroeze, Jan (2014) On the shortage of engineering in recent information systems research. *Proceedings of the 25th Australasian Conference on Information Systems, ACIS 2014*.
- Hagendorff, Thilo (2020) The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds and Machines*. [Online] 30 (1), 99–120. [online]. Available from: <https://doi.org/10.1007/s11023-020-09517-8>.
- Hajkowicz, SA, Karimi, S., Wark, T., Chen, C., Evans, M., Rens, N., Dawson, D., Charlton, A., Brennan, T., Moffatt, C., Srikumar, S. & Tong, KJ (2019) *Artificial intelligence: Solving problems, growing the economy and improving our quality of life*. [online]. Available from: <https://data61.csiro.au/en/Our-Research/Our-Work/AI-Roadmap>.
- Haynes, Toby (2019) *Brexit*. [online]. Available from: <https://www.imdb.com/title/tt8425058/>.
- Heinrich, L. (1999) *Grundlagen der Wirtschaftsinformatik. Part F in Rechenberg, P., and Pomberger, G., (eds.): Informatik-Handbuch*. 2nd ed. Carl Hanser Publ.
- Hengstler, Monika, Enkel, Ellen & Duelli, Selina (2016a) Applied artificial intelligence and trust-The case of autonomous vehicles and medical assistance devices. *Technological Forecasting and Social Change*. [Online] 105105–120. [online]. Available from: <http://dx.doi.org/10.1016/j.techfore.2015.12.014>.
- Hengstler, Monika, Enkel, Ellen & Duelli, Selina (2016b) Technological Forecasting & Social Change Applied artificial intelligence and trust — The case of autonomous vehicles and medical assistance devices. *Technological Forecasting & Social Change*. [Online] 105105–120. [online]. Available from: <http://dx.doi.org/10.1016/j.techfore.2015.12.014>.
- Hesenius, Marc, Schwenzfeier, Nils, Meyer, Ole, Koop, Wilhelm & Gruhn, Volker (2019) Towards a software engineering process for developing data-driven applications. *Proceedings - 2019 IEEE/ACM 7th International Workshop on Realizing Artificial Intelligence Synergies in Software Engineering, RAISE 2019*. [Online] 35–41.
- Hirschheim, Rudy & Klein, Heinz K. (1989) Four Paradigms of Information Systems Development. *Communications of the ACM*. 32 (10), 1199–1216.
- Hoffman, Charles (2017) *Accounting and Auditing in the Digital Age*. 1–14. [online]. Available from: <http://bruegel.org/2017/04/do-we-understand-the-impact-of-artificial-intelligence-on-employment/>.

- Houkes, Wybo & Vermaas, Pieter E. (2009) Contemporary Engineering and the Metaphysics of Artefacts. *Monist*. [Online] 92 (3), 403–419.
- Huber, Steffen, Wiemer, Hajo, Schneider, Dorothea & Ihlenfeldt, Steffen (2019) DMME: Data mining methodology for engineering applications - A holistic extension to the CRISP-DM model. *Procedia CIRP*. [Online] 79403–408. [online]. Available from: <https://doi.org/10.1016/j.procir.2019.02.106>.
- Iivari, J. (1991) A paradigmatic analysis of contemporary schools of IS development. *European Journal of Information Systems*. 1 (4), .
- Iivari, Juhani (2017a) Conducting information systems research the old-fashioned way. *Scandinavian Journal of Information Systems*. 29 (1), .
- Iivari, Juhani (2017b) Information system artefact or information system application: that is the question. *Information Systems Journal*. [Online] 27 (6), 753–774.
- Iivari, Juhani & Hirschheim, Rudy (1996) Analyzing information systems development: A comparison and analysis of eight is development approaches. *Information Systems*. [Online] 21 (7), 551–575.
- Iivari, Juhani, Hirschheim, Rudy & Klein, Heinz K. (2000) A dynamic framework for classifying information systems development methodologies and approaches. *Journal of Management Information Systems*. [Online] 17 (3), 179–218.
- Iivari, Juhani, Hirschheim, Rudy & Klein, Heinz K. (1998) A Paradigmatic Analysis Contrasting Information Systems Development Approaches and Methodologies. *Information Systems Research*. [Online] 9 (2), 164–193.
- Iivari, Netta (2019) Discursive struggles within cross-disciplinary design. *40th International Conference on Information Systems, ICIS 2019*. 1–17.
- Iivari, Netta & Kuutti, Kari (2018) Towards Critical Design Science Research. *ICIS 2017: Transforming Society with Digital Innovation*. 0–15.
- Isaias, Pedro & Issa, Tomayess (2015) *High Level Models and Methodologies for Information Systems*. New York: Springer Science Business Media.
- Jamshidi, Pooyan, Camara, Javier, Schmerl, Bradley, Kaestner, Christian & Garlan, David (2019) Machine learning meets quantitative planning: Enabling self-adaptation in autonomous robots. *ICSE Workshop on Software Engineering for Adaptive and Self-Managing Systems*. [Online] 2019-May39–50.
- Jans, Mieke & Hosseinpour, Marzie (2019) How active learning and process mining can act as Continuous Auditing catalyst. *International Journal of Accounting Information Systems*. [Online] 32 (November 2018), 44–58. [online]. Available from: <https://doi.org/10.1016/j.accinf.2018.11.002>.
- Jarrahi, Mohammad Hossein (2018) Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*. [Online] 61 (4), 577–586. [online]. Available from: <https://doi.org/10.1016/j.bushor.2018.03.007>.

- Jayarathna, N. (1994) *Understanding and Evaluating Methodologies*, NISAD: A Systematic Framework. Maidenhead, UK McGraw-Hill.
- Joffe, H. (2012) 'Thematic analysis', in D. Harper & A. R. Thompson (eds.) *Qualitative Research Methods in Mental Health and Psychotherapy*. Chichester: John Wiley. p.
- Kahn, Kenneth B., Barczak, Gloria & Moss, Roberta (2006) Dialogue on Best Practices in New Product Development - Perspective: Establishing an NPD Best Practices Framework. *Journal of Product Innovation Management*. 23 (2), 106–116.
- Kaushik, Chandra Aman & Raj, Utkarsh (2020) AI-driven drug discovery: A boon against COVID-19? *AI Open*. 11–4. [online]. Available from: <https://doi.org/10.1016/j.aiopen.2020.07.001>.
- Kaya, Heysem & Salah, Albert Ali (2018) *Multimodal Personality Trait Analysis for Explainable Modeling of Job Interview Decisions*. [Online] 255–275.
- Kevin Casey (2019) *What is explainable AI?* [online]. Available from: <https://enterpriseproject.com/article/2019/5/what-explainable-ai>.
- Kinoshita, Tomoo, Hayashi, Shinpei & Saeki, Motoshi (2017) Goal-Oriented Requirements Analysis Meets a Creativity Technique. *S. de Cesare and U. Frank (Eds.): ER 2017 Workshops*. [Online] LNCS 10651101–110.
- Klein, H. K. & Hirschheim, R. A. (1987a) A Comparative Framework of Data Modelling Paradigms and Approaches. *The Computer Journal*. 30 (1), .
- Klein, H. K. & Hirschheim, R. A. (1987b) A Comparative Framework of Data Modelling Paradigms and Approaches. *The Computer Journal*. 30 (1), .
- Knight, J., & Jack, K. (1992) *Institutions and social conflict*. Cambridge University Press.
- Koschwitz, D., Frisch, J. & van Treeck, C. (2018) Data-driven heating and cooling load predictions for non-residential buildings based on support vector machine regression and NARX Recurrent Neural Network: A comparative study on district scale. *Energy*. [Online] 165134–142. [online]. Available from: <https://doi.org/10.1016/j.energy.2018.09.068>.
- Koskela, Lauri & Kagioglou, Mike (2005) On the Metaphysics of Production. In: *13th Annual Conference of the International Group for Lean Construction*. July. [online]. Available from: <http://eprints.hud.ac.uk/id/eprint/26011/>.
- Krush, Michael T., Agnihotri, Raj, Trainor, Kevin J. & Nowlin, Edward L. (2013) Enhancing organizational sensemaking: An examination of the interactive effects of sales capabilities and marketing dashboards. *Industrial Marketing Management*. [Online] 42 (5), 824–835. [online]. Available from: <http://dx.doi.org/10.1016/j.indmarman.2013.02.017>.
- Kuhn, T. S. (1970) *The structure of scientific revolutions*. Chicago, University of Chicago Press.

- Kurgan, Lukasz A. & Musilek, Petr (2006) A survey of Knowledge Discovery and Data Mining process models. *Knowledge Engineering Review*. [Online] 21 (1), 1–24.
- Langford, John & Ortega, Ruben (2012) Machine Learning and Algorithms; Agile Development. *Communications of the ACM*. [Online] 55 (8), 10–11.
- Lauterbach, By Anastassia & Bonim, A. (2016) Artificial Intelligence : A Strategic Business And Governance Imperative. *NACD Directorship*. (October), 54–57.
- Lee, Allen S., Thomas, Manoj & Baskerville, Richard (2015) Going back to basics in design science: from the information technology artifact to the information systems artifact. *Information Systems Journal*. 25 (1), 5–21.
- Lee, Boyeun, Cooper, Rachel, Hands, David & Coulton, Paul (2019) Design Drivers: A critical enabler to mediate value over the NPD process within Internet of Things. *4d Conference Proceedings: Meanings of Design in the Next Era. Osaka : DML (Design Management Lab), Ritsumeikan University*. 202596–107.
- Lee, Min Kyung, Kusbit, Daniel, Metsky, Evan & Dabbish, Laura (2015) Working with Machines. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*. [Online] (June), 1603–1612. [online]. Available from: <http://dl.acm.org/citation.cfm?doid=2702123.2702548>.
- Lenfle, S. & Loch, C. (2010) Lost roots: How project management came to emphasize control over flexibility and novelty. *California Management Review*. 53 (1), 32–55.
- Levy, Steven (2018) Inside Amazon's Artificial Intelligence Flywheel. *Wired* [online]. Available from: <https://www.wired.com/story/amazon-artificial-intelligence-flywheel/>.
- Li, Yunpeng, Roy, Utpal & Saltz, Jeffrey S. (2019) Towards an integrated process model for new product development with data-driven features (NPD 3). *Research in Engineering Design*. [Online] 30 (2), 271–289. [online]. Available from: <http://dx.doi.org/10.1007/s00163-019-00308-6>.
- Makridakis, Spyros (2017) *The Forthcoming Artificial Intelligence (AI) Revolution : Its Impact on Society and Firms The Forthcoming Artificial Intelligence (AI) Revolution : Its Impact on Society and Firms*.
- Marbán, Oscar, Segovia, Javier, Menasalvas, Ernestina & Fernández-Baizán, Covadonga (2009) Toward data mining engineering: A software engineering approach. *Information Systems*. [Online] 34 (1), 87–107.
- Mariscal, Gonzalo, Marbán, Óscar & Fernández, Covadonga (2010) A survey of data mining and knowledge discovery process models and methodologies. *Knowledge Engineering Review*. [Online] 25 (2), 137–166.
- Marquardt, Michael; & Yeo, Roland K. .. (2012) *Breakthrough Problem Solving with Action Learning*. Kindle. Stanford University Press.
- Martinez-Plumed, Fernando, Contreras-Ochando, Lidia, Ferri, Cesar, Hernandez Orallo, Jose, Kull, Meelis, Lachiche, Nicolas, Ramirez Quintana, Maria Jose & Flach, Peter

- A. (2019) CRISP-DM Twenty Years Later: From Data Mining Processes to Data Science Trajectories. *IEEE Transactions on Knowledge and Data Engineering*. [Online] 1–1.
- McAfee, Andrew & Brynjolfsson, Erik (2016) Human Work in the Robotic Future. *Foreign Affairs*. 95 (4), 139–150. [online]. Available from: <http://search.ebscohost.com/login.aspx?direct=true&db=ssa&AN=115986098&site=ehost-live>.
- McClure, Paul K. (2018) “You’re Fired,” Says the Robot: The Rise of Automation in the Workplace, Technophobes, and Fears of Unemployment. *Social Science Computer Review*. [Online] 36 (2), 139–156.
- Mehta, Yash, Majumder, Navonil, Gelbukh, Alexander & Cambria, Erik (2019) *Recent Trends in Deep Learning Based Personality Detection*. [online]. Available from: <http://arxiv.org/abs/1908.03628>.
- Meligy, Ali, Dabour, Walid & Farhat, Alaa (2018) The role of ethnography in agile requirements analysis. *ACM International Conference Proceeding Series*. [Online] 27–31.
- Merkel-Davies, Doris M., Brennan, Niamh & McLeay, Stuart (2011) Impression management and retrospective sense-making in corporate narratives : a social psychology perspective. *Accounting, Auditing and Accountability*. [Online] 24 (3), 315–344.
- Meservy, Rayman D., Denna, Eric L. & Hansen, James V. (1992) Application of artificial intelligence to accounting, tax, and audit services: Research at Brigham Young University. *Expert Systems With Applications*. [Online] 4 (2), 213–218.
- Mezirow, J. (1991) *Transformative Dimensions of Adult Learning*. San Francisco, CA: Jossey-Bass.
- Mitchell, Tom M. (1997) *Machine Learning*. McGraw-Hill.
- Mitchell, Tom M. (2006) The Discipline of Machine Learning. *Carnegie Mellon ML Department*. [Online] 17 (July), 1–7. [online]. Available from: <http://www-cgi.cs.cmu.edu/~tom/pubs/MachineLearningTR.pdf>.
- Morgan, Blake (2018) How Amazon Has Reorganized Around Artificial Intelligence And Machine Learning. *Forbes* [online]. Available from: <https://www.forbes.com/sites/blakemorgan/2018/07/16/how-amazon-has-reorganized-around-artificial-intelligence-and-machine-learning/?sh=403ca4aa7361>.
- Moser, Gabriel (1990) In Memoriam, Arthur Samuel: Pioneer in Machine Learning. *AI Magazine (AAAI)*. [Online] 11 (3), 10–11.
- Naqvi, Al & Munoz, J. Mark (2020) Rebuilding a Company with Artificial Intelligence. *California Management Review*. [online]. Available from: <https://cmr.berkeley.edu/2020/03/rebuilding-a-company-with-ai/>.
- Nimdzi Insights Pactera EDGE (2019) *Artificial Intelligence: Localization Winners*,

- Losers, Heroes, Spectators, and You. Technical Report.* [online]. Available from: <https://www.nimdzi.com/wp-content/uploads/2019/06/Nimdzi-AI-whitepaper.pdf>.
- Nobre, Farley Simon (2012) Governing industrial organizations through cognitive machines. *AI and Society*. [Online] 27 (4), 501–507.
- Nugroho, Suryanto, Hadi, Sigit & Hakim, Luqman (2017) Comparative Analysis of Software Development Methods between Parallel, V-Shaped and Iterative. *International Journal of Computer Applications*. [Online] 169 (11), 7–11.
- O'Meara, Richard M. (2012) Contemporary Governance Architecture Regarding Robotics Technologies: An Assessment. *Robot Ethics: The Ethical and Social Implications of Robotics*.
- Orlikowski, Wanda J. & Baroudi, Jack J. (1991) Studying information technology in organizations: Research approaches and assumptions. *Information Systems Research*. [Online] 2 (1), 1–28.
- Perez, Carlota (2016) *Capitalism , Technology and a Green Global Golden Age : The Role of History in Helping to Shape the Future*. (April), .
- Piatetsky, Gregory (2014) CRISP-DM remains the most popular methodology for analytics, data mining, and data science projects. KDnuggets [online]. Available from: <https://www.kdnuggets.com/2014/10/crisp-dm-top-methodology-analytics-data-mining-data-science-projects.html>.
- Pomykalski, James J. & Buzydlowski, Jan (2017) *Using Systems Analysis and Design to Enhance the Business Understanding Stage in CRISP-DM*. 1–7.
- Porter, Michael E. & Heppelmann, James E. (2015) How smart, connected products are transforming companies. *Harvard Business Review*. 2015 (October), .
- Porter, Michael E. & Heppelmann, James E. (2014) How Smart, Connected Products Are Transforming Competition.(Spotlight On Managing The Internet Of Things). *Harvard Business Review*. 92 (11), 64. [online]. Available from: https://www.gospi.fr/IMG/pdf/porter-2014-hbr_how-smart-connected-products-are-transforming-competitionhbr-2014.pdf.
- Potter, J. & Wetherell, M. (1987) *Discourse and Social Psychology: Beyond Attitudes and Behaviour*. London: Sage.
- Putnam, H. (2004) *Ethics without Ontology*. Cambridge, MA.: Harvard University Press.
- Quin, Federico, Weyns, Danny, Bamelis, Thomas, Sarpreet, Singh Buttar & Michiels, Sam (2019) Efficient analysis of large adaptation spaces in self-adaptive systems using machine learning. *ICSE Workshop on Software Engineering for Adaptive and Self-Managing Systems*. [Online] 2019-May1–12.
- Quine, W. & Ullian, J. (1979) *The Web of Belief*. New York: Random House.
- Raelin, Joseph A. & Coghlan, David (2006) Developing managers as learners and

- researchers: Using action learning and action research. *Journal of Management Education*. [Online] 30 (5), 670–689.
- Rahman, Md Saidur, Rivera, Emilio, Khomh, Foutse, Guéhéneuc, Yann-Gaël & Lehnert, Bernd (2019) *Machine Learning Software Engineering in Practice: An Industrial Case Study*. 1–21. [online]. Available from: <http://arxiv.org/abs/1906.07154>.
- Rajan, Zeeshan (2016) *Robotic Process Automation (RPA) within Danske Bank*. [online]. Available from: [https://www2.deloitte.com/content/dam/Deloitte/dk/Documents/finance/Finans Agenda2016/Zeehan_Rajan_Danske_Bank_Robotic Process Automation.pdf](https://www2.deloitte.com/content/dam/Deloitte/dk/Documents/finance/Finans Agenda2016/Zeehan_Rajan_Danske_Bank_Robotic_Process_Automation.pdf).
- Reason, Peter & Bradbury, Hilary (2008) *The SAGE Handbook of Action Research: Participative Inquiry and Practice*. London: Sage Publications Inc.
- Rochkind, Marc J. (2004) *Advanced Unix Programming*. Second. Addison-Wesley.
- Rollins, John (2015) Why we need a methodology for data science. *IBM*. August (24), . [online]. Available from: <https://www.ibmbigdatahub.com/blog/why-we-need-methodology-data-science>.
- Roth, Jonas, Shani, Abraham B. & Leary, Myleen M. (2007) Insider action research: Facing the challenges of new capability development within a biopharma company. *Action Research*. [Online] 5 (1), 41–60.
- Russell, S. J. & Norvig, P. (2016) *Artificial intelligence: a modern approach*. Malaysia: Pearson Education Limited.
- Russell, Stuart & Norvig, Peter (1994) *Artificial Intelligence: A Modern Approach*. 3rd edition. Prentice Hall.
- Salomo, Sören, Weise, Joachim & Gemünden, Hans Georg (2007) NPD planning activities and innovation performance: The mediating role of process management and the moderating effect of product innovativeness. *Journal of Product Innovation Management*. [Online] 24 (4), 285–302.
- Samuel, A. L. (1959) Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*. [Online] 3 (3), 210–229.
- SAS (1997) Data Mining Using SAS Enterprise Miner Software. *SAS Institute*.
- SCAI (2019) The National Artificial Research and Development Strategic Plan. National Science and Technology Council (June). [online]. Available from: www.nitrd.gov.
- Schelter, Sebastian, Biessmann, Felix, Januschowski, Tim, Salinas, David, Seufert, Stephan & Szarvas, Gyuri (2018) On Challenges in Machine Learning Model Management. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering*. 5–13. [online]. Available from: <http://sites.computer.org/debull/A18dec/p5.pdf>.
- Schon, David (1983) *The Reflective Practitioner*. New York: Basic Books.
- Schön, Eva Maria, Thomaschewski, Jörg & Escalona, María José (2017) Agile

- Requirements Engineering: A systematic literature review. *Computer Standards and Interfaces*. [Online] 4979–91. [online]. Available from: <http://dx.doi.org/10.1016/j.csi.2016.08.011>.
- Schwab, Klaus (2015) The Fourth Industrial Revolution: What it means and how to respond. Foreign Affairs [online]. Available from: <https://www.foreignaffairs.com/articles/2015-12-12/fourth-industrial-revolution>.
- Schwab, Klaus (2016) *The Fourth Industrial Revolution*. World Economic Forum (ed.). Geneva: .
- Shepherd, Charles & Ahmed, Pervaiz K. (2000) NPD frameworks: A holistic examination. *European Journal of Innovation Management*. [Online] 3 (3), 160–173.
- Shoham, Yoav, Perrault, Raymond, Brynjolfsson, Erik, Clark, Jack, Manyika, James, Niebles, Juan Carlos, Lyons, Terah, Etchemendy, John, Grosz, Barbara & Bauer, Zoe (2018) *AI Index 2018 Annual Report*. [online]. Available from: [http://cdn.aiindex.org/2018/AI Index 2018 Annual Report.pdf](http://cdn.aiindex.org/2018/AI%20Index%202018%20Annual%20Report.pdf).
- Siau, Keng (2018) *Building Trust in Artificial Intelligence, Machine Learning, and Robotics Supply Chain Management View project*. (March), . [online]. Available from: www.cutter.com.
- Stacey, Ralph. D. .. (2011) *Strategic Management and Organisational Dynamics: The challenge of complexity to ways of thinking about organisations*. 6/e XML Vi. Pearson Learning Solutions. VitalBook file.
- Starkey, Ken & Madan, Paula (2001) *Introduction: the challenge of relevance*. 12. [online]. Available from: <https://s3.amazonaws.com/academia.edu.documents/48240614/1467-8551.12.s1.220160822-21549-1e2jhx6.pdf?AWSAccessKeyId=AKIAIWOWYYGZ2Y53UL3A&Expires=1522757589&Signature=6bc6PhcP8x4BDje1Ga%2BwQ9m%2FtwU%3D&response-content-disposition=inline%3Bfilename%3DBrid>.
- Stein, Mari Klara, Galliers, Robert D. & Markus, M. Lynne (2013) Towards an understanding of identity and technology in the workplace. *Journal of Information Technology*. [Online] 28 (3), 167–182.
- Stoica, Marian, Ghilic-Micu, Bogdan, Mircea, Marinela & Uscatu, Cristian (2016) Analyzing Agile Development – from Waterfall Style to Scrumban. *Informatica Economica*. [Online] 20 (4/2016), 15–25.
- Straub, Kayla M., McGwier, Robert W. & A Beex R Michael Buehrer Bert Huang, Chair A. (2016) *Data Mining Academic Emails to Model Employee Behaviors and Analyze Organizational Structure*. [online]. Available from: https://vtechworks.lib.vt.edu/bitstream/handle/10919/71320/Straub_KM_T_2016.pdf;sequence=1.
- Studer, Stefan, Bui, Thanh Binh, Drescher, Christian, Hanuschkin, Alexander, Winkler, Ludwig, Peters, Steven & Mueller, Klaus-Robert (2020) *Towards CRISP-ML(Q): A*

- Machine Learning Process Model with Quality Assurance Methodology*. [online]. Available from: <http://arxiv.org/abs/2003.05155>.
- Surendra, Nanda C. (2008) Using an ethnographic process to conduct requirements analysis for agile systems development. *Information Technology and Management*. [Online] 9 (1), 55–69.
- Thew, Sarah & Sutcliffe, Alistair (2018) Value-based requirements engineering: method and experience. *Requirements Engineering*. [Online] 23 (4), 443–464.
- Tonn, Bruce & Stiefel, Dorian (2012) The future of governance and the use of advanced information technologies. *Futures*. [Online] 44 (9), 812–822. [online]. Available from: <http://dx.doi.org/10.1016/j.futures.2012.07.004>.
- Torbert, William. R. (2001) The Practice of Action Inquiry. *Handbook of Action Research*. (January 2001), 250–260.
- Törnebohm, H. (1976) ‘Inquiring Systems and Paradigms’, in R.S. Cohen, P.K. Feyerabend, & M.W. Wartofsky (eds.) *Essays in Memory of Imre Lakatos*. Boston Studies in the Philosophy of Science. Dordrecht: Springer. p.
- Trondsen, Marianne & Sandaunet, Anne Grete (2009) The dual role of the action researcher. *Evaluation and Program Planning*. [Online] 32 (1), 13–20.
- UK (2017) Industrial Strategy: Building a Britain fit for the future. Report on Industrial Strategy
- US Government (2019) *THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN: 2019 UPDATE*. [online]. Available from: <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf>.
- Vieira, Sandra, Pinaya, Walter Hugo Lopez & Mechelli, Andrea (2020) ‘Chapter 1 - Introduction to machine learning’, in Andrea Mechelli & Sandra Vieira (eds.) *Machine Learning, Methods and Applications to Brain Disorders*. Academic Press. pp. 1–20.
- Wan, Zhiyuan, Xia, Xin, Lo, David & Murphy, Gail C. (2019) How does Machine Learning Change Software Development Practices? *IEEE Transactions on Software Engineering*. [Online] 1–1.
- Wang, Yan (2016) *AMPD700 CD Journal* [online]. Available from: <https://sites.google.com/site/ampd700/process-framework>.
- Weick, Karl F. (1995) *Sensemaking in organisations*. Thousand Oaks CA: Sage.
- Wheeler, Gregory (2016) *Machine Epistemology and Big Data*. [online]. Available from: <http://philsci-archive.pitt.edu/12804/1/Machine-Epistemology.pdf>.
- White House (2016a) *Artificial Intelligence, Automation, and the Economy*. [online]. Available from: <https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF>.

- White House (2016b) *Preparing for the future of Artificial Intelligence*. [online]. Available from:
https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.
- Whitehead, A. N. (1925) *Science and the modern world*. Macmillan and Co.
- Wiemer, Hajo, Drowatzky, Lucas & Ihlenfeldt, Steffen (2019) Data Mining Methodology for Engineering Applications (DMME)— A Holistic Extension to CRISP Model. *Applied Sciences*. 9 (2407), .
- Willig, Carla (2013) *Introducing Qualitative Research in Psychology*. 3rd edition. Maidenhead Berkshire: Open University Press, McGraw Hill Education.
- Willis, Verna J. (2010) ‘Action Learning and Action Research’, in Robert L. Dilworth & Yury Boshyk (eds.) *Action Learning and its applications*. New York: . p. 167.
- Wirth, Norbert (2018) Hello marketing, what can artificial intelligence help you with? *International Journal of Market Research*. [Online] 60 (5), 435–438.
- Wirth, Rüdiger & Hipp, Jochen (2000) CRISP-DM : Towards a Standard Process Model for Data Mining. *Proceedings of the Fourth International Conference on the Practical Application of Knowledge Discovery and Data Mining*. [Online] (24959), 29–39.
- Wright, David, Witherick, Dupe & Gordeeva, Marina (2017) The robots are ready. Are you? Untapped advantage in your digital workforce. *Deloitte*. 28. [online]. Available from:
<https://www2.deloitte.com/content/dam/Deloitte/tr/Documents/technology/deloitte-robots-are-ready.pdf>.
- Wright, Scott A. & Schultz, Ainslie E. (2018) The rising tide of artificial intelligence and business automation: Developing an ethical framework. *Business Horizons*. [Online] 61 (6), 823–832. [online]. Available from:
<https://doi.org/10.1016/j.bushor.2018.07.001>.