

# Optical Tau Guidance of Unmanned Aerial Systems

A thesis submitted in accordance with the requirements of the  
University of Liverpool for the degree of Doctor in Philosophy

by

Christopher Mark Dadswell

September 2019

## Abstract

The use of Unmanned Aerial Systems (UAS) or 'drones' as they are more commonly known, has increased dramatically in recent years. Innovations in electrical power systems as well as improvements in aircraft autonomy technologies have driven a significant consumer surge in the use of small UAS. At the same time, there has been a simultaneous large increase in commercial and military use of UAS across a range of sizes and missions.

There is significant evidence that unmanned aircraft experience accidents at a much greater rate than is observed in human-crewed aviation. A number of surveys investigating the general public's perception of drones have also reported a general distrust of the emerging technology, in part due to negative media coverage and high-profile mishaps. Though inadequate reporting of UAS accidents seems to mask the scale of the problem somewhat, analysis of accident rates for military UAS reveals unacceptably high failure rates for their platforms. Common problems include an inability to sense and avoid hazards, unstable communication links, and failures during landing.

New technology solutions are required to overcome the problems outlined above. The application of Tau theory to the guidance function of unmanned aircraft offers one such possibility. At the University of Liverpool, ecological Tau theory has been established as the basis to understand and model pilot behaviour, based on the optical parameter 'Tau', or time-to-contact. Tau guidance is inherently reactive to external obstacles and hazards, does not require any external signals or infrastructure to function, and is well suited to soft landing manoeuvres, which have all been identified as problematic functions for UAS. This thesis describes work carried out to implement and analyse Tau-guided UAS in the context of simulated landing manoeuvres. Landing in difficult or dynamic environments are examined in detail to illustrate that the inherent reactivity of Tau guidance is useful in these demanding situations. Two main scenarios are considered: rotary-wing maritime landings, and fixed-wing landings on uneven ground.

Existing computer vision techniques to measure time-to-contact were identified and analysed. Three previously established techniques for estimating Tau with monocular video cameras were compared: a dimension-tracking method; an optical flow divergence method, and a direct gradient-based method. The direct gradient method was selected as the most reliable and widely applicable of the three approaches and was thoroughly investigated to find the limits of its operation. It was found that, for the camera used in testing, the method is most effective at measuring time-to-contact when the ground truth value is between 10 and 1 seconds. Above 10 seconds time-to-contact, measurements became too noisy to be usefully applied, and close to the contact point high optical flow causes divergence of the estimate. The accuracy of the method is closely linked to the magnitude of optical flow perceived by the camera, and it was found that manipulating camera frame rate and resolution through subsampling is useful to maintain accuracy throughout the manoeuvres. High resolution cameras can be used to extend the accurate measurement interval above 10 seconds, and high frame-rate cameras can be used to extend the range closer to the contact point.

One of the key limitations of the gradient method is that constant image brightness is assumed across video frames. This leads to reduced accuracy of the Tau sensor if the illumination of the visual scene changes, such as when artificial lights are turned on or off, or when shadows cast by a vehicle or cloud encroach on the scene. This thesis reports upon the development of a new extension to the direct gradient method that removes this limitation/assumption. This innovation expands the operational envelope of Tau sensors that use gradient methods by increasing their accuracy when compared to existing methods under changing light conditions.

The enhanced Tau perception algorithm was tested in simulation using rotary- and fixed-wing UAS platforms in a number of different scenarios. The algorithm was implemented to use one simulated monocular camera to measure time-to-contact with an approaching object or surface.

The rotary-wing platforms were tested in a frigate deck landing scenario, where the targeted landing point is the ship deck. Landings were analysed in a range of sea state conditions that develop different amounts of deck motion for the aircraft to contend

with during landing. Tau estimates from the enhanced direct gradient method were used as a feedback control variable to the aircraft autopilot to perform successful landings. Using only the time-to-contact as a control variable, the tau guidance system allowed the simulated aircraft to perform smooth landings with lower touchdown velocities when compared to the more commonly used constant descent rate guidance strategy. For a sea state 4 deck landing, the Tau guidance system delivered an average touchdown velocity of 0.2 m/s over 20 different landings, while the softest landing the constant descent rate controller delivered was 0.5 m/s, with an average of 1.3 m/s over 20 landings. However, it was found that poor Tau estimation performance close to the deck increased touchdown velocity, especially in high sea-state conditions. It was therefore recommended that the optical Tau estimation method be augmented with an ultrasonic sensor for increased time-to-contact measurement accuracy during the last 1 – 2 seconds of the manoeuvre.

The algorithm was sufficiently robust to allow deck landings to be performed in any sea state, so long as the aircraft had sufficient heave control power to deal with the heave demands of the moving deck. A heave dynamics model was combined with analytic expressions for the Tau guide spatial parameters to produce a new tool to predict whether an aircraft would be able to perform a Tau guided landing in a specific sea state. The tool was tested using an SH60B Seahawk helicopter, and correctly predicted that the aircraft would be unable to make a successful deck landing in sea-state 6 or above.

Tau guidance systems were also implemented on a simulation model of a small fixed-wing aircraft, the 3DR Aero, with a virtual optical Tau sensor integrated into the system. Use of optical Tau to control a fixed-wing aircraft has not previously been demonstrated in operation. Conventional landings were performed on both a standard runway and unprepared sites with changing terrain elevation. Again, it was found that the aircraft was able to consistently make smooth landings using only Tau as a control variable while using a 'perfect' tau measurement, regardless of landing surface. However, when using tau estimates from the virtual tau sensor, tracking performance was too poor for the Tau controller to reliably deliver a safe landing.

Further investigation of fixed-wing implementation of optical Tau systems is required.

Overall, the research showed that Tau guidance can be usefully and practically applied to address some of the problems identified for UAS operations, chiefly the ability to react to dynamic obstacles and hazards in the environment. This was demonstrated by the landing manoeuvres of a rotary-wing UAS onto a moving ship deck. The system also delivered reduced touchdown velocity in comparison to a common alternative guidance system, reducing the risk of aircraft damage. Similar benefits were exhibited for fixed-wing landings on unknown terrain. The key enabler for these benefits was the development of a new variant of the direct gradient method for Tau perception which expanded the operational envelope of the algorithm.

## Acknowledgements

This thesis would not have been possible without the help of many people. Thanks must first go to my two supervisors, Dr Mike Jump and Dr Mark White who have been a source of constant inspiration. Their ideas, support and mentorship have been invaluable throughout my time at Liverpool.

I must also thank the many colleagues and fellow PhD students that have worked alongside me for the last five years, who helped solved countless problems within my work, and kept me sane outside it.

Special thanks must also go to the members of the Liverpool Raptors American Football team, who probably slowed the delivery of this thesis down by some years, but made it substantially more enjoyable.

Finally, I would like to thank my truly incredible parents Alison and Adrian, and my sister Katherine. I cannot envision any world where I would have got this far without your love and support.

## Nomenclature

### Upper Case

$A$	X direction velocity parameter group for gradient method	$B$	Y direction velocity parameter group for gradient method
$C$	Inverse of time-to-contact (1/s)	$D$	Velocity parameter group for gradient method
$D_\tau$	Constant of integration	$F$	Slope parameter group for gradient method
$G$	Radial gradient	$I$	Image matrix
$I_\omega$	Rotational velocity parameter group for gradient method	$I_x$	Image gradient w.r.t the x direction
$I_y$	Image gradient w.r.t the y direction	$I_t$	Image gradient w.r.t the time direction
$J$	Rotational velocity components in x direction for gradient method	$K$	Rotational velocity components in y direction for gradient method
$K_C$	Crossfeed control gain	$K_D$	Differential control gain
$K_I$	Integral control gain	$K_p$	Proportional control gain
$M$	Multiplicative brightness change matrix	$M_{intrinsic}$	Camera intrinsic parameters matrix
$M_{extrinsic}$	Camera extrinsic parameters matrix	$P$	[x, y, z] co-ordinates of a point in the world (m)
$P_s$	X direction slope parameter group for direct gradient method	$Q_s$	Y direction slope parameter group for direct gradient method
$R$	Rotation matrix	$T$	Rotor thrust (N)

$T_f$	Time period of video frame (s)	$T_g$	Manoeuvre duration (s)
$T_\infty$	Rotor thrust at an infinite distance from the ground plane (N)	$U$	Camera or aircraft body velocity in the X direction (m/s)
$V$	Camera or aircraft body velocity in the Y direction (m/s)	$W$	Camera or aircraft body velocity in the Z direction (m/s)
$X$	Displacement in the inertial frame x direction (m)	$X_0$	Initial displacement in the X axis (m)
$X_M$	Displacement along a generic motion gap (m)	$Y$	Displacement in the inertial frame y direction (m)
$Y_0$	Initial displacement in the Y axis (m)	$Z$	Displacement in the inertial frame z direction (m)
$Z_w$	Heave force stability derivative w.r.t vertical speed	$Z_{\delta_C}$	Heave force stability derivative w.r.t collective deflection

### Lower Case

$c_t$	Image brightness offset field	$f$	Camera focal length (m)
$k$	Coupling constant	$m_t$	Image brightness multiplicative field
$o_x$	Image origin x coordinate	$o_y$	Image origin y coordinate
$p$	[x, y, z] coordinates of a point projected onto the image plane	$p_s$	Surface slope in x direction
$q_G$	Electrical charge	$q_s$	Surface slope in y direction



$s_x$	Camera sensor cell size in x direction (m)	$s_y$	Camera sensor cell size in y direction (m)
$t$	Time (s)	$u$	Optical flow in the x direction (m/s)
$v$	Optical flow in the y direction (m/s)	$w$	Aircraft heave velocity (ft/s)
$w_d$	Ship deck heave velocity (ft/s)	$w_g$	Wind gust velocity in the z direction (ft/s)
$x$	Position in the image frame in x direction (m)	$x_{camera}$	Position in the camera frame in x direction (m)
$x_{image}$	Position in the image frame in x direction (m)	$x_{px}$	Image pixel size in the x direction (m)
$x_{res}$	Image resolution in the x direction (pixels)	$y$	Position in the image frame in y direction (m)
$y_{camera}$	Position in the camera frame in y direction (m)	$y_{image}$	Position in the image frame in y direction (m)
$y_{px}$	Image pixel size in the y direction (m)	$y_{res}$	Image resolution in the y direction (pixels)

## Symbols

$\tau$	Time-to-contact (s)	$\tau_0$	Initial time-to-contact (s)
$\tau_X$	Time-to-contact in X direction motion gap (s)	$\tau_Y$	Time-to-contact in Y direction motion gap (s)
$\tau_m$	Time-to-contact in generic motion gap (s)	$\tau_{ref}$	Reference time-to-contact (s)
$\tau_{eRMS}$	Root Mean Square time-to-contact error	$\nabla$	Divergence of vector field
$\omega_x$	Rotational velocity about x axis (rad/s)	$\omega_y$	Rotational velocity about y axis (rad/s)

$\omega_z$	Rotational velocity about z axis (rad/s)	$\delta_c$	Collective deflection
------------	--	------------	-----------------------

### Miscellaneous

$\dot{a} = \frac{da}{dt}$  First derivative w.r.t time

$\ddot{a} = \frac{d^2a}{dt^2}$  Second derivative w.r.t time

$\dddot{a} = \frac{d^3a}{dt^3}$  Third Derivative w.r.t time

$\sum M$   $\sum_{i=1}^{x_{res}} \sum_{j=1}^{y_{res}} M_{i,j}$   
Where  $M_{i,j}$  is any matrix

## Acronyms

AFCS	Automatic Flight Control System
ALFURS	Autonomy Levels For Unmanned Rotorcraft Systems
ALFUS	Autonomy Levels For Unmanned Systems
CAA	Civil Aviation Authority
CAD	Computer Aided Design
CAT	Commercial Air Transportation
CCD	Charge-Coupled Device
CFD	Computational Fluid Dynamics
CMOS	Complimentary Metal-Oxide-Semiconductor
CORSE	Centre Only Relative State Estimation
CPU	Central Processing Unit
CRM	Crew Resource Management
CV	Computer Vision
DERA	Defence Evaluation and Research Agency
DfT	Department for Transport
DMPC	Decentralized Model predictive Control
DSTG	Defence Science and Technology Group
DVE	Degraded Visual Environment
EASA	European Union Aviation Safety Agency
EC	Environmental Complexity
ES	External System

ESI	External System Independence
EU	European Union
FAA	Federal Aviation Authority
FGR	FLIGHTLAB Generic Rotorcraft
FoV	Field of View
GA	General Atomics
GCS	Ground Control Station
GNC	Guidance, Navigation and Control
GPS	Global Positioning System
GPU	Graphics Processing Unit
GVE	Good Visual Environment
HALE	High Altitude Long Endurance
HD	High Definition
HVGA	Half Video Graphics Array
INS	Inertial Navigational System
MALE	Medium Altitude Long Endurance
MAV	Micro Aerial Vehicle
MC	Mission Complexity
MoD	Ministry of Defence
MTE	Mission Task Element
MTOW	Maximum Take-Off Weight
NASA	National Aeronautics and Space Administration
NATO	North Atlantic Treaty Organisation

PID	Proportional Integral Derivative
RAF	Royal Air Force
RAM	Random Access Memory
RGB	Red Green Blue
RMS	Root Mean Square
RPA	Remotely Piloted Aircraft
RPAS	Remotely Piloted Aircraft System
RT	Real Time
RUAS	Rotorcraft Unmanned Aerial System
SA	Situational Awareness
SARSE	Subtended Angle Relative State Estimation
SFM	Structure From Motion
SLAD	Safe Landing Area Detection
SLAM	Simultaneous Localisation and Mapping
SUAS	Small Unmanned Aerial System
SURF	Speeded Up Robust Features
TALS	Tactical Automatic Landing System
TRC	Translational Rate Command
TRN	Terrain Relative Navigation
UA	Unmanned Aircraft
UAS	Unmanned Aerial System
UAV	Unmanned Aerial Vehicle
UCARS	UAS Common Automatic Recovery System

UCAV	Unmanned Combat Aerial Vehicle
UK	United Kingdom
US	United States
USAF	United States Air Force
USPS	United States Postal Service
USS	United States Ship
VGA	Video Graphics Array
VR	Virtual Reality
XML	eXtensible Markup Language

# Contents

Abstract .....	i
Acknowledgements.....	v
Nomenclature .....	vi
Acronyms.....	x
Contents .....	xiv
1 Chapter 1 – Introduction .....	1
1.1 Current and Future Usage of Unmanned Aerial Systems .....	1
1.2 Unmanned Aircraft Safety.....	4
1.3 Biological Guidance Mechanisms as Autonomous Solutions.....	6
1.4 Tau Theory in Autonomous Systems.....	7
1.5 Research Aims and Objectives .....	9
1.6 Thesis Scope, Structure and Content .....	10
1.6.1 Thesis Scope .....	10
1.6.2 Thesis Structure and Content.....	11
1.7 Original Contribution to Learning.....	12
2 Chapter 2 – Technical Review .....	15
2.1 Overview of UAS Technology .....	15
2.1.1 UAS Terminology.....	15
2.1.2 UAS Autonomy .....	17
2.1.3 State of the Art UAS Guidance, Navigation and Control .....	20
2.1.4 UAS Accident Statistics.....	27
2.2 Tau Theory.....	33
2.2.1 Visual Perception .....	33
2.2.2 Tau Theory Guidance .....	34

2.2.3	Constant Rate of Change of Tau Strategy .....	36
2.2.4	Tau Coupling.....	39
2.2.5	Intrinsic Tau Guidance.....	41
2.2.6	Alternatives to Tau Theory.....	48
2.3	Tau in Aviation.....	49
2.4	Tau Perception .....	50
2.4.1	Spatial Tau Perception Methods.....	50
2.4.2	Optical Tau Perception.....	53
2.5	Tau in Guidance and Control.....	60
2.6	Technical Review Summary.....	66
3	Chapter 3 – Experimental Setup.....	67
3.1	Simulation Software and Hardware .....	67
3.2	Aircraft Simulation Models.....	70
3.2.1	3DR X8 .....	70
3.2.2	3DR Aero .....	74
3.2.3	RUAS Linear Models.....	76
3.3	Camera Modelling .....	78
3.3.1	Camera Modelling.....	79
3.3.2	Camera Model Parameters .....	81
3.3.3	Virtual Helipad Landing Scene .....	82
3.3.4	Virtual Runway Landing Scene.....	83
3.4	Matlab Simulink Virtual Reality Worlds .....	85
3.4.1	Frigate Deck Landing Scene .....	86
3.4.2	Desert Runway Landing Scene.....	87
3.5	Mission Task Elements .....	88
3.5.1	Frigate Deck Landing.....	88



3.5.2	Runway Landing .....	90
4	Chapter 4 – Tau Perception through Computer Vision .....	91
4.1	Tau Perception through Dimension Tracking.....	92
4.1.1	Overview .....	92
4.1.2	Helipad Landing Scene – Tau Analysis with Dimension Tracking .....	95
4.1.3	Runway Landing Scene – Tau Analysis with Size Method.....	99
4.2	Tau Perception through Optical Flow Divergence .....	101
4.2.1	Overview .....	101
4.2.2	Helipad Landing Scene – Tau Analysis with Optical Flow Divergence Method 102	
4.2.3	Runway Landing Scene – Tau Analysis with Optical Flow Divergence Method 107	
4.3	Tau Perception through Direct Gradient Method.....	109
4.3.1	Overview .....	109
4.3.2	Rotary-Wing Platform Analysis .....	113
4.3.3	Factors Affecting Tau Estimate Accuracy.....	118
4.3.4	Simulink Implementation.....	131
4.3.5	Fixed-Wing Platform Analysis .....	136
4.4	Direct Gradient Method Extension for Brightness Changes .....	145
4.4.1	Overview .....	145
4.4.2	Derivation.....	146
4.4.3	Testing.....	150
5	Chapter 5 –Tau Guided UAS Landings .....	156
5.1	Rotary Wing Tau Guidance System Implementation .....	156
5.1.1	Tau Control Law Analysis .....	157
5.2	Frigate Deck Tau Guided Landings .....	161

5.2.1	Benefits of Tau Guidance for Deck Landings .....	164
5.2.2	Start Point Variation.....	167
5.2.3	Sea State Variation .....	169
5.3	Tau Manoeuvre Limitations .....	173
5.4	Optical Tau Guided Frigate Deck Landings Results .....	181
5.4.1	Sea State 1.....	181
5.4.2	Higher Sea States .....	185
5.5	Fixed-Wing Tau Guided Landings .....	188
5.5.1	Tau Guidance System Implementation.....	188
5.5.2	Flat Runway Tau Landings.....	189
5.5.3	Landings at Unprepared Sites .....	190
6	Chapter 6 – Conclusions .....	194
6.1	Conclusions of the Research .....	194
6.1.1	Tau Perception .....	194
6.1.2	Tau Guidance of Rotary UAS .....	197
6.1.3	Tau Guidance of Fixed-Wing UAS .....	198
6.1.4	Overall Conclusions.....	198
6.2	Potential Further Work .....	199
	References.....	226
	Appendix A – Aircraft Linear State-Space Models.....	200
	Appendix B – Tau Perception Derivations.....	203
	Appendix C – Simulation Model Flight Controller Systems.....	221

# 1 Chapter 1 – Introduction

## 1.1 Current and Future Usage of Unmanned Aerial Systems

There are increasing signs that the aviation industry is on the brink of a significant upheaval. Innovations in electrical power systems have driven a surge in the use of Small Unmanned Aerial Systems (SUAS) that has democratised aviation in an unprecedented way. This rapid increase in unmanned aircraft use seems set to continue. On December 14<sup>th</sup>, 2015, the Federal Aviation Administration of the United States issued a new rule requiring all Unmanned Aerial Systems (UAS) weighing between 250 grams and 25 kilograms to be registered using a new online system and registration became a legal requirement one week later.

By the end of December 2016, a year later, 626,245 aircraft owner-hobbyists in the US had registered one or more aircraft with the FAA, with an estimated 1.1 million hobbyist aircraft units in the country [1]. By the end of 2018, that number of registered hobbyists had grown to over 900,000 [2]. For comparison, the FAA reported that there were 633,317 certificated civil pilots in the US in 2018, split across various license types [3]. Evidently, the number of unmanned ‘pilots’ has already outstripped the sum of their manned equivalents in the United States. This is perhaps unsurprising given that there are no certification requirements for hobbyist UAS pilots, but it is indicative of this paradigm shift in the aviation community.

A similar online registration system for commercial UAS operators was introduced in April 2016 and is also reported on in reference [1]. By early 2017, 44,000 businesses had registered unmanned aircraft with the FAA. The organisation forecast huge growth in this sector, with over 400,000 commercial operators predicted to register by 2021. The FAA report also categorised the various sectors these commercial users occupy, shown in Figure 1-1. Aerial photography accounts for over a third of the market but applications are clearly varied, speaking to the versatility of UAS platforms.

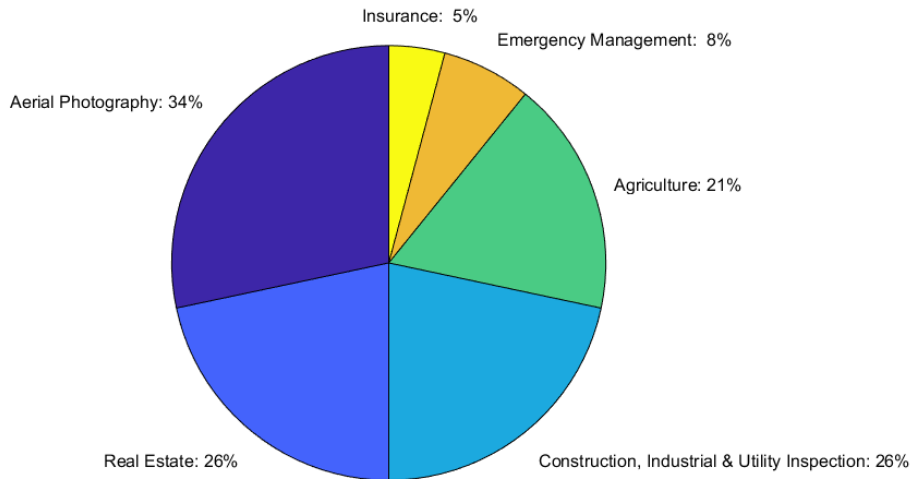
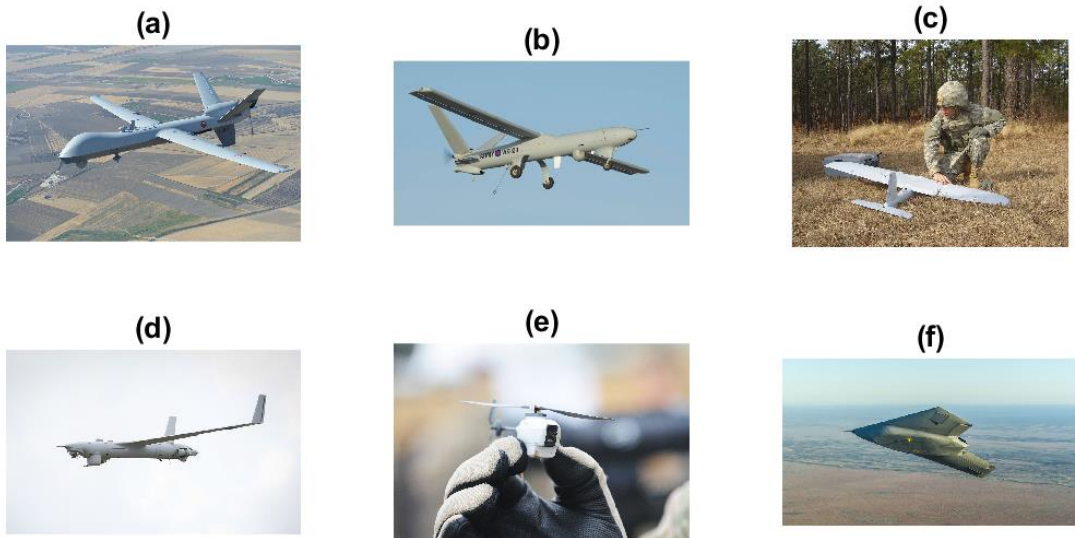


Figure 1-1: Usage of commercial UAS platforms in the US in 2016 [1]

Unmanned aircraft are set to fill yet more commercial roles with their well-publicised integration into logistics and distribution networks, likely the most public-facing application for them to date. Amazon first announced their Prime Air service in 2013 and with it their intention to begin delivering packages within 30 minutes of ordering [4]. The service delivered its first package to an address in Cambridge via quadcopter in 2016, but the service is still yet to see an official rollout [5]. Post offices and parcel services around the world have stated their intentions to offer similar services [6], along with food delivery companies such as Domino’s Pizza [7]. Google Wing became the first company to obtain an air operator’s certificate from the FAA in April 2019 to deliver takeaway food and beverages by air, and are conducting a pilot program in Queensland, Australia [8]. However, years of development by even these most powerful of corporations have thus far failed to materialise into usable services.

The use of UAS platforms by military services also continues to grow. The United Kingdom armed forces reported to operate 585 unmanned aircraft in 2015, comprised of five different platforms, shown in Figure 1-2. Reference [9] describes the aircraft and their roles. They include full-sized aircraft like the General Atomics MQ-9 Reaper (a) and Thales Watchkeeper (b) operated by the Royal Air Force (RAF), as well as smaller man-portable aircraft like the Desert Hawk III (c), Boeing Scan Eagle (d) and Black Hornet Nano (e) used by other services.



*Figure 1-2: Current and future UAS operated by the British Armed Forces: (a) GA MQ-9 Reaper [10] (b) Thales Watchkeeper [11] (c) Lockheed Martin Desert Hawk III [12] (d) Boeing ScanEagle [13] (e) Prox Dynamics Black Hornet Nano [14] (f) BAE Taranis [15]*

All operate in surveillance and reconnaissance roles except the Reaper, which is the only Unmanned Combat Aerial Vehicle (UCAV) among them. Watchkeeper and Reaper are both full-scale, runway-launched aircraft, while the Desert Hawk and Black Hornet can both be hand-launched and recovered in the field from unprepared sites. The Scan Eagle is the only UA operated by the Royal Navy and is launched from a ship-borne catapult. It is also recovered using a ship-borne net system. Procurement of further unmanned aircraft is also planned, with the purchase of 20 Protector UAVs announced in the 2015 Strategic Defence and Security Review [16]. A program to acquire a derivative of the BAE Systems Taranis UCAV (f) has also been announced, set to enter service in the 2030s [17]. The US Department of Defence (DoD) outlines its approach to unmanned systems use and development in reference [18]. Their spending on UAS has increased year on year, and over \$4.2 billion was spent on UAS in 2017 across various branches of military.

Clearly there is a large market for UA in both the civil and military domains, and many planned applications require flight in urban, dynamic or cluttered environments. However, public desire for these kinds of operations remains an open question. Public perception of unmanned aircraft, or ‘drones’ as they are commonly known in the press, is generally negative. Reference [19] details a public dialogue conducted

by the UK Department for Transport (DfT) through a series of workshops. It was found that participants generally had low awareness and knowledge of drones, with the knowledge they did have focussed on military hardware and small consumer quadcopters. They raised concerns about the safe use of drones in public spaces, the quality of the aircraft themselves and potential for technological failures. However, as the workshops progressed and participants learned more about the commercial uses of drones, current legislation and safety procedures, they became more positive and engaged with potential benefits for citizens.

Similar opinions were also found to be prevalent in a survey by the United States Postal Service (USPS) on the public perception of drone delivery services [20]. Though 75% of respondents expected drone delivery to be a reality by 2021, 46% said that potential malfunction was their primary concern about such services, and only 32% believed that the system would be safe. These statistics raise the obvious question: how safe are unmanned aircraft in reality?

## 1.2 Unmanned Aircraft Safety

Air travel is often referred to as one of the safest modes of transport available, and this is borne out by the literature. Reference [21] presents an analysis of fatal accident statistics for various modes of transport for the year 2000, demonstrating that by the measure of fatalities per distance travelled, air travel was ten times safer than rail transport, and over 2000 times safer than road travel. However, the safety of unmanned aircraft is a more open question; slow development of legislative frameworks has resulted in lower reporting of incidents and accidents involving UAS. Accidents in manned aviation will always carry a greater risk to life by virtue of carrying human cargo; the potential for fatal accidents involving UAS is limited to mid-air-collision or ground injury. As a result, a smaller portion of accidents involving UAS lead to any human injury, and as such are often not examined or documented in the same exhaustive detail of traditional air accident investigations.

Though accident data relating to consumer and hobbyist drone use is scarce, recent high-profile incidents have damaged the public perception of UAS. A prospective parcel delivery service, operated by Swiss Post, was forced to suspend its services

indefinitely in early 2019 after two of its vehicles crashed [22]. Swiss Post operate quadcopters produced by Matternet that can carry a 2 kg payload up to 10 km, with a total aircraft weight of up to 10 kg. The system is designed to deploy a parachute if a failure occurs in flight to slow the fall of the vehicle and avoid injury to anyone on the ground. The first crash occurred after a GPS failure. The backup parachute deployed and the vehicle landed in a body of water safely. However, after a second failure due to an unidentified fault, the tether that connected the backup parachute to the vehicle broke, and the aircraft suffered an uncontrolled crash. It crashed within 50 m of a group of children playing, though no one was hurt.

Accidents involving UAS in use by military forces have generally been better documented. One of the earliest investigations of UAS accident statistics was published by the Washington Post in reference [23]. After manually searching through thousands of declassified documents, they found evidence of over 400 major accidents involving large US military drones between 2001 and 2014. Of the 269 MQ-1 Predator aircraft the United States Air Force (USAF) acquired during this period, 48% had been involved in a major accident. Between 2009 and 2014, the mishap rate per 100,000 hours of flight for the Predator dropped by two thirds, from 13.7 to 4.79, but this still represents an alarming failure rate. Similar aircraft, like the MQ-9 Reaper and RQ-2 Pioneer (Figure 1-3), suffered similar accident rates. The Washington Post's analysis identified several fundamental problems that contributed to the high failure rate in USAFs remotely piloted aircraft:

1. Persistent mechanical and electrical defects due to insufficient testing before being brought into service
2. Unreliable communication links
3. Pilot errors, particularly during landings
4. Limited ability to detect and avoid obstacles



*Figure 1-3: The aftermath of a collision between an RQ-2 Pioneer aircraft and a C-130 Hercules [24]*

This poor safety record for military UAS is supported by an FAA report analysing similar military drone accidents that suggested that the mishap rate for UAVs was 100 times higher than for manned aircraft [25]. It also highlighted the landing phase as especially problematic for UAS due to lower situational awareness caused by the removal of the pilot. The FAA report further emphasised that approximately 50% of these were attributed to aircraft failures associated with relaxed design methods and system reliability.

Evidently, there are still significant hurdles related to UAS capability and reliability that must be overcome to earn the trust of the general public and drive higher adoption of UAS. There is a clear gap between accident rates in manned and unmanned aviation, so technologies must be developed to close this gap. Since humans still seem to be better suited to flight control than machines (at least in some situations), it is useful to examine the mechanisms they employ for control and guidance.

### 1.3 Biological Guidance Mechanisms as Autonomous Solutions

The natural world has often been used as inspiration for the development of technology, and autonomous guidance is no exception. After all, animals, including humans, can successfully guide themselves through their environments autonomously in a seemingly effortless manner, so it stands to reason that emulation of these behaviours would be effective for autonomous systems.



One nature-inspired avenue of investigation for guiding motion has used Tau Theory as its basis. Tau Theory, proposed by David Lee in 1976 [26], offers an ecological mechanism for the guidance of purposeful motion through the cluttered environment of the Earth's surface. By examining motorist's control strategies during braking manoeuvres, Lee found that drivers modulated their brake inputs to keep a constant rate of change of instantaneous time-to-contact with their desired stopping point. This strategy allowed for smooth deceleration to a stop so that relative position and velocity both closed to zero simultaneously. Tau was posited to be perceived directly through the driver's visual system, i.e. its perception did not require any cognitive processing. Existing research suggests that  $\tau$  is indeed a fundamental control variable responsible for the guidance of many actions, including for flight control [27]. For example, ongoing work at the University of Liverpool has established the strong role that  $\tau$  plays in manned flight: in the landing flare manoeuvre for fixed-wing aircraft [28], rotorcraft decelerations and low-level flying [29].

Tau theory offers a simple and widely applicable mechanism for both obstacle avoidance and interception, actions that UAS struggle to perform according to presented evidence. Perception of Tau seems an intuitive way of increasing situational awareness in autonomous systems and allows for manoeuvres to be performed without the need for external command or reference, reducing the communication link reliance also cited as a problem area for UAS. The work presented in this thesis will explore the use of Tau-based control as a tool for increasing autonomy in UAS.

#### 1.4 Tau Theory in Autonomous Systems

The use of Tau Theory for flight control has also gained some traction in the unmanned aviation arena already. Reference [30] presents a hardware-based study into the use of Tau for the control of a SUAS for various 'docking' manoeuvres, which includes landing, where the 'docking' is, in that case, with the ground. The study shows the utility of Tau for such flight manoeuvres but relies upon the use of the onboard Inertial Navigation Systems (INS) and Global Positioning Systems (GPS). While this is a valid approach (such positioning systems are likely to be present on all

forthcoming new air vehicles), these systems are prone to errors and can be spoofed: false signals can be used to trick GPS units into providing erroneous position estimates [31]. In the case of GPS, it simply may not be available in some urban or indoor environments. Reference [32] reports on a simulation study of the vision-based detection and tracking of a moving ground-based rover by a quadrotor UAV followed by a landing of the latter on the former using tau-based guidance. Again, good results are reported, but the study is limited in the sense that the vision-based element of the experiments rely on AprilTags, a template detection system [33], to provide the required visual information.

While it cannot be claimed that vision-based systems in the natural world cannot be spoofed or fooled – if this were the case, optical illusions would not exist – it is argued that the evolutionary process has resulted in a robust and yet simple means to provide reliable guidance in cluttered environments. Similarly, whilst it is recognized that some texture/structure is required to make sense of a visual scene (a white-out in a snowy landscape might be an example of where these are not present), it is clear that observers in the natural world do not require specific markers in their visual field to be able to perceive their motion. As such, the work reported in this thesis builds on these prior works by showing how Tau Theory can be used to guide aerial vehicle motion using vision-based sensors without the need for any special or ‘abnormal’ elements in the visual scene.

Time-to-contact can be easily estimated using active ranging sensors such as radar, Lidar or sonar. However, the hardware required can be expensive, heavy, or possess limited accuracy at range. Passive imaging sensors are preferable in this application due to their availability and obvious analogue to the natural vision mechanisms used by humans and animals. Several image processing methods exist that yield time-to-contact estimates using only monocular vision, but no comprehensive review of their individual benefits and drawbacks has been found in the literature. This thesis will analyse Tau perception mechanisms for aerial sensor platforms.

Work on Tau-guided autonomous systems has also been exclusively limited to rotary-wing platforms such as quadcopters thus far. These vehicles are prevalent in the consumer-hobbyist UAS arena, but military UAS usage tends to be concentrated on

fixed-wing platforms. No data has been found to analyse the platform composition of the commercial drone market, but several prospective drone delivery services have presented fixed-wing or hybrid systems, such as Amazon Prime Air [5] and Google Wing [8]. As a result, investigation of Tau guidance of fixed-wing UAS is required to facilitate the use of Tau as a control mechanism throughout the unmanned aviation arena.

## 1.5 Research Aims and Objectives

The application of Tau theory-based control to both manned and unmanned aircraft is not a new concept, but implementations found in literature leave a number of research questions unanswered. The work presented in this thesis aims to address some of those questions, which are detailed below:

1. Can optical Tau-based navigation guidance, navigation and control systems be used to address the common problems identified for UAS? Current Tau implementations have failed to formalise their benefits in the context of the problems faced in the wider UAS community, identified through accident reporting.
2. What is the most effective method for perceiving time-to-contact for Tau control? The vast majority of Tau-controlled systems detailed in literature either leverage spatial state information from GPS-INS systems, or rely on visual template-matching systems to localise themselves relative to a target to estimate time-to-contact, neutralising some of the benefits of Tau control.
3. Is optical-Tau guidance equally useful for fixed- and rotary-wing unmanned aircraft? No examples of fixed-wing implementation of Tau-control have been found in hardware, though research suggests that Tau is applied by pilots across both platforms.

A number of associated project objectives were established to explore the above research questions:

1. To identify and analyse currently available techniques for time-to-contact perception and explore the limits of their effectiveness. There are already

several techniques for perceiving Tau, but no comparison of them has been found.

2. To create a set of virtual tools to be used for the analysis of Tau perception techniques in simulation.
3. To implement Tau guidance systems on various rotorcraft in simulation, and analyse their performance in landing manoeuvres, with a focus on difficult and dynamic environments.
4. To implement Tau guidance systems on fixed-wing aircraft in simulation and investigate their performance in landing manoeuvres.

## 1.6 Thesis Scope, Structure and Content

At the outset of this thesis, it is useful to describe what will be included within its pages and how it will be structured, as well as what is considered beyond the scope of this document and will not be discussed.

### 1.6.1 Thesis Scope

Ecological Tau theory has been written about extensively since its inception in the 1970s, both in support and opposition of its fundamental hypotheses. Many examples can be found in the literature of human and animal actions that seem to adhere to the temporal and spatial trajectories predicted by Tau theory, as well some that do not. It is not the aim of this thesis to present any evidence for or against the validity of Tau theory, though some of the existing data will be reviewed in due course. Instead, for the purposes of this thesis, Tau theory is viewed as a paradigm of biological guidance that can be usefully applied to unmanned systems to effect particular behaviours; the veracity of the theory is irrelevant to whether or not its methods can be usefully applied to autonomous systems.

Tau-guided behaviours can be attributed to a multitude of actions, which include both evasive and interceptive actions. This thesis will focus on applying Tau theory in the context of interceptive action to achieve safe landings of UAS. This phase of flight has been chosen as a focus in response to the evidence in literature that emphasises approach and landing manoeuvres as problematic for UAS. Many of the techniques

described can be equally applied for obstacle avoidance by actuating different controls based on Tau perception.

While this thesis details several problems that have been identified in UAS operations, the work presented is not intended as a comprehensive solution to all of these problems. Instead, this thesis provides analysis and novel implementations of one technique that may be helpful in solving some of these problems. The application of Tau guidance systems is also restricted to single axis heave control throughout. Control of the other aircraft axes is accomplished using PID feedback control.

### 1.6.2 Thesis Structure and Content

This document is divided into six chapters, each with a different and specific focus. Chapter 2 reviews the relevant technical material that forms the basis of the work described in later chapters. First, a short review of the terminology applied to unmanned systems is presented, as well as some established measures of vehicle autonomy, and the current state-of-the-art guidance, navigation and control techniques that can be found on modern unmanned aircraft. A thorough review of ecological Tau theory and the work that led to its inception is given next. Particular focus is given to the various guidance laws that have been proposed as elements of the theory, as these form the basis for the proposed UAS operations. Next, Tau behaviours in manned aviation are explored in both rotary- and fixed-wing contexts. Technologies for the perception of Tau are examined with a focus on computer vision. Appendix A gives more detail on some of these techniques, including derivations for their mathematical equations. Finally, operational examples of Tau-controlled systems are surveyed.

Chapter 3 describes work carried out to create the tools and experimental setups used during the subsequent chapters. A major component of this work was creating simulation models that could be used to test Tau control systems. The development of a SUAS model is described, as well as details of a number of existing models used during the work. Since computer vision techniques for Tau perception were to be investigated, a number of tools were developed to create custom virtual visual scenes for testing image processing techniques. The tools and visual scenes

developed are described here. Finally, several Mission Task Elements (MTEs) are described that were used for the clinical and repeatable testing of Tau guidance and control systems.

Chapter 4 details a comprehensive analysis of several computer vision techniques for Tau perception: dimension tracking, optical flow divergence analysis and direct gradient analysis. They are tested against the virtual visual scenes described in Chapter 3 in both rotary- and fixed-wing contexts. The limitations of each are assessed and a recommendation made for which is the most useful. Techniques for mitigating weaknesses of the technique are also presented and tested. A modification of the direct gradient method for Tau perception is also proposed with results from initial testing.

Chapter 5 provides an analysis of Tau guidance applied to the rotary-wing MTEs detailed in Chapter 3 using simulation. The implementation of the Tau guidance system is detailed with analysis of several control laws. Tau guided landings on ship decks are examined in detail with investigation of various sea states and aircraft. Simulations of Tau landings using optical Tau are presented. Tau guidance is also applied to the fixed-wing MTEs detailed in chapter 3. Fixed wing landings using optical Tau are presented.

Chapter 6 presents the conclusions of the work described in the preceding chapters and some recommendations for how optical-Tau-based control can be applied and further investigated.

## 1.7 Original Contribution to Learning

A PhD thesis must make an original contribution to learning to satisfy the requirements of that degree. An original contribution to learning is defined in reference [34] as any of the following acts:

1. Discovering new knowledge;
2. Explanation or connection of previously established facts;
3. Developing new theories;
4. Revising established views.

The original contributions of the work described in this thesis are described in the points below for the sake of clarity:

1. An investigation and comparison of optical Tau perception techniques is presented in chapter 4 that sheds new light on the limiting factors of a Tau perception technique known as the direct gradient method. Specifically, new information about where estimates provided by the method start to diverge and how this divergence relates to the parameters of the camera used for Tau perception are presented.
2. A new extension to the direct gradient method for Tau perception is presented in chapter 4. It extends the capability of the already established method to allow for changing brightness in the visual scene, a problem which would previously have reduced the accuracy of the Tau estimate. This new method mitigates the error introduced by changing scene brightness and is demonstrated to be effective with virtual reality tools.
3. An implementation of optical Tau perception and guidance on a fixed-wing aircraft platform is described and tested in simulation in chapters 4 and 6. The system uses a downward looking camera to measure time-to-contact while descending towards a runway, and controls the pitch attitude of the aeroplane to effect a smooth flare and touchdown manoeuvre, which has not previously been demonstrated with optical Tau perception.
4. An analytic prediction of whether a rotorcraft will be able to perform Tau guided landings on a frigate deck in high sea states is presented in chapter 5, which is a novel technique. The approach uses an existing heave dynamics model to relate the demands of a defined Tau manoeuvre over a heaving ship deck to the stability derivatives from a linear aircraft model.

These original contributions aim to address some of the UAS issues which have been briefly discussed in chapter 1.2 and further examined in section 2.1.4. Primarily, these contributions establish that optical time-to-contact is a viable control variable for UAS landings in difficult and dynamic environments and provide a mechanism for expanding the range of environments that an optical Tau sensor can operate in effectively. These Tau guidance methods can increase autonomy levels of UAS by

providing a mechanism to improve reactivity of the system to its external environment; time-to-contact is inherently reactive to surrounding terrain and dynamic entities. This capability can be used to improve the reliability of UAS landing manoeuvres, and also be applied in obstacle avoidance, which was identified as one of the problems contributing to high UAS failure rates. Tau guidance does not require any external signals, so can also improve reliability of these functions in the absence of communication links or external navigation systems, which has also been identified as a problem for UAS reliability.



## 2 Chapter 2 – Technical Review

The work presented in this thesis is based on a range of technical fields where much work has already been carried out, including contemporary UAS technology, the techniques currently used for their guidance, navigation and control, and theories of biological guidance and their applications to the aerospace domain. This chapter will examine technical information relevant to these areas to expose the gaps in knowledge that have led to the research described in the rest of this thesis.

### 2.1 Overview of UAS Technology

#### 2.1.1 UAS Terminology

The terms used to classify unmanned aircraft are wide-ranging and often differ between operators, manufacturers and governing bodies. For the purposes of this thesis, it is useful to clarify the meanings of the terminology typically used to describe various aircraft types. It is important to clearly define what is meant by various terms to avoid ambiguity between remotely piloted vehicles and vehicles which are, to some degree, autonomous, as the systems have different legal and operational implications.

Terms such as ‘Unmanned Aerial Vehicle’ or ‘Drone’ are often used in relation to unmanned systems, but these terms are often used non-specifically and can refer to any number of platforms with or without the presence of a remote pilot. The UK Ministry of Defence (MoD) has defined its terms as shown in Table 2-1.

*Table 2-1: UK MoD UA, UAS, RPA, RPAS definitions [35]*

<b>Unmanned Aircraft (UA)</b>	An unmanned aircraft is defined as an aircraft that does not carry a human operator, is operated remotely using varying levels of automated functions, and is normally recoverable.
<b>Unmanned Aircraft System (UAS)</b>	An Unmanned Aircraft System is defined as a system, whose components include the Unmanned Aircraft and all

	equipment, network and personnel necessary to control the Unmanned Aircraft.
<b>Remotely Piloted Aircraft (RPA)</b>	A Remotely Piloted Aircraft is defined as an aircraft that, whilst it does not carry a human operator, is flown by a pilot and is normally recoverable.
<b>Remotely Piloted Aircraft System (RPAS)</b>	A Remotely Piloted Aircraft System is the sum of the components required to deliver the overall capability and includes the pilot, sensor operators (if applicable), Remotely Piloted Aircraft, ground control station, associated manpower and support systems, satellite communications links and data links.

The MoD assigns a class to unmanned aircraft according to their Maximum Take-Off Weight (MTOW), based on a NATO classification system [36]. The different classes are described in Table 2-2, and aircraft will be referred to by their class throughout this thesis.

Table 2-2: MoD & NATO UAS Classification System [36] [37]

<b>Class</b>	<b>Category</b>	<b>MTOW</b>	<b>Operating Altitude</b>
Class I	Nano	< 0.2 kg	< 100 feet
	Micro	0.2 – 2 kg	< 200 feet
	Mini	2 – 20 kg	< 3,000 feet
	Small	20 – 150 kg	< 5,000 feet
Class II	Tactical	150 – 600 kg	< 10,000 feet
Class III	Medium Altitude Long Endurance (MALE)	> 600 kg	< 45,000 feet
	High Altitude Long Endurance (HALE)	> 600 kg	< 65,000 feet

### 2.1.2 UAS Autonomy

The MoD also draws a distinction between automated and autonomous systems; an automated system can be described as one programmed to logically follow a pre-defined set of rules to provide an outcome in response one or more sensor inputs [35]. In this way, the response of an automated system can always be predicted as long as the input is known. On the other hand, autonomous systems are capable of higher-level understanding and can take appropriate action to bring about a desired state on their own. Without human oversight, they can decide on a course of action from a number of alternatives; therefore, individual actions performed by an autonomous system may not always be predictable [35]. This is a relatively simplistic definition that is general to all automated or autonomous systems.

A more detailed approach to autonomy definition for RUAS is described by Kendoul [38] in the form of an eleven-point scale named the Autonomy Levels For Unmanned Rotorcraft Systems (ALFURS) framework. It is an extension of the ALFUS framework for rotorcraft systems specifically, first proposed by Huang [39]. The different levels of this scale are defined by increasing levels of Guidance, Navigation and Control capability. Before discussing the ALFURS framework it is important to define what is meant by these three terms.

Control in the context of UAS can be defined as the process of manipulating the inputs to a dynamical system to obtain a desired effect on the system outputs, with or without a human operator in the loop. In real terms, this translates to an algorithm or control law computing commands for vehicle actuators to affect a change in the vehicles 3D motion. This function is performed by an Automatic Flight Control System (AFCS), sometimes also termed an autopilot [39].

Navigation is described as the process of monitoring and controlling the movement of a vehicle from one place to another through the processes of data acquisition, data analysis and extraction and inference of information about the vehicle and the surrounding environments states. Huang further breaks down the task of navigation into four subcategories: sensing, state estimation, perception and situational

awareness. A sensing system is one that uses one or more devices that respond to a physical stimulus and generate signals that reflect some information about an object or physical phenomenon. Typical sensing systems include gyroscopes, accelerometers, pressure sensors and cameras. State estimation is the processing of raw signals provided by a sensing system to estimate variables related to the vehicle or environment state. State estimation systems are often used to approximate a vehicle's attitude, position or velocity from sensor data. Perception is the ability to use inputs from sensors or state estimation systems to build an internal model of a surrounding environment and recognise entities, events or situations using a priori knowledge. Typical UAS perception functions include environment mapping, obstacle or target detection and recognition. Situational awareness is the comprehension of perceived objects or events and the projection of their status into the near future. In other words, a system that exhibits situational awareness is able to understand objects or events perceived by a perceptual system, and with this understanding can make inferences about their behaviour in the future [39].

A Guidance system exercises planning and decision-making functions to drive a system toward an assigned mission or goal. In the context of a UAS, a guidance system replaces the cognitive process of a human operator, using inputs from a navigation system to make high-level decisions and generate reference trajectories and commands for a control system. Huang defines five typical tasks of a guidance system: trajectory generation, path planning, mission planning, decision making and reasoning [39]. A trajectory generator has the role of computing the different motion functions that are physically possible for a RUAS to complete and providing these as reference trajectories for a flight control system to follow. These trajectories can be pre-programmed, sent to the RUAS in real-time or dynamically generated when required. Path planning is the process of using accumulated navigation data and a priori information to find the best or safest way to reach a goal state. Mission planning is the process of generating tactical goals, a route, command structure, coordination and timing for a UAS. Planning can be performed in advance of a mission or in real-time. Decision making is the ability of a UAS to select a course of action from several alternative scenarios based on available information and analysis. The

decisions made are relevant to achieving mission goals efficiently and safely. Reasoning is the UAS's ability to analyse and reason using contextual associations between different entities. These functions are the highest level of autonomy possible and are equivalent to fully replacing the human cognitive process [39].

The ALFURS scale spans the whole range of possible autonomy levels, from RPAS that are completely controlled by a remote operator to fully autonomous systems that are capable of human-level decision making and planning without external input. The stages in between define notable guidance, navigation and control milestones that increase autonomy level. The scale can be seen in full in Figure 2-1. The MoD definitions of automation and autonomy are most directly comparable to levels 1 and 6, respectively, demonstrating the deficiencies in this approach to autonomy classification.

LEVEL	LEVEL DESCRIPTOR	GUIDANCE	NAVIGATION	CONTROL	ESI	EC	MC
10	Fully Autonomous	Human-level decision-making, accomplishment of most missions without any intervention from ES (100% ESI), cognizant of all within the operation range.	Human-like navigation capabilities for most missions, fast SA that outperforms human SA in extremely complex environments and situations.	Same or better control performance as for a piloted aircraft in the same situation and conditions.	approaching 100% ESI	extreme environment	highest complexity; all missions
9	Swarm Cognizance and Group Decision Making	Distributed strategic group planning, selection of strategic goals, mission execution with no supervisory assistance, negotiating with team members and ES.	Long track awareness of very complex environments and situations, inference and anticipation of other agents intents and strategies, high-level team SA.	Ability to choose the appropriate control architecture based on the understanding of the current situation/context and future consequences.	high level ESI	difficult environment	collaborative, high complexity missions
8	Situational Awareness and Cognizance	Reasoning and higher level strategic decision-making, strategic mission planning, most of supervision by RUAS, choose strategic goals, cognizance.	Conscious knowledge of complex environments and situations, inference of self/others intent, anticipation of near-future events and consequences (high fidelity SA).	Ability to change or switch between different control strategies based on the understanding of the current situation/context and future consequences.	mid level ESI	moderate environment	mid complexity, multi-functional missions
7	RT Collaborative Mission Planning	Collaborative mission planning and execution, evaluation and optimization of multi-vehicle mission performance, allocation of tactical tasks to each agent.	Combination of capabilities in levels 5 and 6 in highly complex, adversarial and uncertain environment, collaborative mid fidelity SA.	same as in previous levels (no-additional control capabilities are required)	low level ESI	simple environment	low level tasks
6	Dynamic Mission Planning	Reasoning, high-level decision making, mission driven decisions, high adaptation to mission changes, tactical task allocation, execution monitoring.	Higher-level of perception to recognize and classify detected objects/events and to infer some of their attributes, mid fidelity SA.	same as in previous levels (no-additional control capabilities are required)	0% ESI	lowest EC	lowest MC
5	RT Cooperative Navigation and Path Planning	Collision avoidance, cooperative path planning and execution to meet common goals, swarm or group optimization.	Relative navigation between RUAS, cooperative perception, data sharing, collision detection, shared low fidelity SA.	Distributed or centralised flight control architectures, coordinated maneuvers.			
4	RT Obstacle/Event Detection and Path Planning	Hazard avoidance, RT path planning and re-planning, event driven decisions, robust response to mission changes.	Perception capabilities for obstacle, risks, target and environment changes detection, RT mapping (optional), low fidelity SA.	Accurate and robust 3D trajectory tracking capability is desired.			
3	Fault/Event Adaptive RUAS	Health diagnosis, limited adaptation, onboard conservative and low-level decisions, execution of pre-programmed tasks.	Most health and status sensing by the RUAS, detection of hardware and software faults.	Robust flight controller, reconfigurable or adaptive control to compensate for most failures, mission and environment changes.			
2	ESI Navigation (e.g., Non-GPS)	Same as in Level 1	All sensing and state estimation by the RUAS (no ES such as GPS), all perception and situation awareness by the human operator.	Same as in Level 1			
1	Automatic Flight Control	Pre-programmed or uploaded flight plans (waypoints, reference trajectories, etc.), all analyzing, planning and decision-making by ES.	Most sensing and state estimation by the RUAS, all perception and situational awareness by the human operator.	Control commands are computed by the flight control system (automatic control of the RUAS 3D pose).			
0	Remote Control	All guidance functions are performed by external systems (mainly human pilot or operator).	Sensing may be performed by the RUAS, all data is processed and analyzed by an external system (mainly human).	Control commands are given by a remote ES (mainly human pilot).			

Figure 2-1: Autonomy Levels for Unmanned Rotorcraft Systems (ALFURS) Framework [35]. Acronyms: ESI (External System Independence), EC (Environmental Complexity), MC (Mission Complexity), ES (External System), SA (Situational Awareness), RT (Real-Time).

### 2.1.3 State of the Art UAS Guidance, Navigation and Control

It is also necessary to examine what technologies are currently in use on state-of-the-art UAS to contextualise these autonomy levels in today's UAS market. In the consumer SUAS market, many aircraft are controlled by propriety GNC systems, and the techniques in use are not advertised. However, there are also a huge number of consumer aircraft that use open source autopilot software, such as PX4.

PX4 is a deeply embedded robotics middleware and programming environment designed specifically for use on Micro Aerial Vehicles (MAVs) [40]. The software allows for multiple operation modes, including remote control and autonomous flight due to its modular nature. PX4 consists of three main tiers: low-level sensor and actuator drivers, a mid-level Object Request Broker, and a high-level Flight Stack. The system is centred on the Object Request Broker that manages all data used in the firmware with a 'publish and subscribe' model, allowing any module in the firmware to access any variable it requires. The flight stack performs all guidance, navigation and control functions using data published by sensor drivers, then publishes control inputs for use by actuator drivers. The default PX4 flight stack contains attitude and position state estimators, attitude and position rate controllers, and attitude and position controllers for autonomous operations.

The modular nature of the PX4 firmware also allows new guidance, navigation and control functions to be easily added to the system by registering new modules in the firmware configuration file. This architecture allows for new controllers or guidance routines to be tested with no need to modify system architecture and is commonly used in research to do so. It is also the firmware used on a number of unmanned aircraft owned by the University of Liverpool, which are detailed in section 3.1. PX4 is also the intended platform for hardware implementation of any technologies developed during this research, though that implementation is not a component of this thesis. As a result, it was just used as an example autopilot throughout the work described in this thesis as an (open) source of inspiration.

#### *2.1.3.1 Visual Guidance Technologies*

According to the ALFURS scale introduced in section 2.1.2, some form of sensor beyond GPS and inertial navigation sensors is necessary for any autonomy level above 2, to reduce reliance on external signals. These sensors are necessary to allow for continued operation in environments where GPS is inadequate or denied and can be either active or passive. Active sensors such as radar or Lidar are valuable for increasing autonomy but can be expensive and limited in range or resolution. Passive sensors such as video cameras are inexpensive and are an intuitive option, since vision forms the basis of natural human and animal state estimation. Visual state

estimation can be divided into 6 distinct types: on ground vision, visual odometry, target relative navigation, terrain relative navigation, concurrent estimate of motion and structure, and bio-inspired optic flow navigation [38].

On ground vision systems use video cameras to track UAS from a third person perspective and estimate aircraft position and attitude in a local or image frame of reference. This removes the need for heavy, power consuming cameras onboard, and reduces the computational requirements on the UAS processor. However, aircraft are greatly restricted in operating area with on ground vision systems. As a result, they are often used as supplemental systems when greater positional accuracy is required, such as for landing or docking manoeuvres.

Visual odometry estimates an aircraft's attitude and relative position by analysing successive images captured by a video camera. This process is performed in three steps: detection of features to be tracked, feature correspondence between images and motion parameter estimation. This has been accomplished using a fusion of stereo visual odometry and inertial navigation data on aircraft such as a CMU autonomous helicopter [41] and the AVATAR helicopter [42]. Visual guidance problems are more complex when using a monocular vision base since depth cannot be recovered from a single image. However, monocular visual odometry has been achieved in 2009 using a mathematical homography based technique to associate and track point features [43]. Researchers at MIT have developed a monocular vision-based method of estimating position and velocity when GPS signals are unavailable by establishing absolute reference points before GPS signal is lost [44]. Another visual odometry system proposed by Kendoul calculates the position of a rotorcraft by computing optic flow, tracking visual features and integrating image displacement to find distance travelled [45].

Target relative navigation uses similar methods to visual odometry but is predicated on the existence of a target. It can be separated into four categories: vision-based landing on a known target, vision-based landing on an unknown target, vision-based static target tracking and vision-based mobile target tracking. Vision-based landing on a known target uses a-priori knowledge of the target landing zone to identify and track the objective. A common example is a helipad consisting of a large 'H' symbol



or other defined shape. These known shapes can be identified within an image and their attitude and size can be used to infer the position and attitude of the aircraft to be landed relative to those shapes. This approach has been performed with monocular vision systems, as the known size of the object to be tracked removes scale factor ambiguity in the image. Reference [46] reports on an example of this type of system. Vision based landing on an unknown target can use similar techniques, but a target point to be tracked must first be identified by through Safe Landing Area Detection (SLAD). SLAD has been accomplished by the NASA/US Army Autonomous Rotorcraft Project with the Palace experiment [47]. Static target tracking shares many features with landing approaches with a modified objective of flying toward, or hovering near, a target. This target can be on the ground but is not necessarily so. An example of static target tracking was given by Proctor and Johnson in 2005, in which an R-MAX helicopter was made to calculate its position and attitude relative to a dark square with known composition, computed from images and inertial data using an extended Kalman filter [48]. Mobile target tracking presents a greater challenge since only relative motion can be estimated and recovering the velocities of target and rotorcraft becomes more difficult without an exterior reference. Mobile target tracking requires a robust image tracking algorithm to track dynamic and unpredictable target behaviour in successive images. A Georgia Tech experiment in 2007 used an active contour method to calculate the location of the target using two distinct strategies [49]: Centre Only Relative State Estimation (CORSE) tracks the target centre position from monocular images using an extended Kalman filter to estimate relative position and velocity. To maximise range accuracy, optimal control was used to generate a sinusoidal trajectory, mimicking stereo vision by moving the viewpoint around the target to give more information. The second approach, Subtended angle relative state estimation (SARSE) built on this approach by also tracking the wing tips of an aircraft as well as the centre to allow for future inference of the targets flight path.

Terrain or landmark relative navigation infers an aircraft's position by comparing terrain measurements from an onboard sensor with an a-priori terrain map. The reference map can be a digital elevation map, satellite images or topographical maps

with landmarks. Terrain Relative Navigation (TRN) systems were originally created for cruise missile guidance before the advent of satellite navigation systems but have also been used for accurate landing of spacecraft on the moon and Mars. Though more uncommon, there are isolated examples of TRN systems implemented on RUAS. In 2010 a system that used probabilistic template matching to compare flight images with a georeferenced map was successfully implemented on a Yamaha R-MAX helicopter [50]. Shortly afterward, a similar system was demonstrated on a small quadcopter, which generated a visual route through an environment database and organised them as ordered target images, then navigated through the area by comparing the target images with onboard sensor output [51].

Concurrent estimation of motion and structure techniques both estimate the pose of a UAS and map the local environment. They are commonly split into two variants: Structure From Motion (SFM) and Simultaneous Localisation and Mapping (SLAM). Both approaches present significant challenges due to issues with robust feature tracking, unknown scale factors in monocular vision, estimate drift and high computational requirements. An example of an SFM method was presented by Kendoul in 2009 which used three nested Kalman filters to compute optic flow, fuse optical and inertial data and recover motion parameters [45]. SLAM methods create a map of an unknown environment while localising the vehicle on that map. The first successful monocular implementation of SLAM, a system named MonoSLAM, was presented in 2010. It used a sparse, persistent map of natural landmarks within a probabilistic framework to localise the camera in real time, citing an active approach to mapping and measurement, the use of a general motion model for smooth camera movement and new solutions for monocular feature initialisation orientation estimation as significant elements of the system. However, this system has only been demonstrated on a desktop computer and not onboard an aircraft due to the high computational requirements [52].

The final category of visual state estimation technologies is bio-inspired optic flow navigation systems. These take cues from natural systems to infer information about motion and environment. These technologies were originally derived from the way flying insects appear to use cues from optical flow for navigation, without explicitly

estimating motion or environmental structure. Optical flow can be defined as the pattern of apparent motion of objects, surfaces or edges in a visual scene caused by relative motion between the observer and scene and was pioneered by James Gibson in the 1940s [53]. Calculation of optic flow centres around estimating the motion between two image frames at every pixel position, with differential methods, phase correlation and block based methods all successfully applied for optical flow calculation. The principles of optical flow can be applied to accomplish many tasks; altitude control or terrain following has been achieved by regulating vertical optical flow to a constant zero value in the Octave autopilot using a small optic flow sensor and simple PID controllers [54]. Another system maintains vertical optical flow constant at a non-zero value, forcing an aircraft to slow down as it descends to maintain the optical flow, allowing for smooth landings based on the behaviour of honey bees. Optical flow has also been used for the purpose of obstacle avoidance, as areas of high optic flow in an image can be interpreted as imminent collisions and so can be avoided without explicitly mapping the environment.

The maritime environment is a challenging even for human pilots, especially during launch and recovery [55]. As a result, use of UAS in the maritime environment is currently limited to only a handful of platforms. The Boeing ScanEagle entered service with the US Navy in 2005 and was also trialled for use by the Royal Navy at the same time, though it did not enter service with the UK forces. The ScanEagle is a small, long-endurance, fixed wing unmanned aircraft used for autonomous battlefield surveillance. At sea it is launched with a catapult and recovered with a 'skyhook' retrieval system [46]. The Northrop Grumman X-47B is a demonstration unmanned combat air vehicle (UCAV) that has been flown by the US Navy since 2011. The aircraft has been successfully integrated into aircraft carrier operations with automated take-off and landings demonstrated [18]. The only RUAS platform currently in active service in a maritime role is the Northrop Grumman MQ-8 Fire Scout, shown in Figure 2-2. The Fire Scout first landed autonomously on an amphibious transport ship in 2006, with no pilot controlling the aircraft and has since flown for over 12,000 hours. For autonomous landings the Fire Scout uses a UAS Common Automatic Recovery System (UCARS) that requires a ground tracking

station and aircraft transponder to find the aircrafts relative position to the deck. This information is then relayed to the aircraft by a secure uplink and used to land the aircraft [56].



*Figure 2-2: MQ-8B Fire Scout RUAS landing onboard the USS McInerney in 2009 [57]*

## 2.1.4 UAS Accident Statistics

### 2.1.4.1 Civil UAS Accidents

The European Union Aviation Safety Agency (EASA), and by extension, the UK Civil Aviation Authority (CAA), established mandatory occurrence reporting for UAS in Regulation No. 376/2014, published in May 2014. The CAA defines a reportable occurrence as 'Any incident which endangers or which, if not corrected, would endanger an aircraft, its occupants or any other person' [58]. Specific to SUAS, reportable occurrences may include:

- Loss of control/datalink – where that loss resulted in an event that was potentially prejudicial to the safety of other airspace users of third parties;
- Navigation failures;
- Pilot station configuration changes/errors:
  - Between pilot stations;
  - Transfer to/from launch control/mission control stations;
  - Display failures.
- Crew Resource Management (CRM) failures/confusion;
- Structural damage/heavy landings;
- Flight programming errors;
- Any incident that injures a third party.

All commercial operators of UAS must comply with these incident reporting rules, but hobbyist pilots are likely not even aware of them since there is no certification or registration procedure for amateur pilots in the UK or the wider EU aviation authority. The FAA's small unmanned aircraft rule (part 107) also establishes mandatory accident reporting requirements, but only in cases where serious injury to a person occurs, or the cost of repair or replacement of any property damage exceeds \$500. Rule 107 only came into force in June 2016 [59]. As a result, many civil SUAS accidents still likely go unreported. Note, all references to legislation are correct at the time of writing, but the rapid regulatory change associated with the UAS may render some statements obsolete in the near future.

Reference [60] reports on a study of 152 civil drone accidents and incidents and incidents between 2006 and 2015, with data collected from 19 different sources in the wake of an alleged collision between a drone and a British Airways Airbus A320 at Heathrow Airport. Statistical analysis was used to determine how distribution of occurrence type, phase of flight and safety issues compare to Commercial Air Transportation (CAT). While the sample size was too small to say whether civil UAS accidents were more likely than CAT accidents, it was shown that technology issues, rather than human factors, were the key contributor in UAS events. This contrasts with CAT accidents where human factors problems are the primary driver behind accidents.

#### *2.1.4.2 Military UAS Accidents*

The most documented UAS accidents are those that occur on US military aircraft since they generally operate the most UA platforms of any single operator. Reference [25] reports on RPAS accidents in the US army between 1986 and 2004. The data was sourced from the Army Risk Management Information System, and a total of 74 accidents were identified. Additional data from the US Naval Safety Centre reported 239 UA mishaps between 1986 and 2002, and a further 15 mishaps were reported by USAF. All these accidents were classified in a two-step process by initial cause (human factors, maintenance, aircraft or unknown), and any human factors causes were further classified into alerts/alarms, display design, procedural error, skill-based error, and other. Extra categories were added in response to airframe specific concerns. Five aircraft were examined in the analysis: the RQ-2 Pioneer, RQ-5 Hunter, RQ-7 Shadow, MQ-1 and MQ-9 Predator and Global Hawk. The number of accidents in each category per airframe is shown in Figure 2-3.

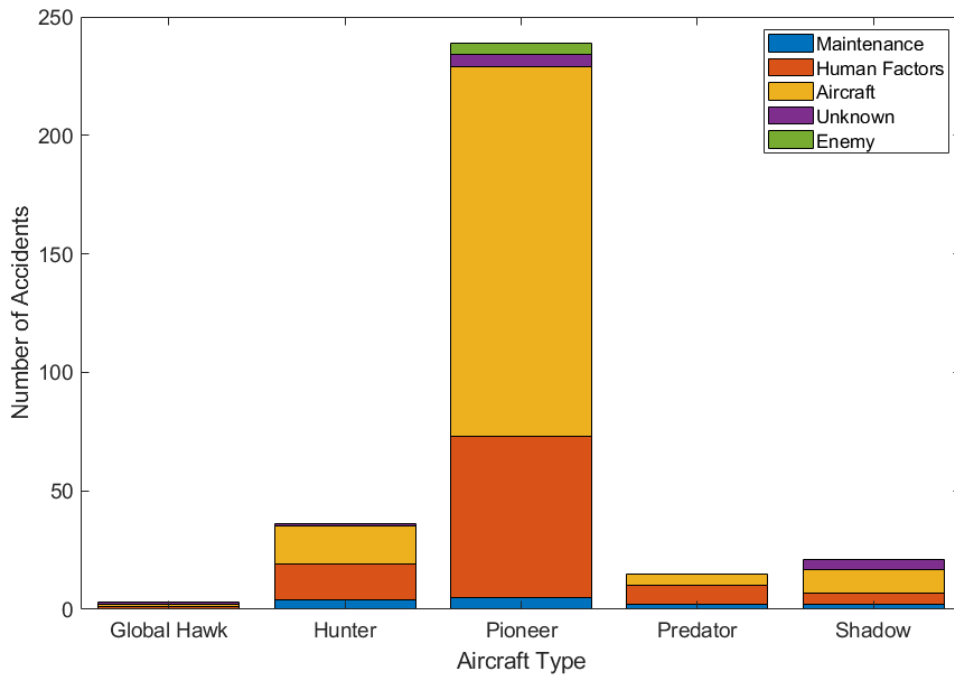


Figure 2-3: US military RPAS accident causes between 1986 and 2004 for the five airframes operational in this period [25]

The RQ-2 Pioneer suffered the highest accident rate by far with 239 total major accidents. The Pioneer was the first RPAS to enter service with any US military force and was in service for the whole 18-year period examined in the report. All of the other aircraft entered service between 1986 and 2004. Lessons were clearly learnt from the Pioneer, and subsequent platforms experienced fewer problems. Aircraft faults make up a large proportion of accident causes, and this is attributed to the more relaxed design processes and standards applied to unmanned aircraft since there is no pilot to protect. Human factors are the second biggest cause of accidents in all aircraft. The identified human factors issues are further categorised in Table 2-3. Some accidents had multiple human factors causes attributed to them.

Table 2-3: Human factors issues causing US military RPAS accidents between 1986 and 2004 [25]

	Hunter	Shadow	Pioneer	Predator	Global Hawk
<b>Pilot-In-Command</b>	1	2			
<b>Alerts and Alarms</b>	2	2		1	
<b>Display Design</b>	1	2		2	
<b>External Pilot Landing Error</b>	7		46		
<b>External Pilot Take-off Error</b>	3		7		
<b>Landing Error</b>		6		1	
<b>Procedural Error</b>	3	2		6	1
<b>Aircrew Coordination</b>			9		
<b>Weather</b>			6		

The Hunter and Pioneer are both operated from Ground Control Stations (GCS) while in the air, but during take-off and landing are operated by an external pilot stood at the edge of the runway. This system caused a large portion of the accidents, especially on Pioneer. These issues are at least partially explained by the cross-control problem faced by external pilots; when the aircraft is moving towards the pilot, control inputs appear reversed when looking at the aircraft. The Pioneer faced even greater challenges since it used arresting cables for runway landings that required very accurate positioning for capture. It also operated from ships where recovery was performed using a net.

Landing was also a problematic flight phase for the Shadow, which used a Tactical Automated Landing System (TALS) to eliminate the external pilot problems faced by its predecessors, though to limited success. The Predator encountered fewer issues during landing, but the newer system was much more vulnerable to procedural errors; one accident occurred when a checklist was not completed and the stability augmentation system was turned off mid-flight, initiating an uncontrolled dive. The most advanced system, the Global Hawk, suffered relatively few accidents, with procedural error the lone culprit in human factors incidents. In this case, the taxi-speed of the aircraft was erroneously inputted into the GCS as 155 knots. The aircraft ran off the runway, collapsing the nose gear and extensively damaging the airframe.



While the falling accident rates suggest that many faults were identified and fixed on the newer airframes, autonomy levels of the vehicles were still low enough that they could not protect from relatively basic operator errors.

The Washington Post article discussed in section 1.2 built on this survey, unveiling that 400 large US military RPAS crashed in major accidents between 2001 and June 2005 [23]. 102 of these aircraft were Predator MQ-1 or -9s, 40% of the total fleet. This suggests that either 87 Predator aircraft crashed in 2004 and mid-2005, or the statistics reported in reference [25] were incomplete. Either way, the mishap rate for this aircraft is high, and this further supports the notion that UAS accident reporting practices may be inadequate. The article listed four fundamental safety gaps that the author viewed as essential improvements for drone manufacturers:

1. Persistent mechanical and electrical defects due to insufficient testing before being brought into service. This included a large number of electrical malfunctions caused by bad weather.
2. Unreliable communication links. Inconsistent wireless links caused navigational information and commands to be missed. Many connections were carried via satellite link that could be disrupted by large bank angles and rapid altitude changes.
3. Pilot errors, particularly during landings.
4. Limited ability to detect and avoid obstacles. For example, the camera used for the pilot's eye view on Predator drones was limited to 30-degrees field of view, which severely limited the remote pilot's awareness of their surroundings.

Data on RUAS accidents is particularly scarce, since there are a few operational examples of rotary-wing UAS in the military. However, the US Navy's MQ-8B Fire Scout RUAS, described in more detail in section 2.1.1, has also faced problems. Fire Scout operations were suspended in 2012, after two major accidents in the space of a week [61]. One aircraft was intentionally ditched in the ocean off the coast of West Africa after failing to lock on to the automated recovery beacon on the USS Simpson. Multiple approaches were attempted amid exhaustive troubleshooting, but

operators were forced to perform a water landing as fuel ran out. The aircraft was later recovered with only minor damage incurred. A second Fire Scout was destroyed after crashing in Afghanistan a week later, prompting the Navy to ground all the remaining aircraft. Details on this second incident are scarce, but the circumstances of the first incident support the hypothesis that unreliable communication links are a major problem for UAS. This link is even more critical for maritime operations due to the need for accurate localisation relative to the moving ship deck, and the apparent lack of backup systems.

Reference [37] contains further information about military drone crashes across the world between 2009 and 2018. The data set included 254 accidents across 21 international operators, though 69% of the aircraft were operated by the US. Again, the MQ-1 and MQ-9 accounted for a majority of these problems with 164 mishaps on these platforms over the 9-year period. This was significantly more than the number of accidents on manned US military attack aircraft over the same period, with only 88 mishaps in this category. The UK armed forces had the second most UAS accidents, with 14 aircraft damaged or destroyed. Analysis of the dataset yielded that 64% of accidents occurred mid-flight, 20% at the point of landing, 8% during take-off, and the remaining 8% occurring at unknown points. This more recent dataset suggests that the rate of drone accidents is declining but is still significantly higher than that of manned aircraft in similar roles. The landing phase is also highlighted as problematic again.

## 2.2 Tau Theory

Tau theory is a key underlying element of all the work presented in this thesis, so it is useful to have some appreciation of its origin, as well as the details of the theory itself. The following sections will outline its genesis and examine the key elements of the theory.

### 2.2.1 Visual Perception

Tau theory has its roots in the ecological approach to psychology and visual perception that was pioneered by J.J. Gibson in the mid-20<sup>th</sup> century. This was largely brought about through his work developing technology for military pilot training in the wake of the Second World War, detailed in reference [53]. Gibson considered how humans perceive the world around them, with the view that perception was inextricable from action; information that humans gained from the world around them was vital in all actions they performed in that world. Gibson theorised that a human was not a detached observer from their environment, simply gathering information in a retinal image to feed into some cognitive process later. Instead, they were immersed in a rich visual environment from which much information could be directly perceived.




Figure 2-4: Optical flow field vectors as would be seen from an aircraft flying over an airfield [62]

Perhaps the most enduring concept of his work is optical flow, defined as “*the pattern of apparent motion of objects, surfaces, and edges in a visual scene caused by the relative motion between an observer and the scene*” [53]. The concept of optical flow is diagrammed in Figure 2-4, where the arrows represent the speed and direction of the visual field motion. It was hypothesised that this perception of motion occurred directly through the light rays entering the observer’s eye, rather than being later constructed from some cognitive process on the retinal images stored in the brain. Gibson also stressed that optical flow was important for his concept of affordance perception, which he defined as the ability to evaluate opportunities for action within an environment.

### 2.2.2 Tau Theory Guidance

Tau theory was first formally proposed by Lee in reference [26] in 1976, where he considered the complex behaviours required by motorists to successfully manoeuvre a vehicle through traffic while avoiding collisions with both moving vehicles and stationary objects. It built on Gibson’s work by again considering information that could be directly perceived from a visual scene. To avoid hitting a stationary object, the driver must be aware of their own position and velocity, and able to extrapolate that motion to register whether or not they are on a collision course. If so, they must then initiate braking in a timely manner, and modulate their braking to an appropriate level throughout the manoeuvre. If following another moving vehicle then they must also monitor their relative motion and maintain a safe distance margin between the two to for allow for the possibility that the other brakes. This process raises questions about both what the driver is visually perceiving, and how whatever visual information gained is applied in his control methodology.




Figure 2-5: A car braking as it approaches another stationary car ahead of it, with the spatial variables labelled

When, or how much, braking should be applied to avoid a collision does not depend simply on either spatial proximity, relative velocity or acceleration, but some combination of these variables, as illustrated in Figure 2-5. From simple kinematics, the instantaneous distance to stop at any point will be:

$$X = \frac{\dot{X}^2}{2\ddot{X}} \quad (2-1)$$

Where  $X$  is the distance to be closed,  $\dot{X}$  is the speed of the car and  $\ddot{X}$  is its acceleration. Hence it follows that to avoid a collision with an object a distance  $X$  away, the driver must modulate their acceleration such that:

$$\frac{X\ddot{X}}{\dot{X}^2} > \frac{1}{2} \quad (2-2)$$

However, distances, speeds and accelerations are hard to judge intuitively, and it is unlikely that all motorists are constantly judging these spatial parameters and performing fast unconscious mental mathematics to avoid a collision. Lee proposed that a simple relationship that could account for all these factors was temporal proximity, otherwise referred to as time-to-collision. He assigned the symbol Tau,  $\tau$ , to represent this variable.

In spatial terms, instantaneous time-to-collision is simple to compute as the ratio between the distance to the obstacle, and the closure rate of that gap. However, it would again seem difficult for a person to constantly estimate these parameters and divide one by the other. Lee instead considered the apparent motion of the image of the approaching obstacle on the motorist's retina; the driver could estimate their time-to-collision from the inverse of the apparent rate of expansion of the object in their eye. These relationships are expressed in equation (2-3).

$$\tau(t) = \frac{x}{\dot{x}} = \frac{1}{\text{rate of retinal image expansion}} \quad (2-3)$$

This parameter gave simple mechanisms for both initiation of braking, and control of the deceleration. Lee presented experimental evidence that motorists initiated

braking when time-to-collision reached a certain threshold value. This temporal margin varied slightly between drivers with an average of 5.0 seconds but did not change with the speed the vehicle was travelling.

### 2.2.3 Constant Rate of Change of Tau Strategy

The rate of change of time-to-collision can be found by differentiating equation (2-3) using the quotient rule, and the result is given in equation (2-4). It bears obvious similarity to that stopping distance defined in equation (2-2) and formed the basis of Lee's hypothesis on how the motorists he studied modulated their decelerative braking inputs; the drivers maintained a constant rate of change of time-to-collision throughout braking manoeuvres less than 0.5 to avoid collision with a looming obstacle.

$$\frac{d(\tau)}{dt} = 1 - \frac{X\ddot{X}}{X^2} \quad (2-4)$$

It is useful to derive equations that give the vehicle position, velocity and acceleration throughout one of these 'constant  $\dot{\tau}$ ' approaches. For a constant rate of change of Tau, 'k':

$$\frac{d}{dt}\left(\frac{X}{\dot{X}}\right) = k \quad (2-5)$$

Both sides can be integrated with respect to time, as in equation (2-6). By convention, at  $t = 0$ ,  $\tau = \tau_0$  so the constant of integration is also  $\tau_0$ .

$$\frac{X dt}{dX} = kt + \tau_0 \quad (2-6)$$

Rearranging to separate spatial and temporal variables gives:

$$\frac{1}{X} dX = \frac{1}{kt + \tau_0} dt \quad (2-7)$$

After evaluating integrals and rearranging:

$$\ln(Cx) = \ln(kt + \tau_0)^{\frac{1}{k}} \quad (2-8)$$

To find the constant of integration at  $t = 0$ ,  $x = x_0$ , so  $C = \frac{1}{x_0}$ . Therefore:

$$X = X_0 \left(1 + \frac{kt}{\tau_0}\right)^{\frac{1}{k}} \quad (2-9)$$

To find expressions for velocity and acceleration, equation (2-9) can be differentiated with respect to time once and twice respectively. The velocity during a constant  $\dot{\tau}$  manoeuvre is:

$$\dot{X} = \dot{X}_0 \left(1 + \frac{kt}{\tau_0}\right)^{\frac{1}{k}-1} \quad (2-10)$$

And the acceleration is:

$$\ddot{X} = \ddot{X}_0 \left(1 + \frac{kt}{\tau_0}\right)^{\frac{1}{k}-2} \quad (2-11)$$

The initial acceleration required for the manoeuvre is unknown, but rearranging equation (2-4) and substituting in that and equation (2-5) yields:

$$\ddot{X} = \frac{\dot{X}_0^2}{X_0} (1 - k) \left(1 + \frac{kt}{\tau_0}\right)^{\frac{1}{k}-2} \quad (2-12)$$

From equation (2-6) it is also possible to derive an equation for the total manoeuvre duration by setting  $\tau = 0$  to find the manoeuvre duration,  $T$ :

$$\tau = 0 = kT + \tau_0 \quad (2-13)$$

$$T = -\frac{\tau_0}{k} \quad (2-14)$$

An example of a constant rate of change of Tau deceleration manoeuvre is shown in Figure 2-6 for a range of coupling constants. Note that for  $k = 1$ ,  $\dot{\tau} = 1$ , so time-to-contact changes at the same rate as time. Therefore, there is no deceleration and collision will occur at initial speed. For  $1 > k > 0.5$ , an infinite deceleration is

required at the moment of contact to close velocity to 0 at the same time. Any real system will be incapable of this infinite deceleration and therefore collide with the obstacle at some speed above 0. For  $k = 0.5$  deceleration is constant, and velocity will close to 0 at the same time as distance. For coupling constants of  $0.5 > k > 0$  initial deceleration will be greater and lead to stopping at the endpoint with a velocity of 0.

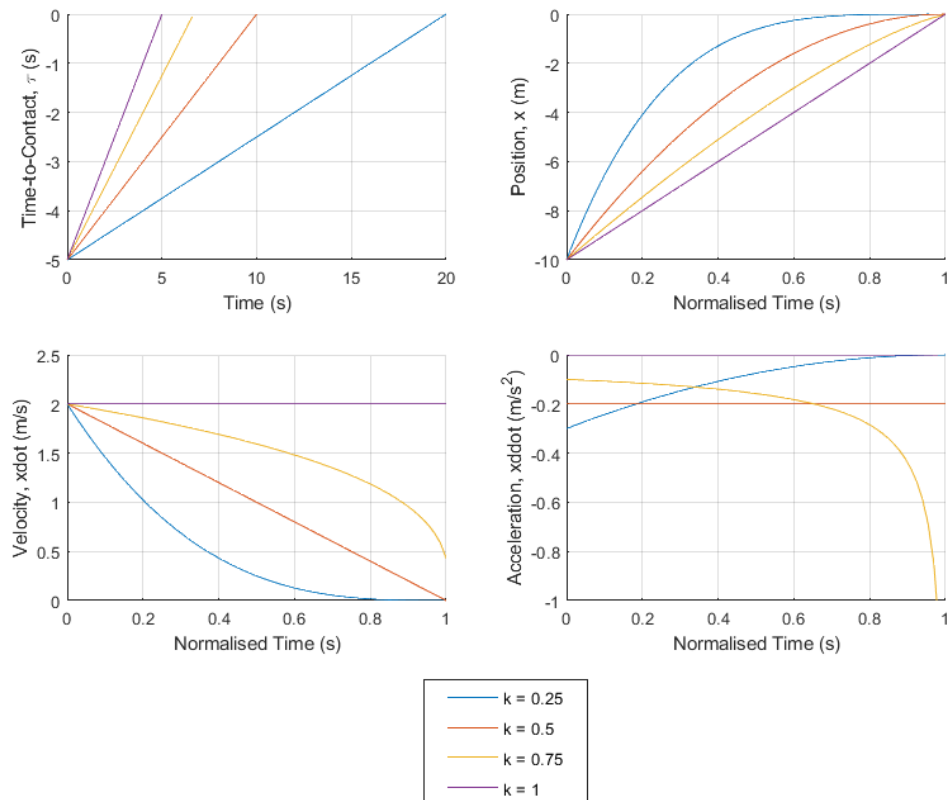


Figure 2-6: A constant rate of change of Tau deceleration manoeuvre: (a) Tau trajectory (b) Motion gap closure (c) Motion gap closure rate (d) Motion gap acceleration profile

Over the next two decades, Lee published a large number of papers demonstrating Tau theory in action, generalising it in the process to more often refer to ‘time-to-contact’, rather than the ‘time-to-collision’ referenced in his first paper. In humans, he demonstrated the use of Tau in many athletic motions including somersaulting [63], regulation of gait in the approach to a long jump [64], control of step length while running over irregular terrain [65] and catching a ball [66].

Tau guided movements were also demonstrated in a range of birds performing various actions. Gannets, a British seabird, search for fish from heights of up to 30m.



They dive toward their prey at high speed and flare their wings just before hitting the water to slow their descent. Analysing the flare manoeuvre with high-speed video cameras through the lens of Tau theory allowed prediction of the flare time when the height of the dive was varied [67]. Similar behaviours were also observed in pigeons landing on a perch; approach velocity appeared to be regulated with the rate of change of Tau, and the extension of their feet to grasp the perch occurred at a certain threshold value of time-to-contact [68]. Hummingbirds also exhibited this behaviour when ‘docking’ with flowers to feed [69]. Several papers also focussed on bats, examining whether they also exhibited Tau guided movements using a different sensory mode: echolocation. In one of these experiments, bats flew through an aperture with their eyes covered to force reliance on acoustic sensing [70]. Strong evidence for them following a constant  $\dot{\tau}$  approach during the manoeuvre was found. This result is interesting as it suggests that time-to-contact is not reliant on vision, and as such can be considered sensor agnostic.

#### 2.2.4 Tau Coupling

Lee built on this work in reference [71] that proposed that the coupling of Taus of multiple action gaps was a key element of Tau guidance. An action gap is a separation between the current state of a system and a goal state to be achieved through some action. Action gaps are usually a spatial dimension, like the distance between a car and a desired stopping point, but the term can cover a range of variables. This work recognised that gaps often need to be closed synchronously and this can be achieved by coupling the time-to-contacts of multiple action gaps together in some ratio with a coupling constant, as in equation (2-15).

$$\tau_X = k \tau_Y \quad (2-15)$$

The constant  $\dot{\tau}$  approach strategy can also be seen as an example of coupling a distance gap to a velocity gap; as the distance to the target reaches zero, so must the velocity. This is illustrated in equation (2-16), the spatial parameters of which can be easily rearranged to the form of equation (2-4). The coupling constant is therefore linked as shown in equation (2-17) for constant  $\dot{\tau}$  manoeuvres.

$$\frac{X}{\dot{X}} = k \frac{\dot{X}}{\ddot{X}} \quad (2-16)$$

$$k = 1 - \dot{\tau} \quad (2-17)$$

Multiple distance gaps can be coupled together to create different approach strategies by varying the coupling constant. This methodology allows for control of both approach velocity and approach direction by varying two different coupling constants. From equation (2-15):

$$\frac{X}{\dot{X}} = k \frac{Y}{\dot{Y}} \quad (2-18)$$

Inverting and integrating with respect to each variable:

$$\ln(Y) = \frac{1}{k} \ln(X) + \ln(C) \quad (2-19)$$

Which can be reduced to:

$$Y = C X^{\frac{1}{k}} \quad (2-20)$$

The constant  $C$  can be found by considering the  $t = 0$  boundary condition, where  $X = X_0$  and  $Y = Y_0$ :

$$C = \frac{Y_0}{X_0^{\frac{1}{k}}} \quad (2-21)$$

The velocity and acceleration equations for the  $y$  direction can also be obtained by differentiating with respect to time once and twice respectively:

$$\dot{Y} = C \frac{1}{k} X^{\frac{1}{k}-1} \dot{X} \quad (2-22)$$

$$\ddot{Y} = C \frac{1}{k} X^{\frac{1}{k}-2} \left[ \left( \frac{1}{k} - 1 \right) \dot{X}^2 + X \ddot{X} \right] \quad (2-23)$$

An example of a multiple-axis, coupled, constant rate of change of Tau manoeuvre is shown in Figure 2-7. A constant  $\dot{\tau}$  deceleration is performed in the  $X$  direction with a coupling constant set at  $k = 0.25$ , while  $\tau_Y$  is coupled to  $\tau_X$  with varying coupling constants between 0.25 and 1. Note that for  $k = 1$  the traces will be identical for

both axes. Plot (b) shows that changing the coupling constant between the Tau of each gap can create different approach angles to the target point. Meanwhile, the other plots illustrate that contact with the target occurs in both axes at a velocity of 0 as in a single axis constant rate of change of Tau manoeuvre; the gap closure is synchronised across the coupled gaps.

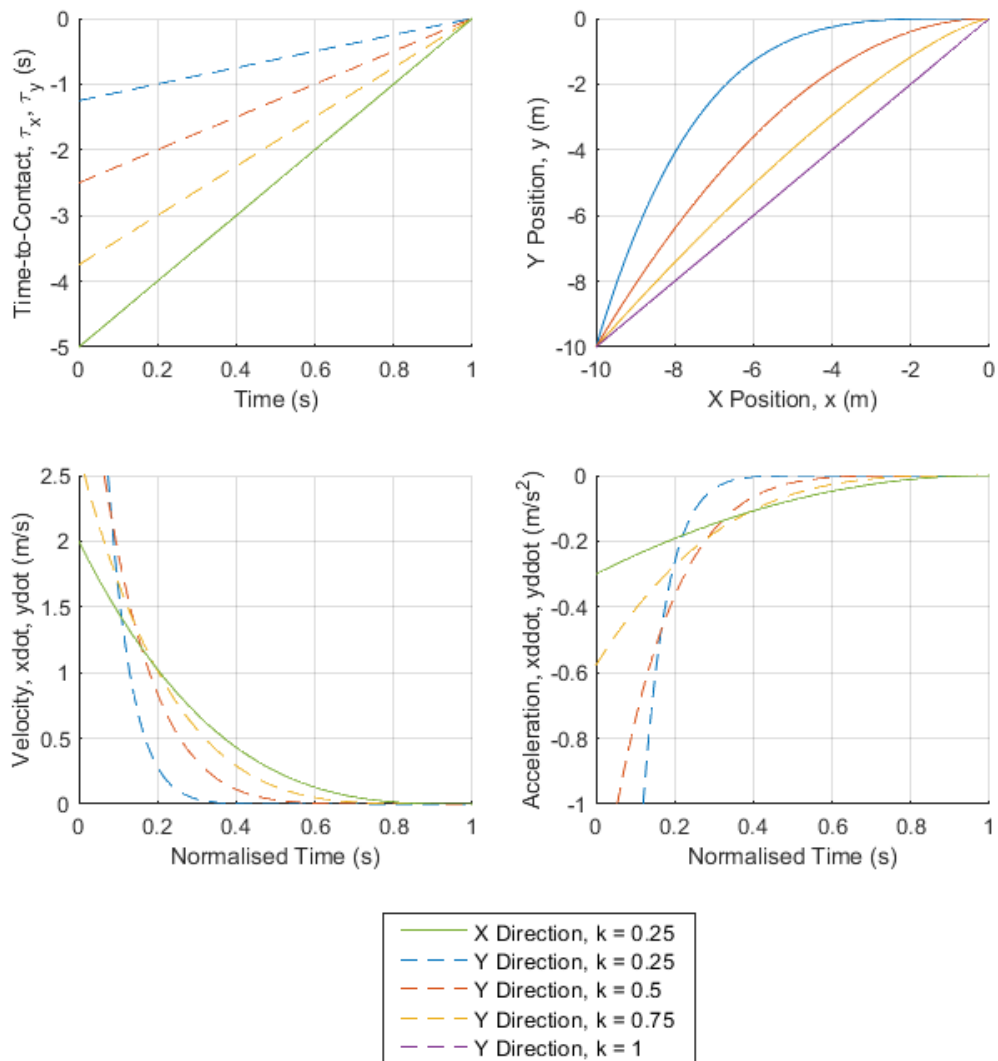


Figure 2-7: A coupled-Tau multi-axis Tau manoeuvre: (a) Tau trajectory (b) Motion gap closure (c) Motion gap closure rate (d) Motion gap acceleration profile

## 2.2.5 Intrinsic Tau Guidance

Thus far, motion gaps have only been coupled to other physical motion gaps, but this is not always an adequate mechanism to explain certain behaviours. Lee considered

a musician in a band, waiting for the right moment in a piece of music to play a chord. They must play the right notes at exactly the right time to fit in with the beat of the music, with the right amount of force to create the appropriate dynamics from their instrument. In this case, there is no motion gap that the player can couple their action to, so Lee hypothesised that the action of playing the chord was being coupled to some intrinsic Tau guide generated by their nervous system. The framework for coupling would be the same; the information about the extrinsic Tau gap would come through the senses, while the intrinsic Tau guide information would be entirely generated within the nervous system. It is assumed that simple control processes evolve over unnecessarily complex ones, so it is reasonable to also assume that any intrinsic Tau guide would have the simplest form possible.

#### 2.2.5.1 Second-Order Intrinsic Tau Guides

Lee conceived a general intrinsic Tau guide as a time-varying value generated by some bodily process. As it was theorised to be some function of the nervous system, he proposed that this value would be generated by a flow of electric charge from one area of the brain to another, from a full 'reservoir' to an empty one. The rate would be such that the second-order time derivative of the flow was constant. At  $t = 0$ , the reservoir will be full of electrical charge, here termed  $q_G$ , and the flow rate will be 0. After some time,  $T_G$ , the reservoir will be empty. From simple kinematics equation (2-24) can be constructed to give the amount of charge remaining in the reservoir at some time before emptying.

$$q_G = \frac{1}{2} \ddot{q}_G (T_G^2 - t^2) \quad (2-24)$$

Again, the value of the Tau guide will be the ratio of distance to close the gap, and gap closure rate. The rate of emptying at a time  $t$  will be  $-\ddot{q}_G t$ , therefore the value of the Tau guide will be:

$$\tau_G = \frac{\frac{1}{2} \ddot{q}_G (T_G^2 - t^2)}{-\ddot{q}_G t} = \frac{1}{2} \frac{T_G^2 - t^2}{t} \quad (2-25)$$

This expression for the Tau guide can then be coupled to an extrinsic motion gap:

$$\tau_M = k \tau_G = \frac{k t^2 - T_G^2}{2} \quad (2-26)$$

Instead of considering the flow of charge in the nervous system, the same result could also be achieved by analysing an object falling in earth's constant gravitational field. It is again useful to derive equations to define the position, velocity and acceleration trajectories of a coupled intrinsic Tau guide. From equation (2-26), a coupling intrinsic Tau guide can be expressed as:

$$\frac{X_M}{\dot{X}_M} = k \frac{q_G}{\dot{q}_G} \quad (2-27)$$

Rearranging and integrating with respect to each motion gap yields:

$$\ln(X_M) = \ln(q_G)^{\frac{1}{k}} + \ln(C) \quad (2-28)$$

Raising to the power  $e$  and substituting in equation (2-24):

$$X_M = C \left[ \frac{1}{2} \ddot{q}_G (T_G^2 - t^2) \right]^{\frac{1}{k}} = D_\tau (T_G^2 - t^2)^{\frac{1}{k}} \quad (2-29)$$

Differentiating with respect to time once and twice produces equations for velocity and acceleration respectively:

$$\dot{X}_M = -2D_\tau \frac{1}{k} t (T_G^2 - t^2)^{\frac{1}{k}-1} \quad (2-30)$$

$$\ddot{X}_M = 2D_\tau \frac{1}{k} \left[ \left( \frac{2}{k} - 1 \right) t^2 - T^2 \right] (T_G^2 - t^2)^{\frac{1}{k}-2} \quad (2-31)$$

The value of  $D_\tau$  can be found by assessing the boundary condition at  $t = 0$  with equation (2-29):

$$X_0 = D_\tau (T^2 - 0^2)^{\frac{1}{k}} \quad (2-32)$$

Therefore:

$$D_\tau = \frac{X_0}{T^{\frac{1}{k}}} \quad (2-33)$$

An example of a second-order intrinsic Tau guide coupled to a motion gap is shown in Figure 2-8. The main difference between a second-order intrinsic Tau guide and a

first-order constant  $\dot{\tau}$  guides lies in the initial conditions: for a constant  $\dot{\tau}$  guide, the entity must have some initial velocity. An intrinsic Tau guide allows for the entity to accelerate from rest and then decelerate to a stop as evidenced by the velocity trace. Varying the coupling constant will again affect how acceleration is distributed through the manoeuvre. For  $k = 1$ , the entity will accelerate at a constant rate throughout the manoeuvre and collide with the target at speed. For  $1 > k > 0.5$ , it will accelerate toward the target for a greater proportion of the manoeuvre and then decelerate hard at the end to attempt to contact with a velocity of 0. However, this will again require an infinite deceleration that any real system will be unable to provide, so will make contact with the target at some speed above 0. For  $0.5 > k > 0$  the entity will accelerate quickly in the first half of the manoeuvre and begin its deceleration early to shed most of its velocity before the terminal phase. It will, therefore, make contact with the target with 0 velocity.

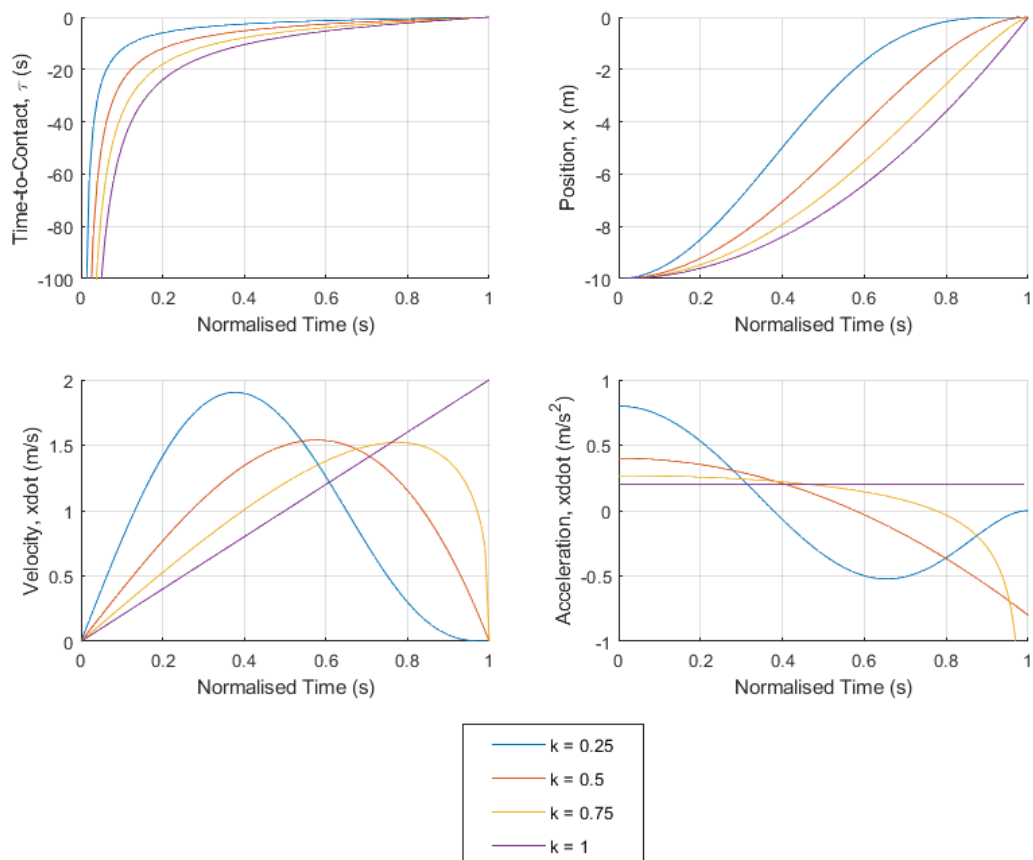


Figure 2-8: A coupled second-order intrinsic Tau guide manoeuvre: (a) Tau trajectory (b) Motion gap closure (c) Motion gap closure rate (d) Motion gap acceleration profile

### 2.2.5.2 Zero-Second-Order Intrinsic Tau Guides

The constant rate of change of Tau strategy can also be explained through intrinsic Tau guides if a zero-second-order (first-order) version of equation (2-24) is constructed. Since  $\ddot{q}_G = 0$ , the amount of charge left must instead be  $\dot{q}_G t$ :

$$\tau_G = \frac{\dot{q}_G(T_G - t)}{-\dot{q}_G} = t - T_g \quad (2-34)$$

This guide can then be coupled to an extrinsic motion gap:

$$\tau_M = k(t - T_G) \quad (2-35)$$

Substituting in equation (2-14) for manoeuvre duration, equation (2-35) can be easily rearranged to the form of equation (2-6), proving that a constant  $\dot{\tau}$  approach can also be attributed to an intrinsic Tau guide.

### 2.2.5.3 Third-Order Intrinsic Tau Guides

Though third-order intrinsic Tau guides have not been explicitly folded into the Tau theory framework, motions that follow their expected movement pattern have been observed in various arm movements. In 1984 Hogan theorised that primates would attempt to minimise jerk (the 3<sup>rd</sup> order derivative of displacement) during certain voluntary movements to avoid large accelerative transients [72]. He presented a solution to the optimal control problem of forearm motion using a jerk cost function, but this approach required symmetric velocity profiles that were not always present when performing faster movements. A later study by Nagasaki examined arm movements over a wide range of speeds and used Hogan's minimum jerk model to analyse the trajectories observed, including an 'asymmetry index' to explain asymmetric velocity profiles that were observed during more aggressive motions [73]. This bears immediate similarity to the coupling constants of Tau theory. There is no literature definitively comparing any natural movement to a defined third-order intrinsic Tau guide but the equations that such a guide would necessitate are defined in reference [74]. These third-order guides were derived using the same process as second-order guides, modelling the flow of charge from one area of the brain to

another. However, a third-order guide assumes that both  $\dot{x} = 0$  and  $\ddot{x} = 0$  at  $t = 0$ , so from simple kinematics the amount of charge left to flow is:

$$q_G = \frac{1}{6} \ddot{q}_G (T_G^3 - t^3) \quad (2-36)$$

And the flow rate at a given time:

$$\dot{q}_G = -\frac{1}{2} \ddot{q}_G t^2 \quad (2-37)$$

Therefore, the reservoir of charge's time-to-empty is:

$$\tau_G = \frac{q_G}{\dot{q}_G} = \frac{1}{3} \frac{t^3 - T_G^3}{t^2} \quad (2-38)$$

The guide can then be coupled to an external action gap as usual:

$$\tau_M = k\tau_G = \frac{k}{3} \frac{t^3 - T_G^3}{t^2} \quad (2-39)$$

The equations for position, velocity and acceleration can be found in the same manner as for second-order tau guides, using the variant equation for electrical charge instead (equation (2-36)). This results in:

$$X_M = D_\tau (T_G^3 - t^3)^{\frac{1}{k}} \quad (2-40)$$

$$\dot{X}_M = -D_\tau \frac{3}{k} t^2 (T_G^3 - t^3)^{\frac{1}{k}-1} \quad (2-41)$$

$$\ddot{X}_M = D_\tau \frac{3}{k} t \left[ \left( \frac{3}{k} - 1 \right) t^3 - 2T_G^3 \right] (T_G^3 - t^3)^{\frac{1}{k}-2} \quad (2-42)$$

Where:

$$D_\tau = \frac{X_0}{T_G^{\frac{k}{3}}} \quad (2-43)$$

Third-order intrinsic guides can be useful because they allow for movements with zero acceleration at the start and end of the manoeuvre. First and second-order intrinsic Tau guides necessarily require an instantaneous non-zero acceleration at the start of the manoeuvre, which the controlled system may not be capable of



developing, leading to tracking errors. An example of the temporal and spatial trajectories generated by a third-order intrinsic Tau guide are shown in Figure 2-9.

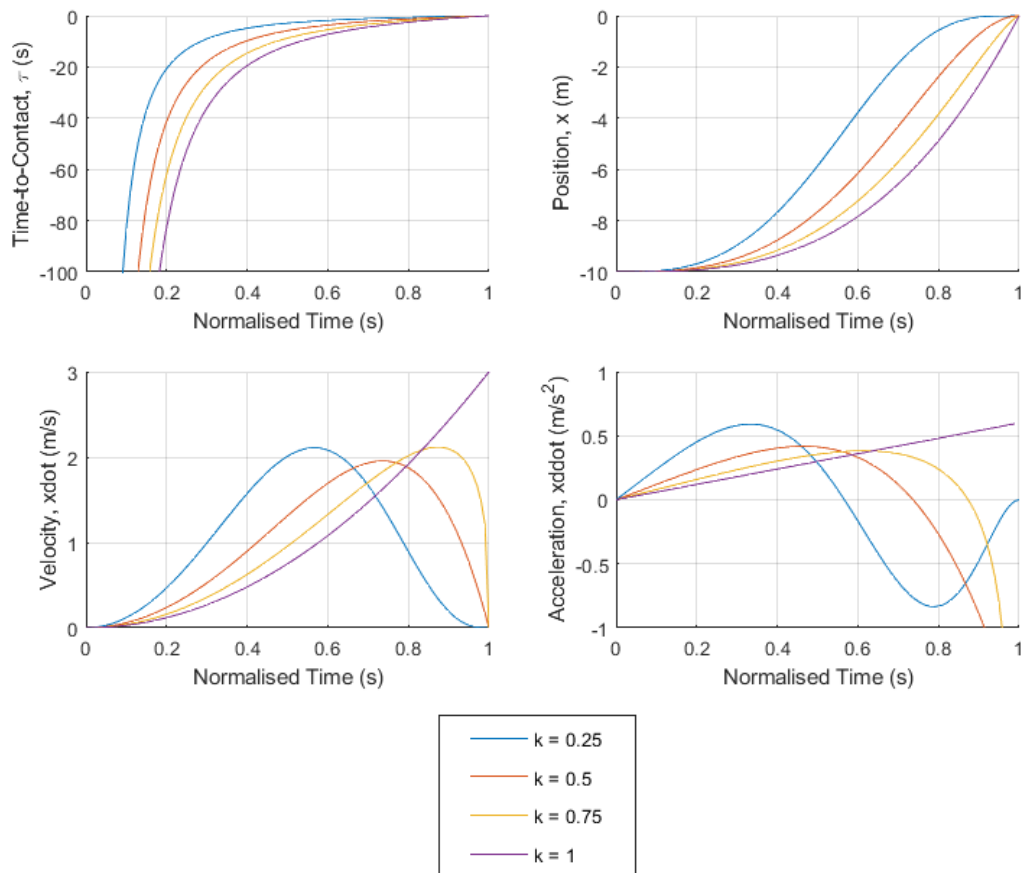


Figure 2-9: A coupled third-order intrinsic Tau guide manoeuvre: (a) Tau trajectory (b) Motion gap closure (c) Motion gap closure rate (d) Motion gap acceleration profile

#### 2.2.5.4 Summary of Intrinsic Tau Guides

Intrinsic Tau guides are a powerful tool that can be used to effect a range of different guidance strategies by changing the order of the Tau guide equation and a few intuitive parameters. If generating trajectories to guide a system, the most appropriate Tau guide to follow seems to depend on the initial conditions of the system that will be coupled to the Tau guide. However, all offer the same key benefits: control of manoeuvre duration, and synchronised deceleration to a stop at a desired point in space. These two factors allow for generation of guidance trajectories in four dimensions (three spatial and one temporal). The different order Tau guides are summarised in Table 2-4, with the velocity and acceleration conditions that are associated with the use of each one.

Table 2-4: Summary of different order Tau guides and the initial conditions that necessitate their use

Initial Conditions		Tau Guide Order	Tau Guide Equation
Velocity	Acceleration		
$\dot{X} \neq 0$	$\ddot{X} \neq 0$	1 <sup>st</sup>	$\tau_G = t - T_G$
$\dot{X} = 0$	$\ddot{X} \neq 0$	2 <sup>nd</sup>	$\tau_G = \frac{1}{2} \frac{t^2 - T_G^2}{t}$
$\dot{X} = 0$	$\ddot{X} = 0$	3 <sup>rd</sup>	$\tau_G = \frac{1}{3} \frac{t^3 - T_G^3}{t^2}$

### 2.2.6 Alternatives to Tau Theory

Much has been written in support of Tau theory since its inception, but it is not universally accepted as a mechanism of ecological guidance. Reference [75] presents an alternative perspective on Tau theory with a critique of the body evidence that has been used in favour of Lee's theory. This study raised concerns about a number of experiments that were used as evidence for Tau theory, which used visual scenes presented through flat-screen monitors to investigate motion control. The author argued that this method of presenting a visual scene removed depth perception cues that are of importance in reality. The study also presented an alternative theory for explaining the interceptive behaviours attributed to Tau theory, instead proposing that a relative distance ratio could provide an equally viable explanation for the acrobatic actions investigated.

As mentioned in section 1.6.1, the focus of this thesis is not to provide any evidence in support, or opposition, of Tau theory. Significant evidence has been presented above that certain actions seem to follow the trajectories that result from the Tau theory hypothesis, and that these approaches are useful in some situations. This information is solely used as a basis for assessing the utility of these guidance strategies on UAS.

## 2.3 Tau in Aviation

Since the use of Tau has been inferred in both the activities of humans and birds, it is reasonable to assume that humans may use Tau as a prospective control strategy while operating aircraft. Since the origin of Tau theory lies in the analysis of motorists it is evident that humans can extend their natural control mechanisms to the operation of machines, but it was not clear how the extra degrees of freedom in an aircraft might play into this paradigm.

This question was first addressed by Padfield in reference [29], co-authored by Lee. They examined common manoeuvres performed by helicopter pilots, such as approaching a hover point, climbing to avoid rising terrain and turning along a line feature. Several flight test experiments at the University of Liverpool and the Defence Evaluation and Research Agency (DERA) were carried out to assess whether pilots were following Tau guides during an acceleration-deceleration manoeuvre. The resulting trajectories correlated strongly with constant  $\dot{t}$  acceleration and deceleration guides in each phase, even when task aggression was varied, suggesting that Tau was a key element in safe rotorcraft flight.

Padfield also examined a hill climb manoeuvre in a 2007 paper, in which he analysed the behaviour of helicopter pilots approaching rising terrain using Tau theory [76]. Flight simulation experiments showed that pilots used their time-to-contact with the slope as a measure of when to initiate a pull-up manoeuvre, with look-ahead times of 6-8 seconds regardless of velocity. The experiments were performed in both a Good Visual Environment (GVE) and a Degraded Visual Environment (DVE) to assess whether the quality of visual cues available would affect Tau behaviour. Results showed that in DVE pilots maintained a similar look-ahead margin but were more likely to adopt lower coupling constants in the manoeuvres they performed with low visibility. The pilots followed second-order intrinsic Tau guides when performing the pull-up manoeuvre in both GVE and DVE. This perspective was reinforced by further flight simulation experiments performed by Lockett in reference [77]. In these experiments, an approach to hover manoeuvre was performed in both GVE and DVE, instead of Padfield's hill climb pull-up. The results showed that the pilot effected

constant  $\dot{\tau}$  deceleration approaches in both GVE and DVE with similar Tau margins as before.

The presence of Tau in the piloting of fixed-wing aircraft was demonstrated by Jump in reference [78] in both the take-off and landing phases of a civil airliner flight. During the rotation phase of take-off, it was found that the pilot increased pitch angle to a desired value for a climb by following an intrinsic Tau guide. While interesting, this was not the focus of the research. Instead, the more safety-critical flight phase of flare during landing was examined. More simulation results gave evidence that the flare manoeuvre was conducted via a constant  $\dot{\tau}$  deceleration strategy to arrest the descent and land safely [28]. Jump further suggested that insufficient visual information to perform the flare manoeuvre was a primary cause of some aviation accidents. In response he developed guidelines for the design of novel display technology to provide Tau cues to the pilot when visual information was insufficient in reference [79].

## 2.4 Tau Perception

Although theories have been presented, the exact ecological mechanism for sensing time-to-contact remains unclear. A range of approaches have been proposed for electronically estimating time-to-contact. These methods can broadly be divided into two categories: spatial and optical.

### 2.4.1 Spatial Tau Perception Methods

A common approach is to calculate time-to-contact from spatial information provided by GPS-INS systems, an example of which is reported in reference [30]. This approach operates by using estimates of the aircraft position and velocity to calculate Tau. For example, one can find a simple estimate of a descending aircraft's time-to-contact with a ground plane by leveraging the altitude and descent rate:

$$\tau = \frac{Z}{\dot{Z}} \quad (2-44)$$

However, these approaches neutralise some of the advantages of Tau-based guidance, chiefly their ability to react to changes in the surrounding environment. It

is impossible to always provide accurate values for spatially computed time-to-contact estimates without comprehensive prior knowledge of the environment and its dynamics, which is usually not feasible, or even possible, to collect, store or use. Depending on the source of altitude estimates it is likely that the given value is not the true height above terrain, as this would either require some variety of active ranging altimeter, or a database of terrain altitudes in the local area. The resolution of either is finite, so vulnerable to error.

Another common method of generating spatially computed time-to-contact estimates is to use machine vision to localise the aircraft with respect to some known visual feature or template. This approach was used in reference [32], where April Tags were placed around a desired landing area for a quadcopter. The April tags were used to localise the aircraft with respect to the target, and the distance and velocity of the aircraft used to calculate Tau. This method allows precise control and reactivity to dynamic environments but requires templates or known features to be placed within the area before operations.

#### *2.4.1.1 Active Tau Perception Methods*

There is strong natural precedent for using active sensors in Tau perception, i.e. sensors that emit a signal in order to estimate the range to an obstacle or target and hence calculate Tau. Section 2.2 discussed experiments involving bats, who use active sonar signals to sense obstacles, and reference [70] presented evidence for the use of Tau guidance using Tau from echolocation. This suggests that active sensors could be a useful tool for Tau perception.

Reference [80] details an early implementation of an active Tau perception system implemented on two different wheeled robots. One robot carried a light bulb of known luminance, and the second was equipped with a light sensor that was used to produce an estimate of the distance between the two robots based on the amount of light received from the first robot. The measured distance and its rate of change were used to produce a Tau estimate.

Reference [81] provides analysis of Radar, Lidar and ultrasound systems as prospective tools for next generation driver assistance functions, all of which could

function as prospective Tau sensors. As already established in the discussion of Tau theory, motorists must maintain high situational awareness while driving, so accurate, reliable sensors for gaining knowledge of the environment and other motorists are essential if elements of driving are to be made autonomous in future. The study assessed Radar, Lidar and Ultrasound sensors effectiveness within defined range bands and angular resolutions, taking into additional capabilities. The findings are summarised in Table 2-5.

*Table 2-5: Typical strengths and weakness of Radar, Lidar and Ultrasound automotive sensors. Key: ++: ideally suited, +: good performance, 0: possible but drawbacks expected, -: only possible with large additional effort, --: impossible, n.a.: not applicable [81]*

	Short Range Radar	Long Range Radar	Lidar	Ultrasound
Range Measurement (< 2 m)	0	0	0	++
Range Measurement (2 – 30 m)	+	++	++	-
Range Measurement (30 – 150 m)	n.a.	++	+	--
Angular Resolution	0	0	++	-
Direct Velocity Information	++	++	--	0
Operation in Rain	++	+	0	0
Operation in Fog or Snow	++	++	-	+
Operation if Dirt on Sensor	++	++	++	++

In summary:

1. Radar sensors are useful for medium and long range but have poorer performance at short ranges. They do have the useful capability of being able to measure velocity information directly through Doppler Radar.
2. Lidar is most effective at medium range, but performance is heavily degraded by environmental factors. The short wavelength of laser radiation also prevents use of the Doppler effect to measure velocity information.
3. Ultrasound sensors perform well at short range but very poorly at long ranges. The report does remark that ultrasonic sensors are currently the

most used of the three sensor types in the automotive industry; they are used for parking sensors due to their relatively low cost.

The report concludes that Radar and Lidar are both useful prospective sensors for the Automotive industry, but the weight and cost of high-quality units is currently prohibitive in their implementation.

#### 2.4.2 Optical Tau Perception

Optical methods estimate time-to-contact directly from sequences of images without recovering spatial parameters. Passive imaging provides a more obvious analogue to the mechanism used by humans and animals, and sensors are low cost and readily available. They also have a wide range of applications so can be used by multiple systems.

Computer vision with a monocular camera is typically unable to recover motion parameters due to scale ambiguity; with only one viewpoint it is impossible to tell the difference between a small object close to the camera, and a larger one further away [82]. However, it is not necessary to recover spatial parameters to compute temporal ones. Monocular computer vision methods examined here can be subdivided into three further common categories for Tau perception: dimension tracking, optical flow divergence, and direct gradient-based.

##### *2.4.2.1 Dimension Tracking Methods*

The simplest method for Tau perception uses size and rate information obtained from a sequence of images to find time-to-contact [82]. While it isn't possible to estimate the true size of an object from only an image of it, the ratio of its apparent size and the rate of change of that size can be used to estimate temporal parameters. Dimension-tracking methods leverage this concept simply and intuitively using a simple pinhole camera model as diagrammed in Figure 2-10.




Figure 2-10: Diagram of a camera descending toward an object on the ground plane, with a point on the object represented by a point,  $P$ , with the relevant distances labelled in the  $X$  and  $Z$  axes.

In a pinhole camera model, all light rays entering the camera are assumed to pass through the same point. Therefore, the following equations can be derived using the geometric similarity between the two triangles formed by the light ray:

$$\frac{x}{f} = \frac{X}{Z} \quad (2-45)$$

In this way, the position of a point in the visual scene can be related to the focal length of the camera and its position in the image [82]. If the camera is allowed to move in the  $Z$ -axis, an expression for the movement of the point on the image plane can be derived by differentiating the perspective projection equations above with respect to time:

$$\frac{dx}{dt} = f \frac{\dot{X}Z - X\dot{Z}}{Z^2} \quad (2-46)$$

Since the size of the object does not change, the rate of change of  $X$  will be zero, and therefore some terms can be eliminated from the equations. This also assumes that the point of interest is not moving on the ground plane. By setting  $\dot{X} = 0$ :



$$\frac{dx}{dt} = -f \frac{X\dot{Z}}{Z^2} \quad (2-47)$$

Rearranging equations (2-45) for the image coordinates, then dividing by equation (2-47) produces an expression for time-to-contact:

$$\frac{x}{\dot{x}} = \frac{f \frac{X}{Z}}{-f \frac{X\dot{Z}}{Z^2}} = -\frac{Z}{\dot{Z}} = \tau \quad (2-48)$$

Using this result, a system can be designed to track some dimension in an image and use it to produce an estimate of time-to-contact. This dimension might be a characteristic length of an object in the image, or a distance between some defined features. This method has been widely implemented, but reference [83] details a patent for a time-to-contact sensor that uses this approach.

#### 2.4.2.2 Optical Flow Divergence Methods

A second option for Tau estimation using monocular images stems from analysis of the optical flow field that develops with camera motion. In a computer vision context, optical flow can be considered the apparent velocity of pixels on the image plane.

##### 2.4.2.2.1 Optical Flow Field Estimation

Any calculation of optical flow is based on an underlying assumption of constant image brightness as originally proposed by Horn and Schunck in reference [84]. Any image can be regarded as an intensity or brightness pattern, and expressed as a function in these terms:

$$I(x, y, t) \quad (2-49)$$

This image brightness or intensity,  $I$ , is a function of position on the image plane,  $[x, y]$ , and time,  $t$ . The constant brightness assumption states that the intensity of a particular point in the brightness pattern does not change significantly with time; in other words, the information contained within the images of a video sequence does not change between video frames:

$$\frac{d}{dt} I(x, y, t) = 0 \quad (2-50)$$

The customary constant brightness equation is derived by applying the chain rule for differentiation:

$$\frac{dI}{dx} \frac{dx}{dt} + \frac{dI}{dy} \frac{dy}{dt} + \frac{dI}{dt} = I_x u + I_y v + I_t = 0 \quad (2-51)$$

Where  $u$  and  $v$  are the optical flow components in the  $x$  and  $y$  directions, and the  $I_x$ ,  $I_y$  and  $I_t$  terms are the image gradients with respect to  $x$ ,  $y$  and time directions. A single equation is insufficient for calculating the two optical flow components, so a further constraint equation must be introduced for a solution. This second equation is typically where optical flow methods differ most.

#### 2.4.2.2.2 Horn and Schunck Optical Flow

Horn and Schunck's second constraint equation assumes that optical flow varies smoothly across the image. The constant brightness assumption states that the brightness of a point remains constant between frames but allows for its motion within the image. If all points in an image moved independently, it would be nearly impossible to recover the optical flow velocities. However, in the majority of cases, a video will depict rigid objects undergoing motion so that neighbouring points will move with similar velocities; the optic flow field will have some degree of smoothness. This inference can be used to derive a second constraint equation to use in the optical flow calculations by using the Laplacian of the vector flow fields as a measure of smoothness. The full derivation of the method can be found in Appendix B.1.1, but reduces to a set of equations that must be solved iteratively for the optical flow field.

#### 2.4.2.2.3 Lucas-Kanade Optical Flow

Lucas-Kanade optical flow is a widely used alternative to Horn and Schunck's method and is detailed in reference [85]. It is based on an assumption that optical flow is constant in a local region surrounding a feature or pixel of interest. The constant brightness equation is solved in the local neighbourhood of the pixel through a least squares criterion, so the extra constraint equations used in this case are actually more versions of the constant brightness equation, specific to certain pixels. The

derivation of the equations for the Lucas-Kanade method can be found in Appendix B.1.2.

To implement this method, a feature detection method is first applied to the image to find points of interest, and then the image gradients are determined at the pixel locations in a window of predetermined size around the feature. The image gradients can be estimated using the same methods as described for the Horn and Schunck estimator in Appendix B.1.1. The Lucas-Kanade estimator can estimate flow efficiently but can only operate where features are detected. It also relies on the extra assumption that flow is constant in the region around the feature. This is in direct opposition to the requirements of the Tau estimation method for optical flow fields that uses the divergence of the flow field. L-K flow estimators are widely implemented in image processing packages.

#### 2.4.2.2.4 Brightness Variance in Optical Flow Fields

All of these optical flow estimation methods rely on the base constant brightness assumption. As with all assumptions, it will break down in certain situations. In this case, the assumption fails with changes in image illumination; brightness is not constant. Changes in illumination can be misinterpreted as motion field change and influence optical flow estimates. These changes are difficult to exclude in scenarios with looming objects because shadows cast by the observer can encroach on the visual scene. A solution to this issue is to acknowledge that there will be a brightness change in the image and account for this in the equation. An alternative equation proposed in reference [86] models this brightness change through a multiplicative field,  $M$  and an offset,  $C$ . With this methodology, the brightness of each pixel can vary up and down and be moved around the image. This concept can be used to reformulate the constant brightness equation as:

$$\frac{d}{dt}I(x, y, t) = M(x, y, t)I(x, y, t) + C(x, y, t) \quad (2-52)$$

Which can be expanded by differentiating the image brightness with the chain rule, and using  $m_t$  and  $c_t$  as the multiplicative and offset field as  $t$  tends toward 0:

$$I_x u + I_y v + I_t - I m_t - c_t = 0 \quad (2-53)$$

The same assumption that was proposed by Lucas-Kanade is then used to produce a solution to this variable brightness equation; optical flow is assumed constant in a region about a point or feature in the image. Since Lucas-Kanade solutions are overdetermined, there are enough equations to cover the increased number of unknowns in the variable brightness equation. The method was shown to produce coherent optical flow fields for situations with scene brightness change, where other methods produced no meaningful results.

#### 2.4.2.2.5 Tau from Optical Flow Divergence

Once the flow field has been estimated, time-to-contact can be estimated from that field. It is not possible to recover spatial parameters of the visual scene from optical flow methods, but temporal parameters can be obtained by analysing the vector flow field. Specifically, the divergence of the optical flow field can be manipulated to yield time-to-contact estimates. Again, there are many examples of this method being implemented, but reference [87] provides details of a patent for a Tau sensor based on this method. It used an actuated camera to centre the focus of expansion within the image and then compute the optical flow field and find the divergence.

#### 2.4.2.3 Direct Gradient Methods

The third class of methods examined are known as direct gradient methods and can be seen as an amalgam of dimension tracking and optical flow-based methods. This approach was first proposed by Horn in 2007 [88]. A gradient method combines the perspective projection equations and constant brightness equation to simplify the Tau estimation procedure using physical insights to the situation. The initial proposal of this method by Horn presents three different algorithms for estimating Tau in scenarios of varying complexity, including simple one-dimensional motion, translational motion in three dimensions, and time-to-contact with sloped planar surfaces. A further extension proposed a hierarchical framework for three-dimensional motion relative to sloped planar surfaces in reference [88]. A third extension was proposed in reference [89] that incorporates corrections for rotational motion in three dimensions into the method. The derivations for these methods are

of interest for this thesis since they are heavily analysed, so the derivations for the five variants of the direct gradient method are included in Appendix B.2.

#### 2.4.2.4 Comparison of Spatial and Optical Tau Perception Methods

Several methods for computing time-to-contact from sequences of images have been presented, but not analysed in great detail here. A full comparison of the three optical methods for Tau perception with their benefits and drawbacks is presented in chapter 4. However, it is useful at this point to compare the utility of an optical Tau sensor against the sensors that have been identified as viable for Tau perception in section 2.4.1. Table 2-6 is an updated version Table 2-5 with an optical Tau sensor included for comparison. It is important to note that while an optical Tau sensor is not actually recovering range or velocity, Tau is considered as a combination of these variables, so the ‘Range Measurement’ fields can be considered as ‘accurate Tau measurements at this range’ for the Tau sensor.

*Table 2-6: Typical strengths and weakness of Radar, Lidar and Ultrasound automotive sensors, with Optical Tau sensor also included for comparison. Key: ++: ideally suited, +: good performance, 0: possible but drawbacks expected, -: only possible with large additional effort, --: impossible, n.a.: not applicable [81]*

	Short Range Radar	Long Range Radar	Lidar	Ultrasound	Optical Tau
Range Measurement (< 2 m)	0	0	0	++	-
Range Measurement (2 – 30 m)	+	++	++	-	++
Range Measurement (30 – 150 m)	n.a.	++	+	--	+
Angular Resolution	0	0	++	-	++
Direct Velocity Information	++	++	--	0	++
Operation in Rain	++	+	0	0	0
Operation in Fog or Snow	++	++	-	+	0
Operation if Dirt on Sensor	++	++	++	++	0

It is clear that no single one of these sensors is suitable for all situations and the different options all have different benefits and drawbacks. Optical Tau sensors can

be very effective at mid- to long-range, and have excellent field of view in comparison to other technologies, though performance in adverse conditions may be poorer. They also exhibit poorer performance at very short ranges. It is unlikely that a single sensor would ever be deployed as a sole source of navigational information for any unmanned aircraft, and UAS will always much more likely use a suite of sensors, which an optical Tau sensor could be part of.

## 2.5 Tau in Guidance and Control

Implementations of control systems using time-to-contact as an input exist as far back as 2008, such as the mobile robot system developed by Kai, Shimada and Ito in reference [80]. This system used two robots: one manually controlled 'leading' robot and one 'follower' robot. A system of light bulbs and light sensors was used as to actively estimate time-to-contact between the two robots. The inverse of this Tau value was then used as the error signal in a proportional feedback system that controlled the speed of the follower robot. They also tested the robots in a 'fleeing' task, where the follower robot attempted to maintain a Tau margin in front of the manually controlled robot, instead of behind it. The system performed well in both scenarios but had to be used in a dark room to prevent light pollution from disrupting Tau measurements.

An early example of Tau guidance in an aerospace context was presented by Voskuijl, Walker, Manimala and Gubbels in 2010, described in reference [90]. They created an automatic system to land a nonlinear simulation model of a Bell 412 helicopter onto a static frigate deck. The deck was placed at a known height, and time-to-contact was calculated from spatial variables. An intrinsic Tau guidance system was integrated into the height rate controller of the Bell 412 to effect the landing. The landing was a standard NATO approach and landing as described in reference [91]; the helicopter hovered alongside the deck, laterally repositioned above the deck, then descended straight down to land. Both the lateral reposition and descent to land manoeuvres followed Tau guides. The manoeuvres were assessed against Mission Task Element (MTE) specifications developed in previous piloted flight simulation trials, and the system performed within in desired limits. A test pilot observing the manoeuvres also

commented that the motion seemed natural to him, indicating it provided a good analogue to human pilot behaviour.

Ridgway developed two Tau-based pilot models in his PhD thesis in 2012 and applied them to a light, fixed-wing training aircraft. Similar to [90], he generated reference Tau trajectories and used them to drive a Translational Rate Command (TRC) control system using error minimisation [21]. The aircraft flew a circuit of an airfield, so a range of manoeuvres were performed including a take-off, climb, banked turns, descent and landing flare. A perfect control prediction method was also implemented for roll manoeuvres that computed the perfect control input to follow a Tau guide exactly taking into account atmospheric conditions and the variation of aerodynamic derivatives with aircraft state.

One of the first operational Tau control systems on aerospace hardware was developed by Kendoul and is described in reference [30]. His 'TauPilot' system was implemented on a quadrotor UAV to perform various complex manoeuvres including braking, docking and landing. The system is split into three separate parts:

1. Tau-Navigation system: responsible for computing the Taus of the different gaps required for a desired manoeuvre
2. Tau-Guidance system: generates target or reference Tau trajectories to be tracked by the Tau-Control component
3. Tau-Control system: computes the control inputs that will force the measured Tau values from the Tau-Navigation system to track the reference ones generated by the Tau Guidance System

The Tau-Navigation system estimates time-to-contact from a desired target point and fused GPS and INS data for Tau values in the horizontal plane. In the vertical axis fused barometric altimeter and INS data is used, so all Tau values are computed spatially. The Tau-Guidance system can produce Tau trajectories for constant  $\dot{\tau}$  manoeuvres, intrinsic Tau manoeuvres, and coupled manoeuvres in multiple axes. Kendoul examined a series of control laws for the Tau-Control system to deal with the singularity problem inherent to Tau controlled systems. A simple proportional feedback controller was found to be ineffective for intrinsic Tau guide manoeuvres

due to the singularity in Tau at  $t = 0$ . An inverse proportional feedback controller can perform intrinsic Tau guide manoeuvres but performs poorly in the terminal phase of any Tau manoeuvre when  $\tau \rightarrow 0$ , and hence control input tends toward infinite. To solve this problem, two options were examined: a combination of the inverse and proportional control laws switched at a threshold value of time-to-contact, or a non-linear ratio control law, shown in equation (2-54).

$$u_c = \begin{cases} K_p \left[ 1 - \frac{\tau_{ref}(t)}{\tau(t)} \right] & \text{if } \tau(t) \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (2-54)$$

The non-linear ratio control law was found to perform well in simulations and hardware testing but was still unpredictable very near the target point when distance, velocity and Tau are all close to zero. A saturation function was therefore added to avoid aggressive control inputs and oscillations in this terminal phase. The saturation function used is displayed in equation (2-55).

$$u_c = \begin{cases} K_p \left[ 1 - \frac{\tau_{ref}(t)}{\tau(t)} \right] & \text{if } t < \frac{9.5}{10} T_G \\ K_p c_1 \operatorname{atan} \left[ c_2 \left( 1 - \frac{\tau_{ref}(t)}{\tau(t)} \right) \right] & \text{otherwise} \end{cases} \quad (2-55)$$

Where  $c_1$  and  $c_2$  are positive constants used to tune behaviour in different axes. Kendoul also reported problems with manoeuvre initiation due to the singularity in Tau at  $t = 0$  for second-order intrinsic Tau guides; the singularity caused large error signals in the controller that had to be saturated out to prevent loss of control of the aircraft. The TauPilot system was implemented on a custom-built quadrotor and the paper reports thousands of successful test flights using Tau manoeuvres to stationary target points. The system also performed docking manoeuvres with a moving target, though it is unclear how the position of this moving target was ascertained by the system. Kendoul also implemented a constant optical flow divergence guidance system on the quadrotor for the sake of performing comparative flight tests. The paper reports this alternative performed worse than the TauPilot system at the start and end phases of equivalent manoeuvres and led to fewer successful flights.



A different approach to generating Tau trajectories was proposed in a 2014 paper by Zhang, Zhang, Xie and Ma, based on the idea of the third-order intrinsic Tau guides discussed in section 2.2.5.3 [74]. These Tau guides mitigated the manoeuvre initiation problems faced by Kendoul since a third-order guide specifies there is no acceleration at the start of the Tau guide. This approach was applied to a simulation model of a quadrotor designed to perch on a tree branch like a bird. Time-to-contact was spatially computed, but no details about the controller were given.

A 2016 paper by Yang, Fang and Li applied Tau trajectories to unmanned aircraft in swarms, as a tool to prevent collisions between participants in the swarm [92]. The system generated and followed Tau guidance trajectories locally on each aircraft, rather than on a central control station as is common with swarms. They assessed various Tau strategies for their purposes, including constant  $\hat{\tau}$  approaches, Tau coupling, second-order intrinsic guides, and third-order intrinsic guides. They also assessed a harmonic intrinsic Tau guide which leveraged sine functions to allow for better shape adjustment in the velocity profile of the coupled motion. They settled on using a modified second-order intrinsic guide to allow for non-zero initial and final velocities in manoeuvres. The Tau guide was only used to generate spatial trajectories to be followed by a position controller and did not use time-to-contact as a control variable. They compared their system against a Decentralized Model Predictive Control (DMPC) system in simulation and found their Tau based trajectories provided safer flights at a lower computational cost to each swarm member.

A recent example of robot control using optically sensed time-to-contact was presented by Zhang in 2016 [89]. A wheeled robot with a camera was used to measure time-to-contact with an object on the ground ahead of it and dock with that object. Time-to-contact was estimated using the gradient method detailed in section 2.4.2.2.1, with rotational rate corrections. The robot did not follow a trajectory generated by any Tau theory guidance law. Instead, it attempted to track to a constant time-to-contact near zero.

Zhang developed this system further in 2017 to use a constant  $\hat{\tau}$  approach that he modified slightly for a two-phase manoeuvre. The first phase decelerated with  $\hat{\tau}$

constant to a non-zero Tau value, the second phase maintained a constant velocity to make positive contact with the target object [93]. Zhang applied a similar system to a quadcopter in 2017, creating one of the first examples of an aircraft using optically sensed time-to-contact to follow Tau trajectories. To reduce noise in the Tau estimation process, he added a Kalman filter to the system [94].

Two other complete systems implementing Tau guidance on unmanned aircraft are presented in references [95] (simulation) and [96] (hardware). Both implement Tau guidance on quadrotor aircraft using April Tags [33] to localise the aircraft with respect to a target and follow spatial reference trajectories to that target.

It is clear that Tau theory can be usefully implemented to guide unmanned aircraft and more generally to vehicles in the wider field of robotics. However, examples that sense time-to-contact optically and use it directly in the nature-inspired guidance laws are rare. It is much more common for systems to produce spatial estimates of Tau, or generate Tau-inspired spatial trajectories to follow with classical position control methods. It is possible to gain some of the benefits of Tau guidance with these approaches, the foremost of these being four-dimensional control, the ability to make soft contact with targets and to control acceleration strategy with a single intuitive parameter.

However, they do neutralise one of its greatest benefits; Tau guidance in its natural form is inherently reactive to dynamic situations. Braking motorists, the original inspiration for Tau theory, exhibit this behaviour constantly as they react to other vehicles on the road. This was clearly considered in Lee's original paper discussed in section 2.2. Consider two cars approaching a traffic light at a constant speed, one following the other. If the light changes to red the lead car will initiate a constant  $\dot{\tau}$  deceleration manoeuvre to stop at the light. As long as they stop without crossing the light they may vary the manoeuvre duration and coupling constant to whatever values suit the driver and vehicle, but the following driver is limited to some degree by what the lead vehicle does if they want to avoid a collision. This is not a problem if the driver of the following car can freely observe the lead car and effect their own Tau based deceleration relative to the first car to avoid a collision.

Consider now the same scenario taking place in heavy fog, with all the leading cars lights broken. The follower can still spot the traffic light turning red through the fog and effect a Tau manoeuvre to stop at it, but without being able to sense time-to-contact with the leading car, they cannot guarantee they won't hit the other car. This is analogous to sensing time-to-contact using only spatial means; only using the position and velocity of the following car to compute time-to-contact will not necessarily provide a collision-free course. The same is true for the common method of following spatially computed trajectories in an inertial frame of reference since they will not move if the target does. To gain all the benefits of Tau-based guidance time-to-contact must be sensed in real-time and applied to a Tau controller.

Reaction to environmental dynamics may be considered less important for aircraft than ground-based vehicles, but there are still many situations where this capability is important. It is vital in both fixed-wing and rotary-wing landings to have accurate information about any potential landing location to ensure safe landing. Reference [90] describes a ship deck landing system gives, illustrating one such scenario where knowledge of the landing target is important; the helicopter must be able to react to the heaving, rolling and pitching motion of the deck to land without excessive relative velocity between the two. Smaller UAS often land at unprepared sites, so even if static the altitude or slope of the landing site needs to be known accurately for a safe landing. Sensing Tau in real-time is one method of gaining this information in an intuitive and simple to use way. Referring back to the ALFURS scale in Figure 2-1, obstacle and collision detection and avoidance is vital for increasing autonomy and External System Independence (ESI). Tau guidance provides a framework for performing those functions.

Notably, operational examples in the aerospace domain are exclusively limited to rotary-wing aircraft, despite the evidence of Tau usage in fixed-wing piloting presented by Jump in reference [79]. The only example of Tau control on a fixed-wing airframe is Ridgway's pilot-modelling system used to control a simulated light training aircraft in reference [21]. This is perhaps because of the relatively low adoption of small fixed-wing airframes by the general public when compared to multi-rotors, but military platforms are largely fixed-wing.

## 2.6 Technical Review Summary

In section 2.1 some basic terminology has been established and a number of issues related to the reliability of UAS operations have been identified. In section 2.2 the basics of Tau theory have been established as a guidance mechanism used by humans and animals, and section 2.3 discusses evidence that this mechanism is also applied by humans when piloting aircraft. Section 2.4 discusses a number of computer vision techniques that can be used to measure Tau. Section 2.5 explores existing examples of Tau guidance systems in the literature. The work described in the remainder of this thesis seeks to address some of the issues described in section 2.1.4. Specifically, optical Tau guidance is applied to UAS landing manoeuvres in an effort to:

1. Improve reactivity to external obstacles and hazards, especially dynamic objects
2. Provide a mechanism for reliably delivering soft landings for both rotary- and fixed-wing UAS platforms
3. Be able to execute both the above functions without the need for external signals, thus reducing reliance on sometimes unstable communication links

In the course of demonstrating these capabilities novel improvements to Tau perception methods will also be presented that expand the operational envelope of Tau sensors based on direct gradient algorithms. Techniques for predicting the performance of aircraft following Tau guided manoeuvres will also be developed that allow for better prediction of how effective Tau guidance can be.

## 3 Chapter 3 – Experimental Setup

This chapter describes the various tools and simulation techniques used during the research reported in this thesis. Investigation of Tau guidance was carried out using aircraft simulation models from several sources. Tau perception techniques were tested against virtual visual scenes generated using computer graphics techniques. The development of these tools is described. Finally, some Mission Task Elements are defined for testing of Tau perception techniques and Tau guidance systems.

### 3.1 Simulation Software and Hardware

Flight simulation is a key technology for the work described in this thesis as it allows for detailed testing of Tau guidance systems without risking any real hardware. Several tools for flight simulation were considered, including FLIGHTLAB, a flight vehicle modelling and analysis tool developed by Advanced Rotorcraft Technologies which is used as the basis for the University of Liverpool's Heliflight simulators [97]. FLIGHTLAB is often used when real-time simulation capability is required, but since this functionality is unnecessary for UAS simulation, Matlab Simulink was chosen as the modelling environment for all of the aircraft models used during the course of this thesis.

Simulink is a graphical programming environment for modelling, simulating and analysing dynamic systems [98]. It allows for fast generation and testing of flight models, control systems and image processing methods, so is ideally suited for the work to be carried out. However, Simulink is often not the most efficient tool for execution of simulation models and image processing methods due to large overheads associated with the user experience. This is particularly relevant when examining the image processing components of this thesis related to Tau perception, where algorithms must be capable of running many times a second for the system to be viable for hardware implementation. The Simulink versions of various Tau perception tools in this thesis often run much slower than desired and this will become apparent during the course of chapter 4. The author does not consider this to be a problem that would stop the use of any of these Tau perception mechanisms

in future because Simulink is well known to be an inefficient tool in this regard, and the computer used is not well optimised for the specific tasks required. To contextualise the execution rate of the programs that will be presented in chapter 4, the hardware specifications of the computer all of the simulation work was performed on are given in Table 3-1. The computer is a mid-range desktop computer designed for computer gaming and was built approximately 5 years ago.

*Table 3-1: Hardware specification for main computer used for simulation work detailed throughout chapters 4 and 5. Machine is a mid-range desktop computer approximately 5 years old.*

<b>CPU Model</b>	Intel Core i7-4771
<b>CPU Cores</b>	4 Cores
<b>CPU Clock Speed</b>	3.50 GHz
<b>RAM Memory</b>	8.00 GB DDR3
<b>RAM Speed</b>	1600 MHz
<b>GPU Model</b>	AMD Radeon R9 280X
<b>GPU Clock Speed</b>	1100 MHz
<b>GPU Memory</b>	3.00 GB GDDR5
<b>GPU Memory Clock Speed</b>	1500 MHz
<b>Operating System</b>	Windows 10

To evaluate the performance of Matlab Simulink for this hardware configuration the speed of image gradient calculation, a foundational operation to the Tau perception methods that will be explored in chapter 4, was tested. The gradient computation time was recorded for four, 5 second videos of different resolutions using both the Matlab image processing toolbox and the OpenCV image processing toolbox in Python. OpenCV is an open source library of programming functions for computer vision that is available in a range of programming languages [99]. Python is a high-level interpreted programming language that is similar to Matlab but runs with lower overheads [100]. A second computer was also used a comparison point to the first. The second computer is a mid-range laptop, 6 months old at the time of writing and its specifications are given in Table 3-2.

Table 3-2: Hardware specification for 2nd computer used for comparison of gradient analysis computation time. Machine is a mid-range laptop, 6 months old at the time of writing.

<b>CPU Model</b>	Intel Core i7-10510U
<b>CPU Cores</b>	4 Cores
<b>CPU Base Clock Speed</b>	1.80 GHz
<b>CPU Boost Clock Speed</b>	4.90 GHz
<b>RAM Memory</b>	8.00 GB DDR3
<b>RAM Speed</b>	1600 MHz
<b>GPU Model</b>	Intel UHD Integrated Graphics
<b>Operating System</b>	Windows 10

The average time to compute the image gradients of each frame for each language and each computer are shown in Figure 3-1.

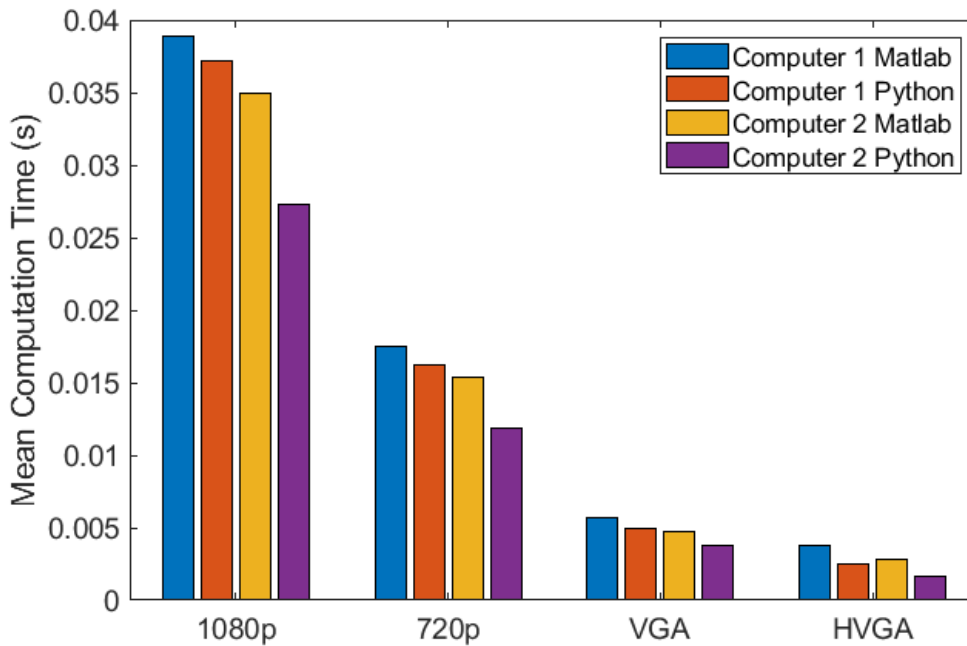


Figure 3-1: Mean frame processing time for image gradient analysis test at four different video resolutions. Results for two computers running two different programming languages are shown to provide points for comparison.

Two main conclusions can be drawn from the data: the gradients are computed slightly faster in Python than they are in Matlab; and the computer with a more

modern processor is also able to compute the gradients faster. As a result, it is safe to conclude that faster execution of image processing routines is possible with up to date hardware and more efficient programming languages. Both Matlab and Python are interpreted languages, so even faster performance should be easily possible with a compiled language like C. However, computation time may be a limiting factor for high resolution images if high frame rates are required.

## 3.2 Aircraft Simulation Models

A number of simulation models of different aircraft were used through the work described in this thesis and all of them will be detailed in this section.

### 3.2.1 3DR X8

The University of Liverpool operates 3DR X8 UAS platforms for the purpose of testing guidance, navigation and control systems. The X8 is a small-scale, commercially available, co-axial octocopter which can be seen in Figure 3-2. The platform is based around the Pixhawk flight controller running the open-source autopilot software PX4.



*Figure 3-2: 3DR X8 co-axial quadcopter RUAS platform operated by the University of Liverpool [101]*

The specifications of the X8 are summarised in Table 3-3.

*Table 3-3: Specifications of 3DR X8 RUAS platform [101]*

<b>Total Mass</b>	2.85 kg
<b>Rotor Diameter</b>	8 x 0.254 m



<b>Endurance</b>	12 minutes
<b>Aircraft Category</b>	0-7kg Multi-Rotor SUSAs
<b>Flight Controller</b>	3DR Pixhawk Autopilot

The Pixhawk controller was specifically designed for efficient operation with PX4 firmware to allow high speed, low latency onboard image processing [40]. The lightweight system design is suited toward MAV applications and allows for the use of up to four cameras for localisation, pattern recognition and obstacle avoidance.

A simulation model of the University of Liverpool’s 3DR X8 UAS platform has been developed for use in this project. The Mathworks SimMechanics toolbox was used to extend the functionality of Simulink for multi-body dynamics simulation in order to create a physics-based model of the aircraft [102]. This approach was chosen to create a representation of the aircraft that was able to physically interact with other objects such as landing decks. Since SimMechanics can be integrated into the Simulink 3D Animation toolbox [103], 3D animations of the simulation models can be generated and actuated.

A 3D model of the X8 platform was generated using CAD package ProEngineer Wildfire 4.0 from physical measurements of individual aircraft components. This model was then exported from ProEngineer using the SimMechanics export utility to an XML file format containing mass, inertia and dimension information. Mobile joints were also defined in the CAD model to allow for motion of motors.

The resulting XML definition file was then imported to SimMechanics to create the basis of the flight dynamics model. This approach accurately depicts the mass and inertia characteristics of the airframe and allows for easy integration of aerodynamic forces into the model. The main body of the aircraft is treated as one entity as it contains no moving parts, though it is comprised of multiple components. Each arm of the multirotor is grouped separately for ease of navigation through the model, with two motors per limb. The camera gimbal and battery were also grouped separately to allow for easy interchange, since these components may be removed or changed for different scenarios. All component positions are referenced from a

point on the main body lower assembly plate, as this was the starting point for construction of the aircraft. The reference point is situated on the plate centreline, at the nose of the aircraft and on the underside of the plate.

The solid components are linked by physical connections, while throttle inputs are provided to each motor as standard Simulink signals. Finally, the height of the aircraft above the terrain is provided to each rotor. Revolute joints with one rotational degree of freedom are actuated by a torque scaled according to a throttle signal. The rotational rate of the joint is measured and fed into a lookup table to determine the amount of thrust produced by the rotor. The thrust lookup table is populated by experimental data obtained from testing of the motor and rotor in a wind tunnel. The thrust signal is multiplied by a gain to ensure correct directionality and then further multiplied by a scaling factor to account for ground effect. Ground effect was modelled by relating a thrust ratio to the size of the rotor and its distance from the ground using the following approximation [104]:

$$\frac{T}{T_{\infty}} = \frac{1}{1 - \left(\frac{R}{4Z}\right)^2}$$

After this correction has been applied, the resultant thrust signal is applied to the rotor as a force normal to the rotor disk. A ground contact modelling system was added to allow for physical landing simulations. This system tracks a reference point on the aircraft and compares it to the height of the ground or deck at each simulation time step. If the reference point passes below the threshold ground value, forces are applied. Vertical force on the aircraft is modelled as a spring and damper with stiffness and damping coefficients selected to give a critically damped response when the aircraft impacts the ground plane with only the minimum amount of deflection possible. Frictional forces are also applied in the horizontal plane to arrest motion. A collision box, shown in Figure 3-3, was defined around the aircraft to better model the effect of different contact points. Forces and moments are calculated based on the position of the collision box vertices. This was important as the unpredictable nature of ship deck motion may lead to certain legs of the aircraft making contact with the deck before others during pitch and roll motion.




Figure 3-3: Collision box generated around the 3DR X8 model airframe. Vertices of the box represent the legs of the aircraft and the extremes of the rotor positions

A simple control system inspired by the flight stack of the PX4 firmware has been constructed to perform initial testing on the flight dynamics model. This will be used for comparison with flight test data of the real-world airframe for validation. The first component of this was a control mixer, used to read conventional control inputs of throttle, pitch, roll and yaw commands, and translate them to motor speed control outputs. For the X8 platform the four control inputs alter motor speeds by adding or subtracting control inputs according to the methodology in Table 3-4 and multiplying by a scaling factor. An idle throttle value is also added to every channel and saturation is applied to limit the upper and lower bounds of each signal. The motors are numbered clockwise from the front right arm of the aircraft, with the first motor above and the second below the arm.

Table 3-4: 3DR X8 Control Mixing Matrix

	Motor Speed Command							
	M1	M2	M3	M4	M5	M6	M7	M8
<b>Throttle</b>	+	+	+	+	+	+	+	+
<b>Pitch</b>	-	-	+	+	+	+	-	-
<b>Roll</b>	-	-	-	-	+	+	+	+
<b>Yaw</b>	+	-	-	+	+	-	-	+

The PX4 flight stack consists of an attitude controller and a position controller. The attitude controller is made up of pitch, roll and yaw rate PID controllers, which are given reference pitch and roll rates by pitch and roll angle PID controllers. Yaw rate is provided directly as a control input. A basic position controller has been added consisting of a translational speed PID and controllers and x, y and z position controllers, fed by a signal building utility to set up basic manoeuvres.

### 3.2.2 3DR Aero

The 3DR Aero is a small fixed-wing UAS platform operated by the University of Liverpool for research purposes. The aircraft is based on the popular Skywalker V6 airframe with a 3DR Pixhawk autopilot system, which runs the PX4 flight controller. An example of the aircraft is shown in Figure 3-4.



*Figure 3-4: 3DR Aero fixed-wing UAS platform operated by the University of Liverpool [105]*

The specifications of the Aero are summarised in Table 3-5.

*Table 3-5: Specifications of 3DR Aero fixed-wing UAS platform*

<b>Total Mass</b>	2.20 kg
<b>Wing Span</b>	1.88 m
<b>Length</b>	1.29 m
<b>Endurance</b>	40 minutes
<b>Aircraft Category</b>	0-7kg Multi-Rotor SUSA
<b>Flight Controller</b>	3DR Pixhawk Autopilot

A simulation model of the Aero was produced by Liverpool's Virtual Engineering Centre using FLIGHTLAB modelling software. A 3D scanner was used to create a model of the aerodynamic surfaces and fuselage section. Computational Fluid Dynamics (CFD) was used to populate lookup tables with aerodynamic data. Since the FLIGHTLAB component library does not include an electrical engine component, one was created using experimental data. The detailed development of this model is described in reference [106]. This FLIGHTLAB model was linearised to produce state-space models of the Aero at a range of speeds for use in Matlab Simulink.

The author created a model of the PX4 fixed-wing flight controller software in Matlab Simulink to link with the FLIGHTLAB model. The autopilot model is split into two subsystems, a stabilisation module and a navigation module. The stabilisation system consists of attitude control loops for each rotational axis and a control mixer to map control inputs to servo outputs. The navigation system implements waypoint following and calculates pitch, roll, yaw and throttle demands to pass to the stabilisation system. Multiple flight control modes are implemented for different levels of autonomy, including automatic waypoint following, airspeed and heading hold, attitude command and manual pass through.

### 3.2.3 RUAS Linear Models

In order to investigate the implementation of Tau control systems on a range of aircraft, four rotorcraft flight dynamics models of different aircraft were used. They were all developed in FLIGHTLAB by other researchers in the course of previous projects. The four rotorcraft models were based upon:

- (i) Sikorsky SH-60B 'Seahawk'
- (ii) Northrop Grumman MQ-8B 'Fire Scout'
- (iii) Yamaha R-MAX
- (iv) Align T-Rex 700

Each helicopter model makes use of a blade-element model for the main rotor with a Peters-He inflow model, a Bailey rotor for the tail rotor, and fuselage and empennage aerodynamic look-up tables. The R-MAX and T-Rex models, in addition, have stability bars which are modelled as rate feedback gain in the roll and pitch channels.

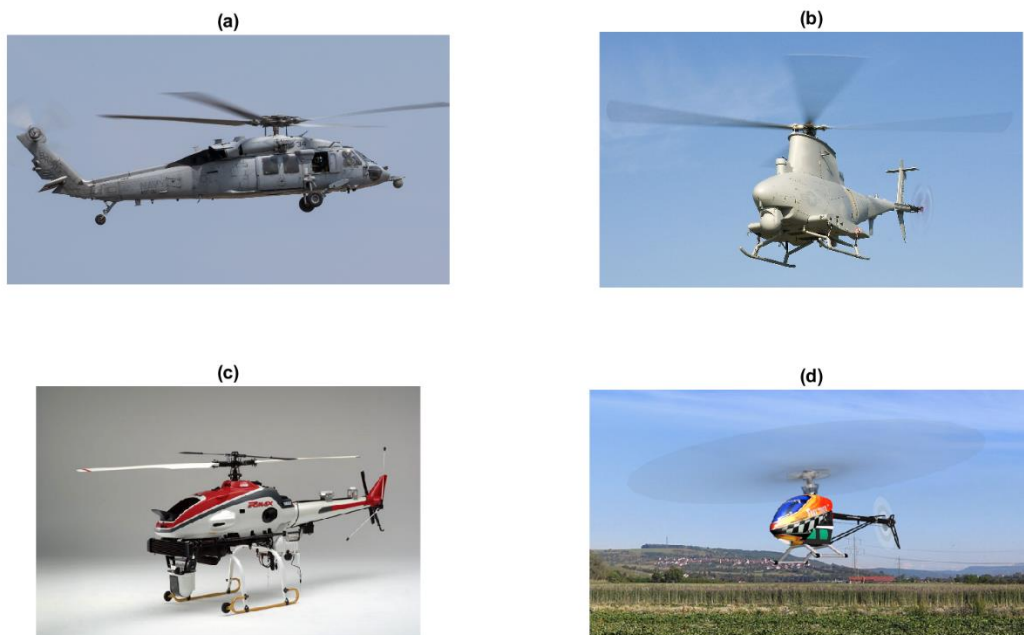


Figure 3-5: The aircraft that the four RUAS models were based on: (a) Sikorsky SH60B [107] (b) Northrop Grumman MQ-8B Fire Scout [108] (c) Yamaha R-MAX [109] (d) Align T-Rex 700 [110]

The Seahawk model, though not typically an unmanned aircraft, is used to represent a full-scale aircraft in the analysis. It was derived from the FLIGHTLAB Generic

Rotorcraft Model (FGR), which, in turn, is based on a UH-60A model and is well documented in the literature [111]. The MQ-8B Fire Scout is an unmanned air and sea support platform based on the Schweizer 333. The FLIGHTLAB model was estimated as a scaled version of the FGR helicopter model to give the same disc-loading as the real aircraft using the main rotor as the scale reference length. The R-MAX is a remotely piloted unmanned helicopter often used in agriculture for crop spraying and as a research platform. The FLIGHTLAB model was developed by DSTG Australia and has had limited internal validation against the on-axis response to controls [55]. The T-Rex 700 is a remotely piloted aerobatic '3D' helicopter. The FLIGHTLAB model was estimated as a scaled version of the R-MAX model and has had limited validation against flight test data provided by the NRC [55]. The salient parameters of these aircraft are displayed in Table 3-6.

*Table 3-6: RUAS simulation model base parameters*

<b>Aircraft</b>	<b>NATO Class</b>	<b>Rotor Diameter (m)</b>	<b>MTOW (kg)</b>	<b>Length (m)</b>
SH60B	Class III	16.35	10,400	19.75
MQ-8B	Class III	8.4	1,430	7.3
R-MAX	Class I	3.115	94	3.63
T-REX	Class I	1.582	5.1	1.32


The full non-linear FLIGHTLAB models were linearised to create nine state linear state-space models of each aircraft for use in Matlab Simulink. All four aircraft were trimmed in a hover condition at an altitude of 100 feet before being linearised, with no incident wind. These models were then interfaced with virtual visual environments and Tau guidance and control systems.

As with the 3DR Aero: only the linear models were used for the work described in this thesis, not the full nonlinear models. This decision was made to ease implementation into a Matlab Simulink testing environment and allow for fast simulation times. As a result, some physical phenomena, such as ground effect, are not adequately modelled. Testing with full non-linear simulation models should be carried out in future to discover any potential problems and inform implementation

and testing. All of the information needed to reconstruct the SH-60B, MQ-8B and 3DR Aero linear state-space models are given included in Appendix A, including a brief description of their implementation. The source of the R-MAX and T-REX models is considered sensitive so those models cannot be released at this time.

### 3.3 Camera Modelling

All the tools used to virtually model video sequences are underpinned by a simple camera model. Cameras collect light rays that have reflected off objects in the outside world and record the information contained in those rays on to some form of media. A typical digital camera can be broken down into several components: an aperture, an optical system and a sensor array, arranged linearly along the optical axis as diagrammed in Figure 3-6.



*Figure 3-6: Simple model of essential camera elements: aperture, optical system and sensor array, arranged along the optical axis*

The aperture controls the amount of light being collected by the camera to prevent overexposure of the image. If the aperture is too large, then multiple rays of light from a single point in the visual scene may enter the camera and cause the image to blur. A common assumption is to treat the aperture as a pinhole, allowing only a single light ray from each point in the visual scene into the camera. The optical system is comprised of one or more lens and serves to focus light rays onto the sensor array. The simplest optical system is a single thin lens. Any light rays entering a thin lens parallel to the optical axis will converge on a single point when leaving the lens, known as the focal point. The opposite is also true; any light rays entering the lens



from the focal point will exit parallel to the optical axis. The distance between the centre of the lens and its focal point is known as the focal length. As a result, the aperture will usually be positioned one focal length from the lens. The sensor array records the intensity of the incoming light with a grid of photoreceptors. Most commercially available cameras use either a Charge-Coupled Device (CCD) or a Complementary Metal-Oxide-Semiconductor (CMOS) sensor array [82]. Both can be regarded as a box that generates a voltage from incident light energy. The voltage produced by each cell in the sensor array is scanned line by line, and the values are quantised on to an integer scale and arranged into a matrix that forms the digital image. For colour images each element the sensor array outputs a vector of three values instead of one; one corresponding to the intensity of the Red, Green & Blue (RGB) light components on the sensor. In video capture, multiple images, known as frames, are captured sequentially many times a second to build the video.

### 3.3.1 Camera Modelling

The fundamental equations of geometric image formation relate the position of a point in the visual scene,  $P = [X, Y, Z]$ , to a corresponding point in the image,  $p = [x, y, z]$  and are detailed in [82]. Using the pinhole camera model, these points can be related to each other in the camera reference frame using similar triangles. It is important to note that in the camera reference frame, the  $Z$  axis is defined along the optical axis and  $z = f$ , the focal length of the camera. Therefore, the fundamental equation of the camera frame is:

$$\frac{P}{Z} = \left[ \frac{X}{Z}, \frac{Y}{Z}, 1 \right] = \frac{p}{z} = \left[ \frac{x}{f}, \frac{y}{f}, 1 \right] \quad (3-1)$$

However, the camera frame is of limited usefulness because objects in a scene are much more likely to be defined by an exterior inertial frame of reference. Therefore, it is necessary to define a transformation between the world and camera frames:

$$P_{camera} = R(P_{world} - T) \quad (3-2)$$

$R$  and  $T$  in equation (3-2) are arrays describing the rotation and translation of the camera relative to the world frame of reference. These are known as the extrinsic

parameters of the camera. A further transformation from the camera frame to the image frame is required to get the position of the image point in pixel coordinates:

$$x_{camera} = -(x_{image} - o_x)s_x \quad (3-3)$$

$$y_{camera} = -(y_{image} - o_y)s_y \quad (3-4)$$

Where  $(o_x, o_y)$  are the coordinates of the image centre in the image frame, and  $(s_x, s_y)$  are the effective size of each pixel on the sensor array, usually measured in millimetres. These parameters, along with the focal length  $f$ , are the intrinsic parameters of the camera. Further intrinsic parameters can be integrated into the model that account for distortion introduced by optical systems, but those will be omitted here. Combining the extrinsic and intrinsic parameters of the camera into a single equation defines the image frame location of an observed point in the world frame:

$$-(x_{image} - o_x)s_x = f \frac{R_1^T(P_{world} - T)}{R_3^T(P_{world} - T)} \quad (3-5)$$

$$-(y_{image} - o_y)s_y = f \frac{R_2^T(P_{world} - T)}{R_3^T(P_{world} - T)} \quad (3-6)$$

Where each  $R_i$  term refers to the  $i$ th row of the rotation matrix  $R$ . Equations (3-5) and (3-6) can be expressed as a linear matrix equation, splitting the intrinsic and extrinsic parameters into two separate terms:

$$\begin{aligned} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} &= \begin{bmatrix} -\frac{f}{s_x} & 0 & o_x \\ 0 & -\frac{f}{s_y} & o_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} R_{11} & R_{21} & R_{31} & -R_1^T T \\ R_{21} & R_{22} & R_{32} & -R_2^T T \\ R_{31} & R_{23} & R_{33} & -R_3^T T \end{bmatrix} \begin{bmatrix} X_{world} \\ Y_{world} \\ Z_{world} \\ 1 \end{bmatrix} \\ &= M_{intrinsic} M_{extrinsic} \begin{bmatrix} X_{world} \\ Y_{world} \\ Z_{world} \\ 1 \end{bmatrix} \end{aligned} \quad (3-7)$$

Where the image coordinates in the  $x$  and  $y$  directions are the ratios of  $x_1$  and  $x_2$  to  $x_3$  respectively. This method also separates the transformations between the different frames of reference neatly, with  $M_{extrinsic}$  converting from the world frame

to camera frame, and  $M_{intrinsic}$  converting from camera frame to image frame. This equation can be used to create a simple computer graphics engine to produce virtual images of imagined scenes, which is very useful for analysing image processing techniques.

### 3.3.2 Camera Model Parameters

It is convenient to define the parameters of a particular camera for use throughout the research described in this thesis for the sake of consistency. A Raspberry Pi camera module, version 2 was chosen for this purpose, shown in Figure 3-7.



*Figure 3-7: Raspberry Pi camera module version 2 used as the source for camera modelling parameters [112]*

This camera is designed for use with the Raspberry Pi microcomputer, which is commonly used for image processing. The specifications of the camera are summarised in Table 3-7.

*Table 3-7: Specifications for Raspberry Pi camera module v2 [113]*

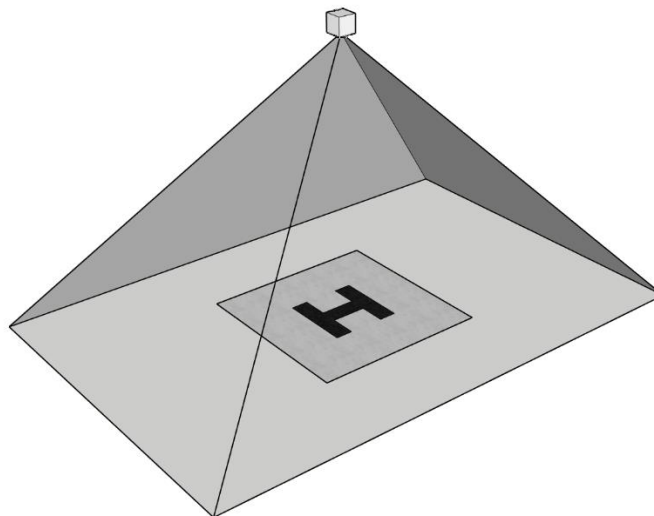
<b>Still Resolution:</b>	8 Megapixels
<b>Video Modes:</b>	1080p30, 720p60, 640x480p90
<b>Sensor Image Area:</b>	3.68 x 2.76 mm (4.7 mm diagonal)
<b>Cell Size:</b>	1.12 $\mu\text{m}$ x 1.12 $\mu\text{m}$
<b>Focal Length:</b>	3.04 mm
<b>Maximum Horizontal Field of View:</b>	62.2 degrees

<b>Maximum Vertical Field of View:</b>	48.8 degrees
--	--------------

### 3.3.3 Virtual Helipad Landing Scene

Since rotary-wing landings are of interest in this research, a simple depiction of a helipad for use in analysing different Tau estimation techniques was created using the mathematics described in section 3.3.1. The imagined helipad is a square of flat colour five metres square on a black background, with an 'H' in the centre. A three-dimensional diagram of the scenario is shown in Figure 3-8.

Of course, for any point to be imaged by the camera, it must be within the field of view of the camera. In 3D computer graphics the field of view of the camera is usually represented by a viewing frustum. This is the region of space in the virtual world that will actually be captured by the camera. The boundaries of the frustum are superimposed over the helipad image in Figure 3-8.

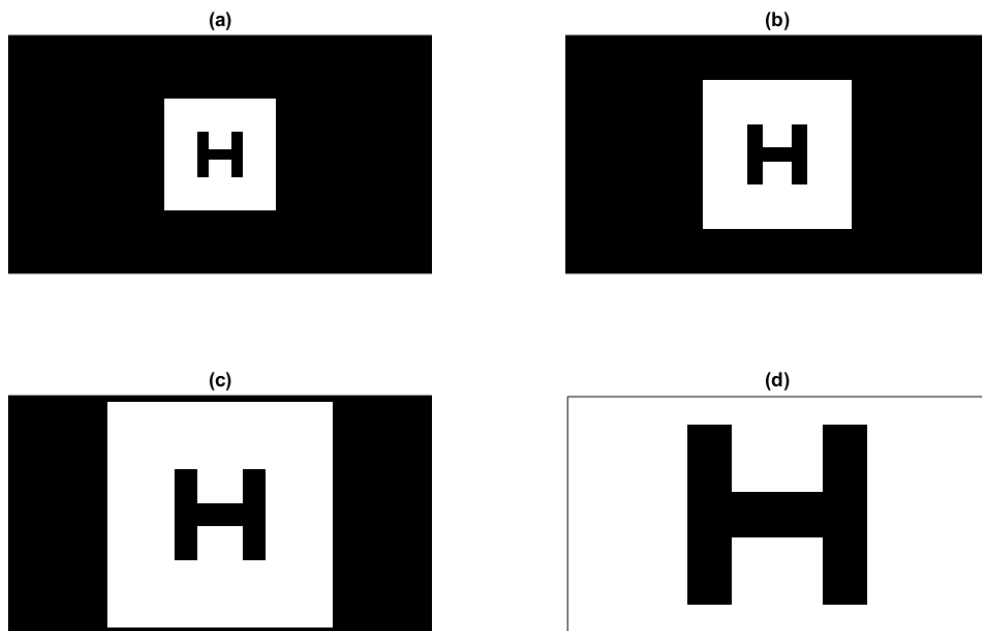


*Figure 3-8: 3D diagram of a camera descending towards a helipad, with the camera frustum superimposed to illustrate how much of the scene is captured by the camera*

A visual scene can be generated for any camera position or rotation in the world frame using this method, so it is simple to produce sequences of images that depict the camera moving relative to the helipad once the points that are visible to the camera have been identified. This can be accomplished by interrogating a database of the terrain to determine which points of the virtual world lie within the volume defined by the viewing frustum by using inequality statements. Once the position of

the helipad vertices is found on the image plane, the interior of the shape can be filled using two-dimensional interpolation over the pixel grid. This model was implemented as a Matlab program that produces video files after defining the camera properties, trajectory and helipad dimensions. The camera resolution used to generate helipad videos can be easily changed, but a 720p resolution with a framerate of 30 frames per second was used for all videos in initial testing.

The dimensions of the helipad can be easily altered, and it is also possible to incline the plane the helipad lies on about both the horizontal axes. This can create complex relative slope situations between the observer and target. Several frames from a video depicting the camera descending straight toward a helipad on flat ground are shown in Figure 3-9 as an example.




*Figure 3-9: Sample frames from a video of a virtual camera descending towards a 5x5 m helipad. All motion is along the optical axis, and the helipad lies on a flat plane. Each frame is recorded at a defined distance above the helipad: (a) 20 m (b) 15 m (c) 10 m (d) 5 m*

### 3.3.4 Virtual Runway Landing Scene

To assess fixed-wing aircraft landings, a virtual runway scene was also developed in a similar manner to the helipad described above. The runway provides some

additional utility over the helipad as it can generate images for videos where the camera is displaced parallel to the ground plane by large amounts. The imagined runway is 1000m long and 26m wide, with typical runway markings included as described in [114]. The runway is imagined to run south to north, and this is reflected in the designation markings. Figure 3-10 shows a plan view of the first 200m of this runway.



*Figure 3-10: Plan view of first 200m of a simple virtual runway markings used for the virtual runway landing scene*

Again, the dimensions of the runway and its markings can be easily altered to suit different situations. The inclination of the plane the runway lies on can also be changed about the X and Y axes. Figure 3-11 shows some frames from a video depicting the camera descending towards the runway while also moving down it as an example. Image (a) shows the runway threshold markings, with the designation markings visible in images (b) and (c).

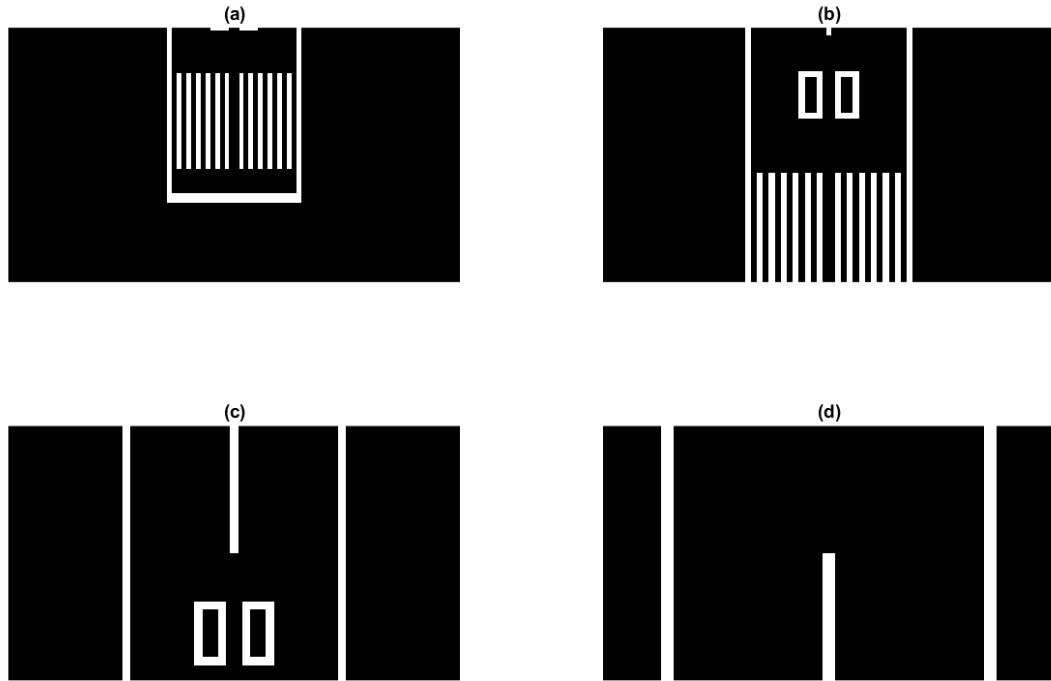


Figure 3-11: Sample frames from a video of a virtual camera descending towards a runway. Motion is downward along the optical axis and forward along the length of the runway. The runway lies on a flat plane. Each frame is recorded at a defined distance above the runway: (a) 100 m (b) 80 m (c) 60 m (d) 40 m

The virtual scenes described in section 3.2.3 are only simple representations of some of the MTEs that will be examined during the course of this thesis. They contain very little information in terms of visual texture, so they are inadequate for some analysis. Improving this simple graphics engine to the point of detailed texturing on three-dimensional objects is far beyond the scope of this thesis. Other tools have been employed to fill this role, primarily the Matlab Simulink 3D Animation toolbox. However, these simple virtual scenes still hold value in the precise control that can be exerted over them, especially in terms of modelling of intrinsic camera properties.

### 3.4 Matlab Simulink Virtual Reality Worlds


The Matlab Simulink 3D Animation toolbox provides tools for linking Matlab programs and Simulink models to 3D graphics objects and scenes. Objects defined in standard modelling languages can be positioned within 3D worlds and viewed using virtual cameras. This toolbox provides an easy to use set of methods for interfacing simulation models with a 3D graphics engine. It can be used to generate video for

defined camera trajectories or can be integrated into real-time simulations to produce video streams for cameras that move dynamically within a 3D world. These video feeds can then be directly processed Matlab and Simulink. A program for creating Virtual Reality (VR) worlds is included with the toolbox called 3D World Editor. Two different VR worlds were created in 3D World Editor for use in this research project: a frigate deck for rotary-wing landings, and a desert runway for fixed-wing landings.

#### 3.4.1 Frigate Deck Landing Scene

A frigate deck landing virtual world was created for use in simulation experiments. A deck was created using Google Sketchup 3D modelling software and textured with generic deck markings inspired by those found on the common frigates. Only the landing deck is modelled as the rest of the ship serves no purpose in any of the experiments described in this thesis. The deck is situated in a flat, blue sea. The position of the deck can be controlled with inputs to the VR simulation, so deck motion in varying sea states can be easily simulated by coupling the world to a Simulink model. A generic helicopter model from the object library in 3D World Editor was used as a platform for a virtual camera. The helicopter visual model is not related to the dynamics of any aircraft simulation attached to the virtual world and serves only to represent the position and orientation of the vehicle, and serve as a reference for the virtual camera pose. It is not visible in the feed provided by the virtual camera. A view of the frigate deck landing scene is shown in Figure 3-12.



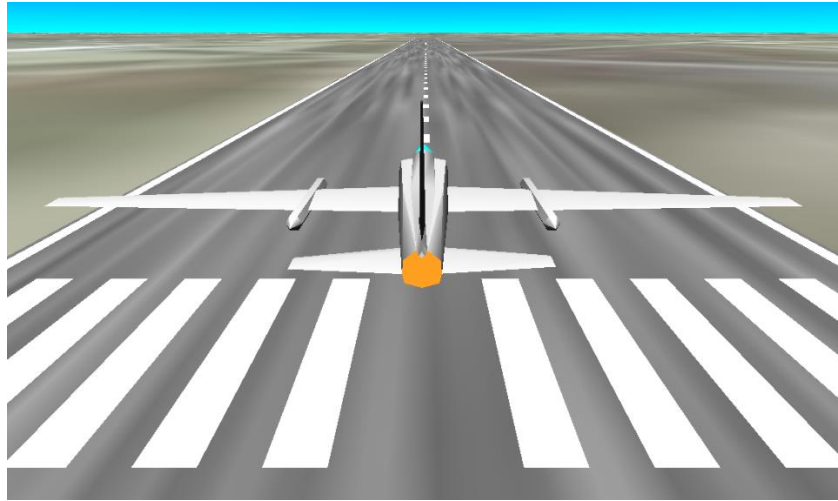


*Figure 3-12: A view of the frigate deck landing scene created in 3D World Editor for the virtual ship deck landing experiments*

Sensors were attached to the landing gear of the helicopter model to indicate when they made contact with the deck model. This contact flag can be fed back to an attached Simulink model for use in a ground contact modelling.

#### 3.4.2 Desert Runway Landing Scene

A second virtual world was used for fixed-wing landing analysis. The desert runway scene was based on a standard example provided with the 3D World Editor with a generic fixed-wing aircraft model from the object library added as a platform for a virtual camera. This VR world was set up in the same manner as the frigate deck landing scene for coupling with a Simulink aircraft model. A view from a chase camera in the VR world is shown in Figure 3-13.



*Figure 3-13: A view of the desert runway scene created in 3D World Editor for the virtual runway landing experiments*

### 3.5 Mission Task Elements

Two MTEs were defined for use during the research. They define short repeatable tasks that represent common operational tasks for UAS. Here, they are a frigate deck landing manoeuvre, and a runway landing.

#### 3.5.1 Frigate Deck Landing

A common approach to helicopter ship deck landings is the fore/aft landing procedure defined in reference [91], diagrammed in Figure 3-14. The procedure is as follows:

1. Approach the port side of the ship and hover alongside in a station keep position. The helicopter's longitudinal axis should be parallel to the longitudinal axis of the ship.
2. Laterally reposition the helicopter over the landing spot on the deck
3. Descend vertically and land




Figure 3-14: Standard helicopter take-off and landing procedures for NATO frigate deck landings [91]

Landings are always performed from the port side of the ship so that the pilot sitting in the right-hand seat has a good view of the ship deck. Take-offs can be performed either way depending on wind conditions. For this research only the third phase of this manoeuvre will be examined. For all of the deck landings examined, the aircraft begins in a hover 10 m above the deck. Deck motion can be varied for different sea states.

### 3.5.1.1 Deck Motion

A physical ground plane was added to all of the rotorcraft simulation models and actuated in order to emulate ship deck motion. Sensors were added to measure its position in relation to the aircraft. Initially, only a single degree of freedom, heave motion in the z-axis, was actuated, but all six degrees of freedom can be actuated for realistic deck motion. Ship motion was generated in ShipMo3D [115] for a generic Type23-like frigate in a variety of sea states use in actuating the ship deck. The heave motion of the landing deck was found by summing the heave motion of the ship as a whole, and the displacement of the landing deck due to pitch motion. Since the temporal resolution of the data was low when compared to that of the simulation model, the data was interpolated using a spline curve to provide smoother data. The vertical displacement of the landing deck over thirty seconds for sea states one to nine is plotted in Figure 3-15.




Figure 3-15: Deck landing spot heave displacement of a Type 23-like ship for varying sea state

### 3.5.2 Runway Landing

The runway landing MTE begins with the aircraft 50 m above the runway threshold, with the aircraft flying on a northerly heading down the runway. The aircraft begins flying straight and level and must descend towards the runway and land. The manoeuvre can be performed at any forward airspeed, but was primarily tested at 15 m/s, the approach speed of the 3DR Aero aircraft.

## 4 Chapter 4 – Tau Perception through Computer Vision

The University of Liverpool's Flight Science and Technology research group has a long history of research into the applications of Tau theory to manned flight, as many of the references mentioned in chapter 2 will attest. However, none of this work has covered systems that sense Tau directly. In section 2.4, it was established that various methods exist for estimating time-to-contact, but it is not immediately clear which of them are viable for use on aerial platforms, or how effective they are. As a result, work was required at the outset of this research project to:

- a) Establish which of the presented methods are most effective for estimating time-to-contact experimentally from aerial platforms
- b) Find the limits of these established methodologies for sensing time-to-contact

Although it has been established that time-to-contact sensing in nature is not always a visual process, the bulk of the literature discussed in section 2.2 posits it as such. Since most animals do not have active ranging senses like echolocation, passive sensing of Tau is of vital importance if it is indeed used as a tool in guidance. The same can be said for the aerospace domain, since UAS are more likely to possess passive sensors like video cameras than active ones like radar, Lidar, or sonar. The methods for sensing time-to-contact with active sensors are also simple without much need for further examination. Though this makes them an obvious choice for the task, they are often expensive, heavy or only useful at limited ranges. Monocular video cameras are an attractive option for any sensor task due to their low cost and wide availability; the absence of a pilot necessitates that the vast majority of UA possess a video camera for basic situational awareness. As a result, this thesis will primarily focus on estimating Tau from sequences of images captured by monocular cameras.

This chapter will primarily analyse the already available techniques for estimating time-to-contact, then demonstrate some techniques for mitigating problems identified during the analysis. All Tau perception techniques will be analysed in the

generalised scenario of a camera descending toward a ground plane, with varying visual scenes projected onto that surface. Fixed-wing platforms introduce an extra layer of complexity to Tau perception due to their high flight speed compared to rotary-wing platforms. This high level of motion parallel to the ground plan conflicts with some of the assumptions required by the Tau estimation techniques, and this will also be explored.

These methods rely on many common image processing techniques, the details of which are beyond the scope of this thesis. Reference [82] provides a good source for this information. The visual modelling tools detailed in sections 3.2.3 and 3.4 are used throughout this chapter to assess the various Tau estimation techniques.

## 4.1 Tau Perception through Dimension Tracking

The dimension tracking method for Tau perception is well-established, as discussed in section 2.4. This section will present the mechanism in greater detail and evaluate its performance and limitations.

### 4.1.1 Overview

The dimension-tracking method for time-to-contact estimation leverages equation (3-1), known as the weak perspective projection equation, to relate a particular dimension in the visual scene to its apparent size in an image of that visual scene. That dimension may be a distance between some point features, or the size of a detected object or shape. Figure 4-1 shows a diagram of a camera looking down toward an object of length  $X$  lying on a ground plane, at a distance  $Z$  along the optical axis.




Figure 4-1: Geometry of a camera descending towards a flat plane with parameters required for Tau estimation through dimension tracking labelled

The size of the object projected on to the image plane  $x$  can be found using the focal length of the camera  $f$ :

$$x = f \frac{X}{Z} \quad (4-1)$$

This diagram shows only  $x$  and  $z$  dimensions, but the third image dimension  $y$  extends into the page, and the size of the object on the image plane in this dimension will be:

$$y = f \frac{Y}{Z} \quad (4-2)$$

If the camera moves along the optical axis, then clearly the size of the projection of the object on the image plane will change. It is assumed that the focal length of the camera is constant; the camera performs no automatic focusing. Differentiating equations (4-1) and (4-2) yields expressions for this change in size in each image axis:

$$\frac{dx}{dt} = f \frac{\dot{X}Z - X\dot{Z}}{Z^2} \quad (4-3)$$

$$\frac{dy}{dt} = f \frac{\dot{Y}Z - Y\dot{Z}}{Z^2} \quad (4-4)$$

Assuming that the object being imaged does not change size and the camera only moves in the  $Z$  axis,  $\dot{X} = \dot{Y} = 0$  and hence:

$$\frac{dx}{dt} = -f \frac{X\dot{Z}}{Z^2} \quad (4-5)$$

$$\frac{dy}{dt} = -f \frac{Y\dot{Z}}{Z^2} \quad (4-6)$$

Dividing equations (4-1) and (4-2) by equations (4-5) and (4-6) respectively:

$$\frac{x}{\dot{x}} = \frac{f \frac{X}{Z}}{-f \frac{X\dot{Z}}{Z^2}} = -\frac{Z}{\dot{Z}} = \tau \quad (4-7)$$

$$\frac{y}{\dot{y}} = \frac{f \frac{Y}{Z}}{-f \frac{Y\dot{Z}}{Z^2}} = -\frac{Z}{\dot{Z}} = \tau \quad (4-8)$$

Thus, it is possible to compute time-to-contact only from the apparent size and rate of change of the size of an object in an image. This calculation can be performed with either the size of a detected object or some characteristic dimension between detected features. For accurate estimates, the dimension detection must be sufficiently robust to persist between image frames with minimal noise. Common methods for detecting the size of an object in an image are through feature detection techniques such as corner or blob detection [82].

#### 4.1.1.1 Dimension Tracking

In computer vision, feature detection is the process of automatically identifying points of interest in a digital image [82]. It is difficult to assign a definition to the fairly amorphous concept of a point of interest since interest can vary wildly between observers. Generally, it is aimed at detecting edges or corners of objects in the image in a repeatable fashion. There are a huge variety of algorithms available, but most rely on some analysis of brightness gradients in the image to evaluate where the parameters of the image vary most. A widely implemented example is the Speeded-Up-Robust-Features (SURF) detector, described in reference [116]. It uses an integer approximation of the determinant of the Hessian operator to detect points of



interest. Feature descriptors are generated by analysing the Haar wavelet response of the area around the point of interest. Feature descriptors can be used to correlate features between multiple images, such as the frames of a video, by matching the descriptors even if the feature has moved in the image between frames. Feature detection tools like SURF are widely implemented in common image processing libraries, such as OpenCV [117] and the Matlab image processing toolbox.

Blob detection is a subset of feature detection that examines regions of an image, rather than singular points. It identifies areas of a digital image that differ significantly from their surroundings in some property, usually brightness or colour. All points within a blob are similar and will have some approximately constant property across the area of the blob that is distinct from the background. This identifying characteristic is what allows blobs to be distinguished from the rest of an image. The most common method for blob detection is to convolve an image with a Gaussian kernel and then apply a Laplacian operator. The dimensions of the blob can then be obtained for further use. Reference [118] describes the development of blob detection algorithms in detail. Blob detection methods are also widely implemented in common image processing libraries.

#### 4.1.2 Helipad Landing Scene – Tau Analysis with Dimension Tracking

Figure 4-2 shows a set of results for the size method applied to a helipad landing video as described in section 3.3.3. The descent is one dimensional, with the camera descending from 50 m altitude at 5 m/s straight towards the centre of the helipad. The dimension used for Tau estimation is the width of the helipad, estimated from blob detection and shown in plot (a). Plot (b) shows the rate of change of that dimension, estimated as the difference between each successive frame divided by the frame length.

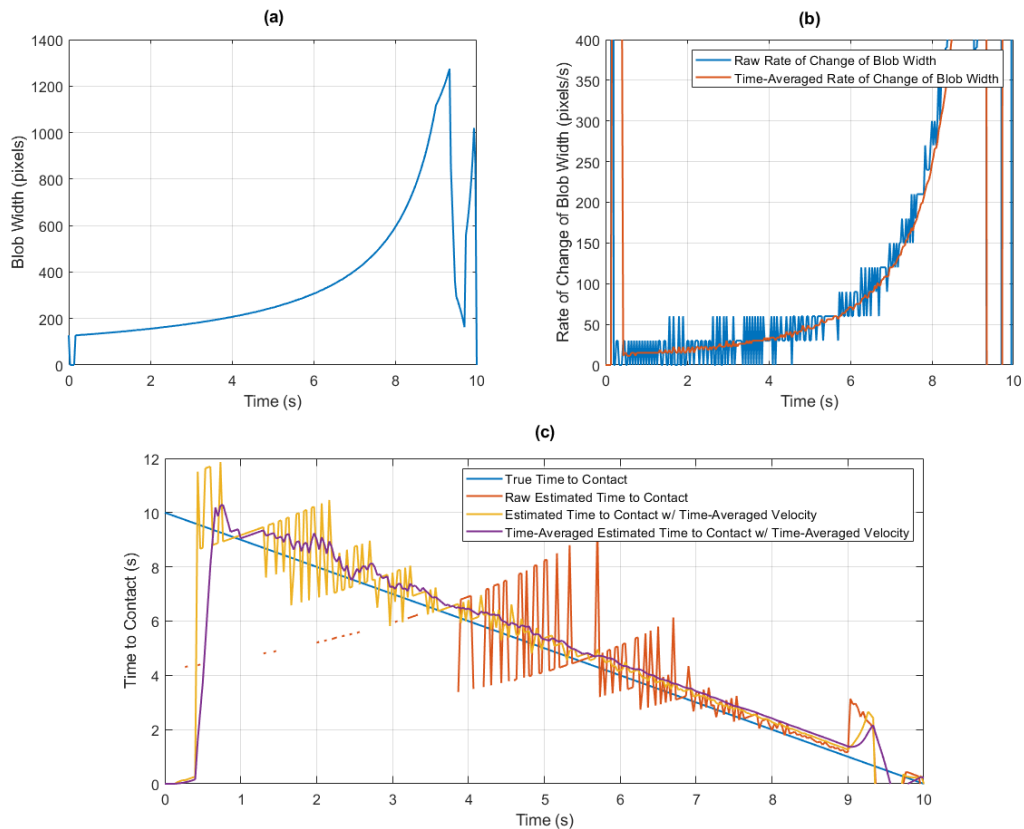


Figure 4-2: Tau estimation plots for a vertical helipad descent using the size method with blob detection: (a) the width of the detected blob (b) the rate of change of the width of the detected blob (c) the estimated time-to-contact as the ratio of (a) and (b)

This rate is noisy and is clearly quantised to a limited number of values, which is in contrast to the measured width of the helipad that seems to change smoothly. Although the size of the helipad will change smoothly as the camera approaches in reality, the accuracy of the measurement is ultimately restricted by the resolution of the camera and the measured size can only change by an integer amount of pixels. This issue is amplified in the rate and is especially problematic early in the descent where the size of the object does not change at all between some frames. The quantisation error results in no time-to-contact estimate through the first 4 seconds of the video using the raw measured values (plot (c)).

There are several ways to mitigate this problem. The simplest is to increase the resolution of the video so that there a greater range of values for the width and therefore, rate. Another solution is to vary the framerate of the video through the manoeuvre, so there is a longer period between measurements and a smaller chance

there is no pixel movement between frames. However, it is difficult to schedule the framerate without already having an accurate estimate of Tau to scale the framerate on. A third option is to apply a filter to smooth the results. In this case, a moving average filter with a window length of 7 frames was applied to the rate to smooth the signal. This filter was also applied to the width of the blob to ensure the width and rate measurements were consistent in time. However, the time-to-contact estimate is still noisy through the first half of the descent. Another identical moving average filter was also applied to the Tau estimate to smooth it further. The estimate tracks the true value well but with a constant offset of approximately 0.5 seconds. This seems to be caused by the effective time delay introduced by the moving average filters; the two nested filters cover 14 frames in total, which correlates to offset since the video frame rate is 30 frames per second.

The computational load of blob detection, filtering and Tau calculation is low enough that this process can be performed in real-time at high frame rates and camera resolutions. The computation time per frame for this example was approximately 0.03 seconds. It was implemented through a Matlab function reading frames from a pre-recorded video with a resolution of 720p. The method can produce relatively accurate results throughout the majority of the manoeuvre; however there is no estimate for the final second of the descent since the edges of the helipad move beyond the camera field of view. No blob can be detected in this phase, so no Tau estimate can be made. This issue highlights the main drawback of the size method; for it to work there must be a characteristic dimension that can be accurately tracked between frames. That dimension can be found through blob detection or feature tracking, but if either fails at any point during a manoeuvre then no accurate Tau estimation can be made.

Another drawback of this approach is that accuracy will be impacted by increased situational complexity. A common factor that adds complexity is the presence of an angle between camera and plane of interest.




Figure 4-3: Diagram of camera descending towards an inclined plane with parameters required for Tau estimation with size method labelled

For a simple case where the ground plane is inclined to the horizontal by an angle, as illustrated in Figure 4-3, the apparent size of an object on the ground plane will be different. Every point on the object is a different distance along the optical axis in this scenario, so the ‘uphill’ half of the object will appear larger on the image than the ‘downhill’ half. The size of the object can still be found by using the perspective projection equations, but is complicated to now become:

$$x = f \left[ \frac{4XZ \cos(\theta)}{4Z^2 - X^2 \sin(\theta)^2} \right] \quad (4-9)$$

Differentiating this with respect to time, and dividing  $x$  by  $\dot{x}$  yields:

$$\frac{x}{\dot{x}} = \tau \left[ \frac{4Z^2 - X^2 \sin(\theta)^2}{-4Z^2 - X^2 \sin(\theta)^2} \right] \quad (4-10)$$

It is still possible to extract Tau from the equation, but it now requires a correction factor that requires knowledge of altitude of the camera, object size and surface inclination. This level of knowledge is not feasible for an operational system on a vehicle that must operate in many different environments.

Another complication arises if the apparent size of the object in the visual scene changes and hence, the lateral rates are non-zero. This may commonly occur if the

object undergoes rotational motion. In this case, it is again no longer simple to extract time-to-contact from the equation and spatial parameters of the scenario will be required as demonstrated in equation (4-11):

$$\frac{x}{\dot{x}} = \frac{XZ}{\dot{X}Z - X\dot{Z}} \quad (4-11)$$

#### 4.1.3 Runway Landing Scene – Tau Analysis with Size Method

To investigate the size tracking method further, it was implemented on the virtual runway landing scene described in section 3.3.4 to ascertain whether the translational motion affects the accuracy of the estimate. The camera descends from a height of 100 m at a rate 5 m/s, while moving down the runway at 15 m/s. Tau is estimated using the size method through blob detection. The results are shown in Figure 4-4.

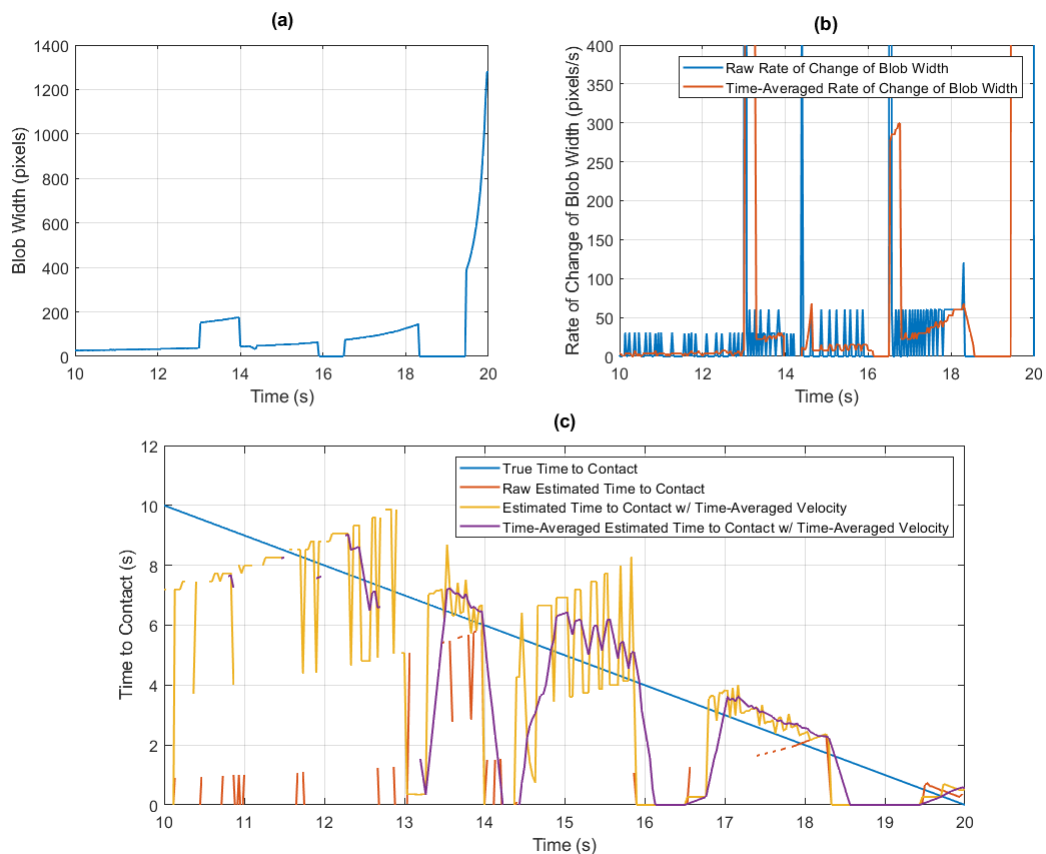


Figure 4-4: Tau estimation plots for a runway landing using the size method with blob detection: (a) the width of the detected blob (b) the rate of change of the width of the detected blob (c) the estimated time-to-contact as the ratio of (a) and (b)

Plot (a) displays the width of the largest detected blob, the width of which is used as the tracked dimension. The width of the blob changes smoothly in some sections but there are significant steps in the data as objects being tracked as blobs move out of frame, and new ones are detected. There are extended periods after the 16 and 18 second marks where no blob is being detected. As a result, there is no Tau estimate in these regions. Where blobs are detected and the rate is consistent, Tau estimates are found and they track the true value well, with a similar offset as seen in the helipad landing analysis. Evidently, the method is still accurate even with large translational motion, but the movement of features out of the camera frame makes it very susceptible to periods where no estimate is possible. Any effective Tau estimator must find a way round this problem.

## 4.2 Tau Perception through Optical Flow Divergence

Another class of time-to-contact estimation methods are based on optical flow analysis of image sequences, discussed earlier in section 2.4.2.2. The optical flow field can be regarded as the velocity field of the features in the image.

### 4.2.1 Overview

The optical flow field can be related to time-to-contact by examining the physical origin of the flow field. Referring back to the weak perspective projection rate equations in section 4.1.1 (equation (4-3) and (4-4)), it is evident that the terms equate to the optical flow components in the  $x$  and  $y$  directions. By substituting the weak perspective projection equations ((4-1) and (4-2)) back in for the  $X$  and  $Y$  terms the following is obtained:

$$u = f \left( \frac{U}{Z} - \frac{xW}{fZ} \right) \quad (4-12)$$

$$v = f \left( \frac{V}{Z} - \frac{yW}{fZ} \right) \quad (4-13)$$

These components of the optical flow field can be used to find time-to-contact by computing the divergence of that vector field across the image. The divergence,  $\nabla$ , of a two-dimensional vector field is defined as:

$$\nabla \cdot \begin{bmatrix} u \\ v \end{bmatrix} = \frac{\delta u}{\delta x} + \frac{\delta v}{\delta y} \quad (4-14)$$

Substituting in the optical flow components and evaluating the partial derivatives:

$$\nabla \cdot \begin{bmatrix} f \left( \frac{U}{Z} - \frac{xW}{fZ} \right) \\ f \left( \frac{V}{Z} - \frac{yW}{fZ} \right) \end{bmatrix} = -2 \frac{W}{Z} \quad (4-15)$$

As a result, time-to-contact can be found by evaluating the partial derivatives of the flow field and applying:

$$\tau = \frac{2}{\frac{\delta u}{\delta x} + \frac{\delta v}{\delta y}} \quad (4-16)$$

If the flow field is known perfectly, then time-to-contact can also be perfectly recovered. Robust and accurate methods of optical flow field estimation are required to facilitate this.

#### 4.2.2 Helipad Landing Scene – Tau Analysis with Optical Flow Divergence Method

To illustrate the results possible using divergence analysis, the ‘perfect’ optical flow field can be calculated using equations (4-12) and (4-13) for the helipad landing scene. The camera state is known at all frames of the video, so the true flow at every pixel position can be found easily. The resulting optical flow vectors are displayed in Figure 4-5, superimposed on the corresponding frames from the helipad landing video. The flow vectors were calculated at every pixel position but are only plotted at a subset of them for the sake of the clarity.

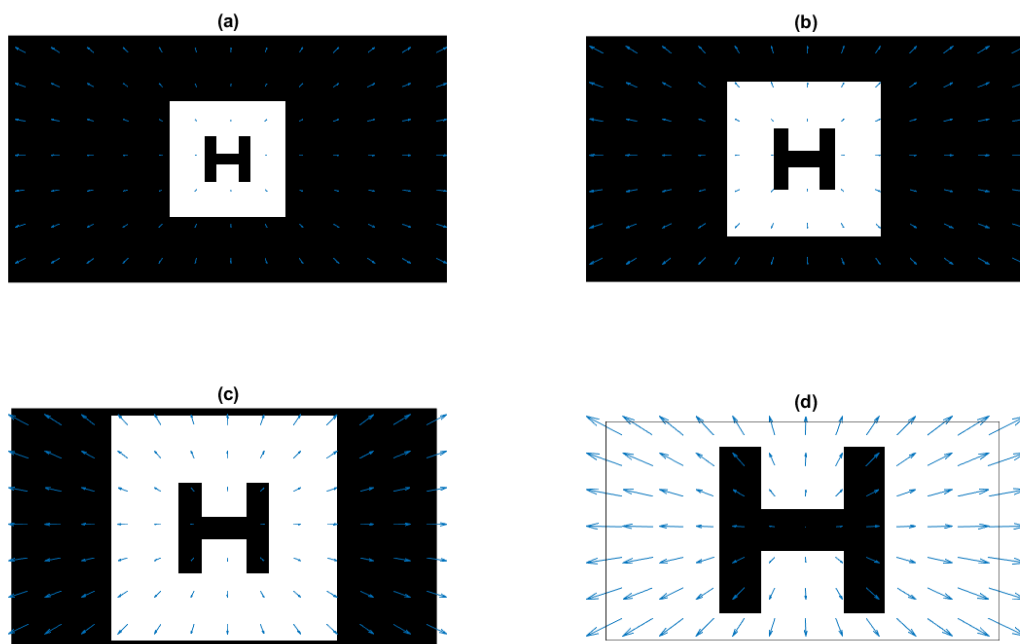



Figure 4-5: Helipad landing scene with perfect optical flow field vectors superimposed calculated on a coarse grid. The helipad lies on a flat plane and in each image, the camera is a different height above the ground: (a) 20m (b) 15m (c) 10m (d) 5m.

By calculating the inverse of the divergence of this flow field at each frame, the time-to-contact of the camera with the helipad can be perfectly reconstructed so that the measured value is exactly equal to the true value. However, this is only possible if the recovered optical flow field itself is also perfect. Hence, an accurate Tau estimate



relies on an accurate optical flow field estimate. Figure 4-6 shows a measured flow field for the helipad scene estimated with a Lucas-Kanade method implemented using the Matlab computer vision toolbox.




*Figure 4-6: Helipad landing scene with a superimposed optical flow field recovered using Lucas-Kanade estimator. Flow is calculated at detected feature locations.*

Evidently, the optical flow vectors are very different from the perfect flow fields shown in Figure 4-5. Mostly, the flow vectors only have a single component normal to the edges of the helipad boundaries. Since these flow vectors are parallel, there is no divergence across these regions. At the corners of the shapes there are vectors with components in both the horizontal and vertical directions that match the expected pattern, but without neighbouring flow vectors that are also correct, there can be no accurate estimate of the field divergence. This is not necessarily due to a fault in the optical flow estimator but may be an artefact of the simplicity of this visual scene; it is very difficult to track the motion of a pixel along an edge when the neighbouring pixels are exactly the same colour. A more complex visual scene was also used for flow field estimation to assess whether scene complexity was affecting the accuracy of the results.

#### *4.2.2.1 Frigate Deck Landing Scene – Optical Flow Divergence Analysis*

To better assess optical flow field estimation techniques, the frigate deck landing scene was used. This virtual environment includes more elements in the scene, orientated in more different directions. There is also richer background texture and greater colour variance. These elements convey more information for linking points

between frames, as required in optical flow estimation. Figure 4-7 shows a frame from a frigate deck landing video with measured optical flow vectors superimposed. In this particular scenario, the camera descends toward the deck at 2 m/s from an initial height of 20 m above the deck.



*Figure 4-7: Matlab VR helipad landing scene with superimposed optical flow field from Lucas-Kanade flow estimator*

The flow vectors are distributed over a wider range of directions but are clearly still concentrated parallel to the two main axes. This is especially clear on the two horizontal lines in the top corner of the image. However, the ring around the central grid gives a more even distribution of flow directions. Calculating time-to-contact using the mean divergence of this field does provide estimates of Tau close to the true value through some of the manoeuvre, as shown in Figure 4-8.

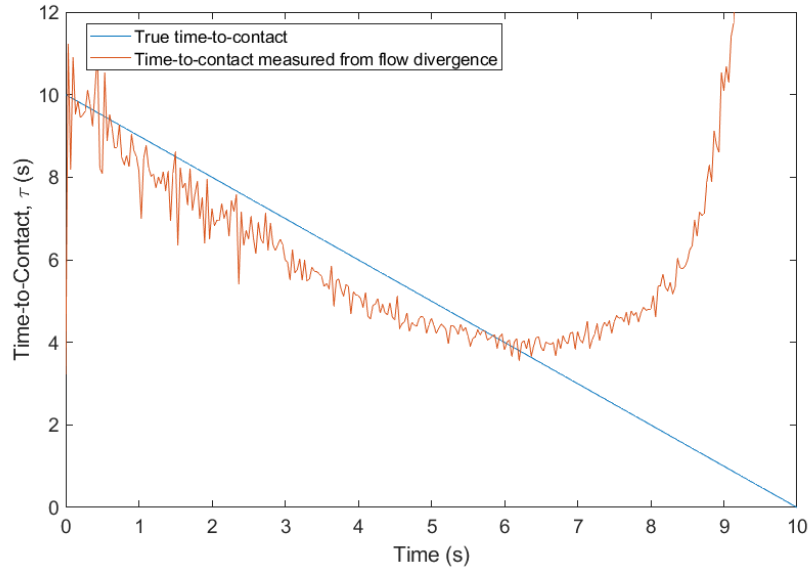



Figure 4-8: Time-to-contact from optical flow divergence for a VR frigate deck landing from a 10m altitude at a constant velocity. Optical flow field estimated using the Lucas-Kanade method. Tau calculated from mean divergence of that field.

Through the first 5 or 6 seconds of this manoeuvre, the estimate of time-to-contact does follow a linear trend of equal gradient to the true value, albeit with a slight offset and some noise. After 6 s the Tau estimate diverges from the true value and towards infinity. The time of this divergence correlates with the time that the ring feature on the deck begins to move out of the camera field of view. The ring seems to be an important feature for accurate flow field estimates, so it is possible that the dominance of perpendicular features at certain scales contributes to the large errors. Another possible contributor is the true amount of optical flow as the camera approaches the deck. Figure 4-9 shows the mean optical flow magnitude in the  $x$  direction during this manoeuvre. The amount of flow will be greater near the edges of the image and smaller near the focus of expansion. Since motion, in this case, is one dimensional, the focus of expansion will be at the centre of the image and the mean flow value corresponds to a point a quarter of the way across the image.




*Figure 4-9: Mean optical flow magnitude in x-direction through a frigate deck landing manoeuvre from a 10m altitude at a constant descent velocity*

The mean flow begins at about 2 pixels per frame and increases steadily to around 5 pixels per frame by the 5-second mark. From this point, the flow increases exponentially. This result is unsurprising when examining the perspective projection equations; the position of any point in the image varies with the inverse of distance along the optical axis, so flow becomes singular as the distance approaches zero. While this is not a problem in the divergence phase of the time-to-contact estimation, it will cause a problem with optical flow field estimation. All optical flow estimations rely on the constant brightness assumption, which states that the brightness information contained in an image does not change with time. In practice this can never be true when motion occurs along the optical axis, but as long as the change is minimal the assumption is serviceable. The exponential increase in optical flow seems to violate this assumption as the manoeuvre progresses, and optical flow field estimates suffer as a result.

The flow divergence method can also be computationally expensive at high video resolutions. The 720p resolution video analysed above took an average of 0.12 s to process each frame, nearly four times the real time period of the frame since it was recorded at 30 frames per second.

### 4.2.3 Runway Landing Scene – Tau Analysis with Optical Flow Divergence Method

The flow divergence problem is compounded in fixed-wing cases, where large translational motion is present due to the high forward speed of the aircraft. The high base level of optical flow caused by this translational motion between scene and camera reduces the effectiveness of optical flow estimators further. The unreliability of feature detection observed in the dimension tracking method is also a problem when using a Lucas-Kanade flow estimator, which relies on consistent feature detection. Figure 4-10 shows a Tau estimate from the optical flow divergence method for the VR runway landing scene described in section 3.4.2.



*Figure 4-10: Tau from optical flow divergence for a fixed-wing runway landing of the 3DR Aero from 100m altitude, descending at 5m/s while travelling down the runway at 30 knots. The optical flow based Tau estimates is noisy and converges to discrete values of Tau that do not track the true value well/*

The estimate does not track the true value well at any point in the manoeuvre and suffers from the same divergence towards the end observed in the deck landing analysis. There is a notable step up in the estimate between 20% and 30% through the manoeuvre, which correlates with the time that the end of the runway (and end

of the threshold markings) disappear from view of the camera. This change in feature content seems to heavily impact the estimate as in the dimension tracking method.

### 4.3 Tau Perception through Direct Gradient Method

The third class of methods that will be examined here, direct gradient methods, borrow elements from both optical flow analysis and dimension tracking. They use the weak perspective projection equations to divine a shortcut to calculating time-to-contact from the constant brightness equation of optical flow.

#### 4.3.1 Overview

The direct gradient method uses the expressions for the optical flow components derived from the weak perspective projection equations, and substitutes them into the constant brightness equation directly to yield:

$$I_x f \left( \frac{U}{Z} - \frac{xW}{fZ} \right) + I_y f \left( \frac{V}{Z} - \frac{yW}{fZ} \right) + I_t = 0 \quad (4-17)$$

This equation is the fundamental equation of the base direct gradient method. The equation can then be formulated as least squares minimisation to yield estimates of time-to-contact. This approach can be simplified and extended in various ways to cover a range of different relative motions between camera and visual scene. The derivations of these methods are given in Appendix B.2. These equations allow for time-to-contact to be estimated for a camera moving in six degrees of freedom relative to a target inclined to the camera in two dimensions, which covers most possible scenarios. This 6 DoF with slope and rotational corrections method will be used throughout the analysis in this section unless otherwise stated. This particular extension results in two linear matrix equations that can be solved iteratively to produce estimates of Tau, which is the inverse of  $C$  in the below equations. The system of equations in the  $Ax = B$  format used is:

$$\begin{bmatrix} \Sigma F^2 I_x^2 & \Sigma F^2 I_x I_y & \Sigma F^2 G I_x \\ \Sigma F^2 I_x I_y & \Sigma F^2 I_y^2 & \Sigma F^2 G I_y \\ \Sigma F^2 G I_x & \Sigma F^2 G I_y & \Sigma F^2 G \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix} = - \begin{bmatrix} \Sigma F^2 I_x I_\omega \\ \Sigma F^2 I_y I_\omega \\ \Sigma F^2 G I_\omega \end{bmatrix} \quad (4-18)$$

$$\begin{bmatrix} \Sigma D^2 x^2 & \Sigma D^2 xy & \Sigma D^2 x \\ \Sigma D^2 xy & \Sigma D^2 y^2 & \Sigma D^2 y \\ \Sigma D^2 x & \Sigma D^2 y & \Sigma D^2 \end{bmatrix} \begin{bmatrix} P_s \\ Q_s \\ C \end{bmatrix} = - \begin{bmatrix} \Sigma D x I_\omega \\ \Sigma D y I_\omega \\ \Sigma D I_\omega \end{bmatrix} \quad (4-19)$$

$A$  and  $B$  are functions of the horizontal velocities  $U$  and  $V$ , altitude of the camera,  $Z_0$  and camera focal length,  $f$ .  $C$  is the inverse of time-to-contact, the ratio of vertical velocity,  $W$ , and altitude.  $P_s$  and  $Q_s$  are functions of the slope of the target plane,  $p_s$  and  $q_s$ ,  $C$ , and focal length.  $J$  and  $K$  are functions of the rotational velocities of the camera,  $\omega_x$ ,  $\omega_y$  and  $\omega_z$ , the camera focal length, and the camera pixel positions,  $x$  and  $y$ . These seven parameters are defined as:

$$A = f \frac{U}{Z_0} \quad (4-20)$$

$$B = f \frac{V}{Z_0} \quad (4-21)$$

$$C = -\frac{W}{Z_0} = \frac{1}{\tau} \quad (4-22)$$

$$P_s = \frac{p_s W}{f Z_0} \quad (4-23)$$

$$Q_s = \frac{q_s W}{f Z_0} \quad (4-24)$$

$$J = \frac{xy}{f} \omega_x - \frac{x^2 + f^2}{f} \omega_y + y \omega_z \quad (4-25)$$

$$K = \frac{y^2 + f^2}{f} \omega_x - \frac{xy}{f} \omega_y - x \omega_z \quad (4-26)$$

These parameters are grouped in four further parameters for simplicity. The radial gradient,  $G$  is a combination of the spatial image gradients and pixel positions.  $I_\omega$  combines the two rotational motion parameters and the image temporal gradient.  $D$  and  $F$  group the unknowns in the equation into a translational motion group, and a surface slope group respectively. These four parameters are defined as:

$$G = I_x x + I_y y \quad (4-27)$$

$$I_\omega = I_x J + I_y K + I_t \quad (4-28)$$

$$D = G + I_x \frac{A}{C} + I_y \frac{B}{C} \quad (4-29)$$



$$F = 1 + \frac{P_s}{C} + \frac{Q_s}{C} \quad (4-30)$$

It is strongly recommended that the reader examines the derivations of the method given in Appendix B.2 to fully understand how it works, because it is difficult to fully conceptualise the process without understanding the underlying mathematics. As a high-level overview: the method uses ratios of spatial gradients in a video frame to temporal gradients between consecutive video frames to estimate time-to-contact with an approaching surface. This is analogous to using the ratio of distance (spatial variable) and velocity (temporal variable) to calculate Tau, as discussed in section 2.2.2. Attempting to avoid the mathematics, a more detailed account of the process can be given as follows:

1. Capture a video frame
  - 1.1. Store the video frame in a frame buffer
2. Calculate the image gradients in the X & Y directions of the captured video frame,
  - 2.1. Store the X & Y gradients in a 2nd buffer
3. Calculate the image gradient in the time direction (between frames) using the current frame and previous frame in the frame buffer
  - 3.1. Store the temporal gradient in a 3rd buffer
4. Calculate the radial gradient of the frame using the X & Y gradients and the pixel positions in the image frame of reference (equation (4-28))
  - 4.1. Store the radial gradient in a 4th buffer
5. Average the spatial, temporal and radial gradients over the length of the 4 respective buffers
6. Assemble parameter groups (equations (4-27) to (4-30))
7. Calculate the expected movement between the current video frame and previous frame caused by rotational motion of the camera; rotational rates from a gyroscope are combined with the pixel positions in the image frame of reference (equations (4-25) and (4-26))

8. Combine the expected movement from rotational motion of the camera with the averaged image temporal gradient to find a corrected temporal gradient
9. Estimate translational motion of the camera from ratios of the spatial and radial gradients to the corrected temporal gradient, as well as an initial guess of relative slope between camera and target (equation (4-18))
10. Use the estimated translational motion of the camera from step 8 and the ratios of spatial, radial and temporal gradients to estimate relative slope between camera and target (equation (4-19))
11. Iterate through steps 9 and 10 multiple times using updated estimates of translational motion and relative slope to sequentially improve estimate accuracy of each quantity
12. Calculate time-to-contact from the estimated translational motion parameters (equation (4-22))

A flow diagram of this process with some example image and gradient frames is provided in Appendix B.3.1 for further clarity.


The method was implemented using a custom class of objects defined by the author in Matlab that can be used to easily store and operate on image matrices in an efficient method using a frame buffer. The class includes methods to leverage multi-threaded processors to run tasks in parallel, as well as offload computations to a Graphics Processing Unit (GPU) to speed up calculations.

The direct gradient method faces some of the same issues as those described for the size tracking and optical flow methods. Primarily, it is difficult to determine accurate temporal gradients due to the quantization error in the image pixel structure. To mitigate these issues, the frame buffer object stores the image gradients for a user-defined number of frames so time averaging can be applied over all of them.

As for the other Tau estimation approaches, this implementation of the direct gradient method was tested for scenarios more applicable to rotary-wing platforms, and others to fixed-wing platforms.

### 4.3.2 Rotary-Wing Platform Analysis

The direct gradient method for Tau estimation was used to analyse the various visual scenes described in section 3.3.3 and 3.4.1, based around helipad and frigate deck landings. Figure 4-11 shows the results of the Matlab implementation of the gradient method applied to a helipad landing video. This example was computed using a 10-frame long buffer, so the Tau estimate should lag the true value by five frames ( $\sim 0.16$  s).



*Figure 4-11: Time-to-contact measured using the direct gradient method from the Helipad landing scene, and true Tau value during the manoeuvre. Camera descends from 50 m at 5 m/s and uses a 720p resolution*

The uncertainty in the temporal gradients early in the manoeuvre cause large amounts of noise in the Tau estimate but this effect begins to dissipate in the middle of the manoeuvre, between 2 and 7 seconds. The estimate diverges completely after 7 seconds, in a similar manner to the optical flow divergence Tau estimate method in section 4.2.2.1. To verify that the extensions to the gradient method described in section 4.3.1 function as expected, a Tau estimate for a helipad landing video with sinusoidal motion in multiple degrees of freedom is shown in Figure 4-12.

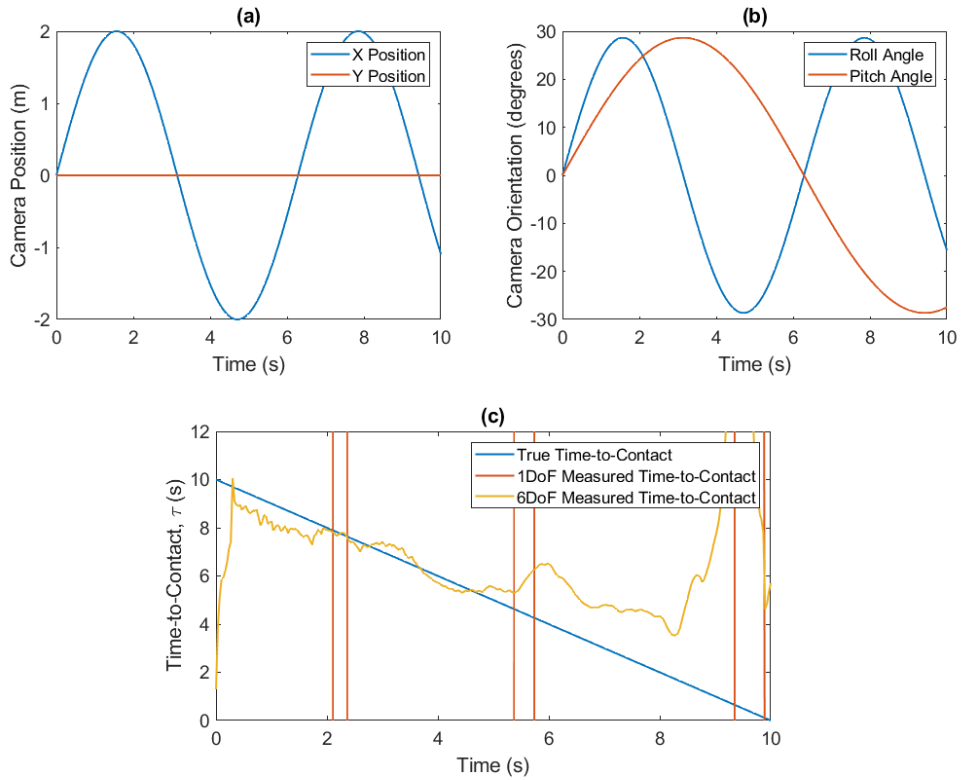


Figure 4-12: Time-to-contact measured using the direct gradient method from the Helipad landing scene. Camera descends from 50 m at 5 m/s with sinusoidal motion in lateral position, pitch and roll angle: (a) camera position in X & Y axes, (b) camera orientation in roll and pitch axes, (c) Tau estimates with 1DoF and 6DoF w/ rotational corrections methods

Interestingly, the amount of noise in the Tau estimate is lower in this 6 degree of freedom case than in the 1 degree of freedom case above. A good Tau estimate is obtained for the first 5 seconds of the manoeuvre before the value starts to diverge. The one degree of freedom method applied to the same video is also shown in plot (c) in red. In comparison to the 6 degree of freedom result also plotted in yellow, this method clearly fails to provide any useful estimate and as such proves that the extensions to the method are valid and useful. However, the 6 DoF estimate begins to diverge at 5 seconds, slightly earlier than the simple 1 DoF approach in Figure 4-11.

#### 4.3.2.1 Motion Parameter Estimates

The direct gradient method also yields estimates of other quantities in addition to time-to-contact.  $A$  and  $B$  contain information about the translational velocities of the camera, and  $P_s$  and  $Q_s$  relate to the relative slope between the camera and surface it is approaching. It is not clear how useful the information from these other products

is, so this section will examine these variables in the context of the simple helipad landing. By rearranging equations (4-20) and (4-21), expressions for the translational velocities of the camera can be found:

$$U = \frac{W A}{f C} \quad (4-31)$$

$$V = \frac{W B}{f C} \quad (4-32)$$

These expressions are based on the ratio of outputs of the gradient method, but also require knowledge of the camera focal length, and vertical velocity of the camera,  $W$ . This velocity parameter could be estimated from an INS, but would not be reactive to environmental dynamics. As a result, it doesn't provide much utility over the translational velocity estimates from an INS. Figure 4-13 shows a velocity estimate from the direct gradient method. This result was obtained from analysis of a helipad landing video where the camera descended from 50 m altitude at 5 m/s, and moved sinusoidally in the X axis at a maximum speed of 2 m/s. This result assumes perfect knowledge of the vertical velocity  $W$ .




Figure 4-13: Camera velocity estimate in X direction from direct gradient method, plotted with true velocity. Video depicts vertical descent with sinusoidal motion in the horizontal X axis.


The velocity estimate tracks the true value well with a small amount of noise, similar to that seen in the Tau estimate. The estimate lags slightly due to the delay introduced by the five frame delay.

The relative slope between the camera and surface being approaching can also be estimated using the  $P_s$  and  $Q_s$  parameters. Rearranging equations (4-23) and (4-24) and taking the inverse tangent allows recovery of these relative angles:

$$p_s = \arctan\left(f \frac{P_s}{C}\right) \quad (4-33)$$


$$q_s = \arctan\left(f \frac{Q_s}{C}\right) \quad (4-34)$$

These quantities only require outputs of the gradient method and the camera focal length, so no spatial parameters are needed as with the translational velocities. Figure 4-14 shows the slope estimate in the X direction for another helipad landing video. The camera descends along the Z axis, but the surface is inclined at 20 degrees to the horizontal plane about the Y axis.



*Figure 4-14: Estimate of relative slope between camera and approaching surface from direct gradient method. Video depicts vertical descent towards surface inclined 20 degrees from the horizontal plane.*

The estimate is highly noisy and varies between -90 and 90 degrees rapidly. Taking the mean of the estimate over the full manoeuvre yields a result of -19.5 degrees, close to the magnitude of the true value, 20 degrees, which is constant throughout. A further case where the camera descends vertically while rotating about the Y axis was examined to examine the slope estimate further, the results of which are shown in Figure 4-15.



*Figure 4-15: Estimate of translational velocity and relative slope between camera and approaching surface from direct gradient method. Video depicts vertical descent towards surface while sinusoidally rotating about the Y axis.*

Despite there being no translational motion of the camera, the direct gradient method erroneously measures velocity in the X axis. This is due to the apparent motion of the surface features caused by the rotational motion. It may be possible

to correct for this using rotational measurements, but this would also require knowledge of the distance between surface and camera along the optical axis. A velocity estimate along the optical axis was already required to generate this signal, so if a reliable distance estimate was also available then Tau could be computed directly with no need for the gradient method in the first place. The slope estimate is again very noisy, and varies between -90 and 90 again, since relative angle is varying in this scenario, taking a mean value is not useful to improve the estimate.

All four of these parameters differ from time-to-contact in one key respect; they are used to recover spatial parameters from the scene, whereas Tau is purely temporal. The difficulty of recovering spatial parameters from monocular videos was discussed in section 2.4.2, and the same concepts apply here; the single viewpoint makes scale and distance ambiguous. Since Tau is defined as a ratio that eliminates spatial variables, the estimates from the direct gradient method are more reliable. However, the velocity and angle information produced is clearly unreliable, as evidenced by the above figures. As a result, it is not recommended to use any of the extra motion parameters from the direct gradient method as control system variables.

#### 4.3.3 Factors Affecting Tau Estimate Accuracy

It is clearly possible to gain accurate Tau estimates in certain conditions, but it is not evident what those conditions are, or whether they can be manipulated for improved accuracy in all situations. If they are extrinsic parameters of the camera then it implies operational restrictions on the user of a sensor using these methods. Intrinsic parameters may be easier to control. The next sections will describe a number of experiments performed to further investigate how estimates can be improved.

##### *4.3.3.1 Approach Speed Variation*

To investigate further, the Tau measurements and optical flow fields of a number of helipad landing videos were analysed. Four different helipad landing videos were generated with varying descent speeds of 2.5 m/s, 5 m/s, 7.5 m/s and 10 m/s. All four began from a height of 50 m above the helipad. Figure 4-16 (a) shows the measured Tau values for each video using the 6 DoF gradient method.



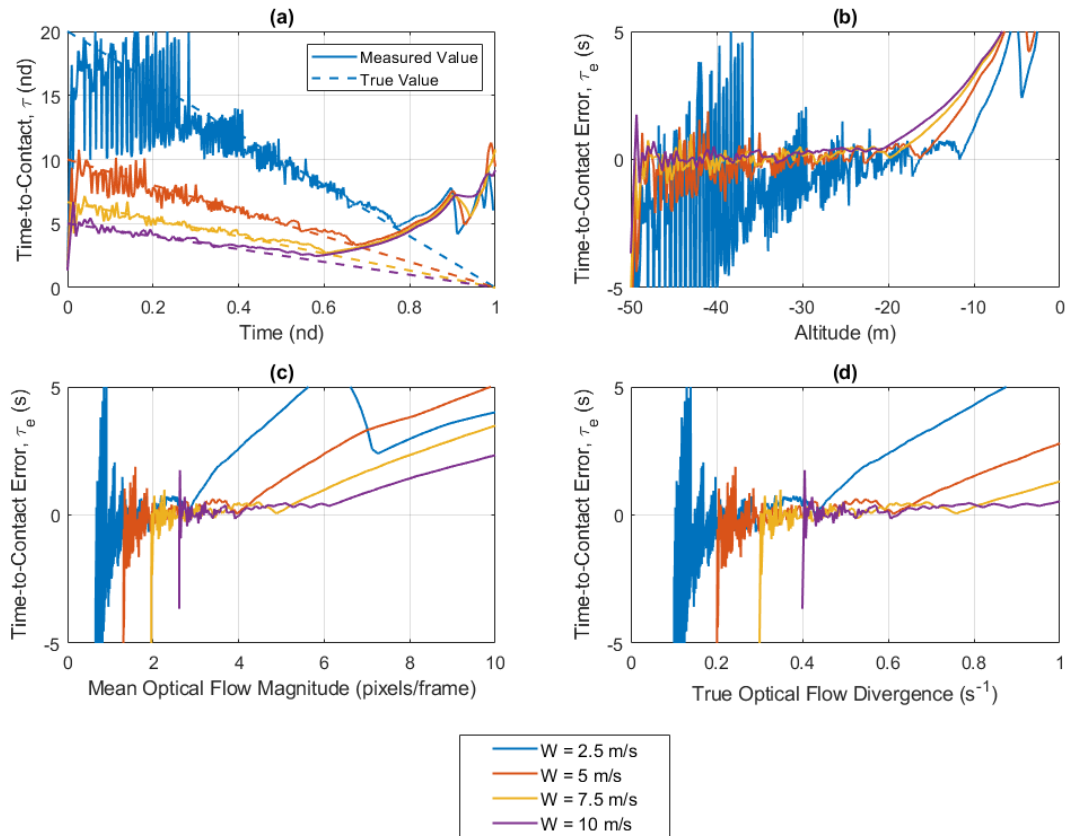


Figure 4-16: Tau estimation analysis for varying descent speeds in the helipad landing scene: (a) measured and true Tau values for the 4 different descent speeds, (b) Tau errors against altitude, (c) Tau error against mean optical flow magnitude, (d) Tau error against optical flow divergence

At higher descent rates, the measurement is less noisy during the first phase of the manoeuvre but diverges away from the true value earlier. At lower descent speeds, the measured value is not accurate initially but has an effective phase further into the manoeuvre and diverges later. To shed light on what may be driving this behaviour, plot (b) shows the error in the Tau estimate against camera height above the helipad. The divergence in measured Tau occurs at a different height for each approach speed, so the divergence does not appear to be related to altitude, or hence what is visible in the scene. Likewise, plotting Tau error against mean magnitude of the optical flow field (c) or optical flow field divergence (d), also does not exhibit an obvious link between either of these quantities and the rapid decrease in measurement accuracy.

The two main assumptions made for the direct gradient method are (1) the camera aperture behaves as a pinhole, and (2), image brightness is constant between frames.

The first assumption is necessarily true in this case since the visual scene was also generated using a pinhole camera model. The second, the constant brightness assumption referenced in section 4.2.1, is more difficult to assess. The constant brightness equation is generally regarded as a subject for minimisation rather than a strict equality, but it is not clear what level of brightness change constitutes a violation of the assumption. To assess whether this is indeed a factor, it's useful to examine the residual of the constant brightness equation at each frame, as well as the above quantities. The mean of the constant brightness equation residuals for the same scenario as above are shown in Figure 4-17 (b). Comparing against the time-to-contact error in (a), it is clear that they both diverge at the same point in the manoeuvre for each speed. Hence, the deviation from the constant brightness assumption is likely the driving factor in the large accuracy decrease.

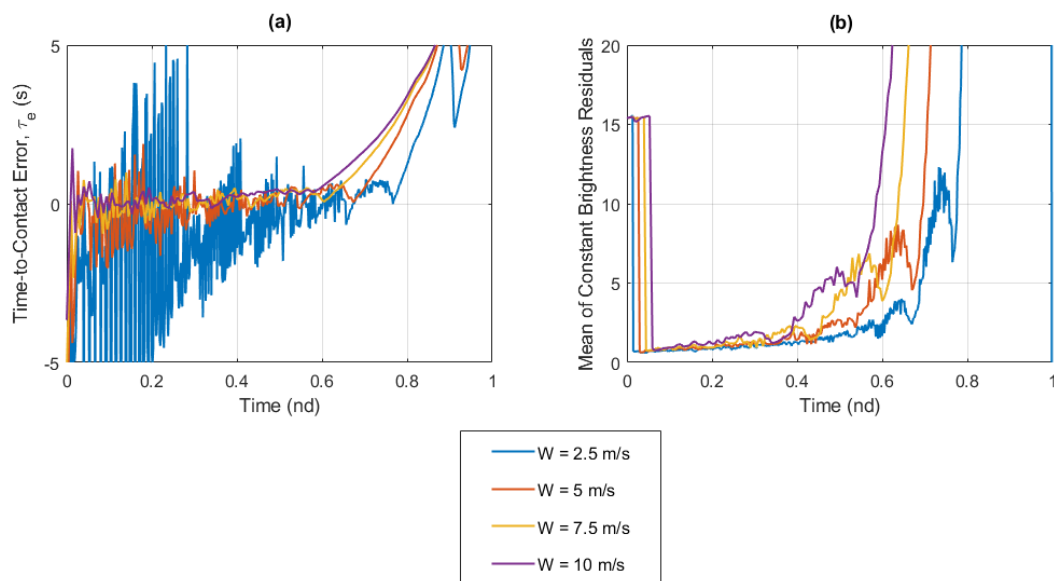


Figure 4-17: Tau estimation analysis for varying descent speeds in the helipad landing scene: (a) Tau measurement error, (b) mean of constant brightness residuals across each frame

Interestingly, the constant brightness violation does not appear to be a factor in the early Tau measurement noise, since the residuals are all low at this point. Instead, this is still likely to be due to the pixel quantization error. This is reinforced by the fact that this noise is amplified in the slower descents, where the flow of pixels per frame is consistently lower than one for the first 40% of the manoeuvre, as shown in Figure 4-18. This corresponds to the bulk of the noise in Figure 4-16 (a).




Figure 4-18: Mean optical flow magnitude for varying descent speeds in the helipad landing scene with camera descending from 50m in each case

Plotting the error in Tau measurements against the constant brightness equation residuals yields an approximately linear relationship between the two while the mean residuals are below a value of 50, regardless of the speed of descent. This is shown in Figure 4-19. The noise in the lower left corner of the plot is again likely due to the pixel quantization error, so has been neglected.




Figure 4-19: Error in Tau measurement using direct gradient method against the residuals of the constant brightness equation for varying descent speeds in the helipad landing scene. Note, the data plotted is not time progressing in any way and only expresses the relationship between Tau and the residual of the constant brightness equation

As a result, it is possible to predict how accurate the Tau measurement is by finding the residual of the constant brightness equation. However, part of the reason for using this gradient method is not having to recover the optical flow field at all. A simple estimation of the flow field based on spatial parameters may suffice to indicate the effective Tau measurement range for a particular manoeuvre and camera. Nevertheless, it is rarely practical to fit the dynamics of a vehicle manoeuvre around the effective operating constraints of a sensor, so the obvious question is: what parameters can be manipulated to keep the residuals of the constant brightness equation low, regardless of the camera dynamics?

There are five variables in the constant brightness equation that contribute to this residual: the three image gradients and the two optical flow components. The spatial image gradients are solely properties of the image and cannot be altered. The temporal gradient is again heavily dependent on the image but can be affected slightly by changing the frame rate of the camera. The values of the optical flow components are also dependent on the camera dynamics and visual scene. However, it has already been illustrated that the quantization of the image over a pixel grid can affect Tau measurements. The optical flow in the X and Y directions for 6 degree of freedom motion in the units of pixels per frame can be defined as:

$$u_{px} = \frac{T_f}{x_{px}} \left[ f \left( -\frac{U}{Z} + \frac{xW}{fZ} \right) + J \right] \quad (4-35)$$

$$v_{px} = \frac{T_f}{y_{px}} \left[ f \left( -\frac{V}{Z} + \frac{yW}{fZ} \right) + K \right] \quad (4-36)$$

All parameters in the square brackets (with the exception of the focal length) are extrinsic parameters of the camera so require movement of the camera to alter. That leaves the dimensions of each pixel and the period of time covered by each video frame as possible avenues of exploration.

#### 4.3.3.2 Frame Rate Variation

Four videos of varying frame rate were generated depicting the same helipad landing scene. The manoeuvre is again performed from a starting height of 50 m, with a descent speed of 5 m/s for all four videos. The only difference between the four

sequences is the number of frames per second rendered, with a separate from 10, 20, 30 and 40 frames per second produced. The Tau measurement Root-Mean-Square (RMS) error for each video is shown in Figure 4-20. The RMS error was computed with a window length of 15 frames.




Figure 4-20: Tau measurement error using the direct gradient method for a constant velocity descent in the helipad landing scene with varying camera framerates: (a) 10 fps (b) 20 fps (c) 30 fps (d) 40 fps

Similar to the results of approach speed variation experiment, the differences lie in the amount of early phase noise, and late phase divergence. The lower frame rate videos exhibit less measurement noise because the increased time between sampling allows for greater change in the image. This is important when movement of image features is small (near a speed of one pixel per frame) and allows for better estimates of the image temporal gradients. However, optical flow increases later in the manoeuvre are amplified by the slow sample rate and lead to violations of the

constant brightness assumption earlier on. The opposite is true for high frame rates; noise is high early in the manoeuvre, but accuracy diverges more slowly. It is also important to note that high frame rates necessarily require more computing resources to perform the required Tau estimation processes in real time. A similar experiment can be performed to assess image resolution changes.

#### 4.3.3.3 Image Resolution Variation

To change the dimensions of each pixel is to either change the field of view the camera, or the resolution of the camera. It is generally undesirable to alter the field of view dramatically, but almost all cameras offer simple ways to change the resolution. To assess the effect of varying camera resolution on Tau measurements, four different helipad landing scenes were generated with different camera properties. The resolutions and properties of each are summarised in Table 4-1. All four are resolutions in widespread use, covering several different aspect ratios. Field of view is approximately maintained by varying pixel size. The focal length is the same for all four resolutions but could be adjusted slightly to create a better field of view match if desired.

Table 4-1: Summary of video resolutions

	<b>Resolution</b>	<b>Number of Pixels</b>	<b>Aspect Ratio</b>	<b>Pixel Size (<math>\mu\text{m}</math>)</b>	<b>Field of View (<math>^{\circ}</math>)</b>
1080p	1920 x 1080	2,073,600	16:9	[1.12, 1.12]	[38.9, 22.5]
720p	1280 x 720	921,600	16:9	[2.24, 2.24]	[50.5, 29.7]
VGA	640 x 480	307,200	4:3	[3.36, 3.36]	[38.9, 29.7]
HVGA	480 x 320	153,600	3:2	[4.48, 4.48]	[38.9, 26.5]

Again, all videos depict a camera descending from 50 m at 5 m/s. The only change is the resolution of each video. The Tau measurement moving RMS error is displayed in Figure 4-21. The RMS error was computed with a window length of 15 frames.

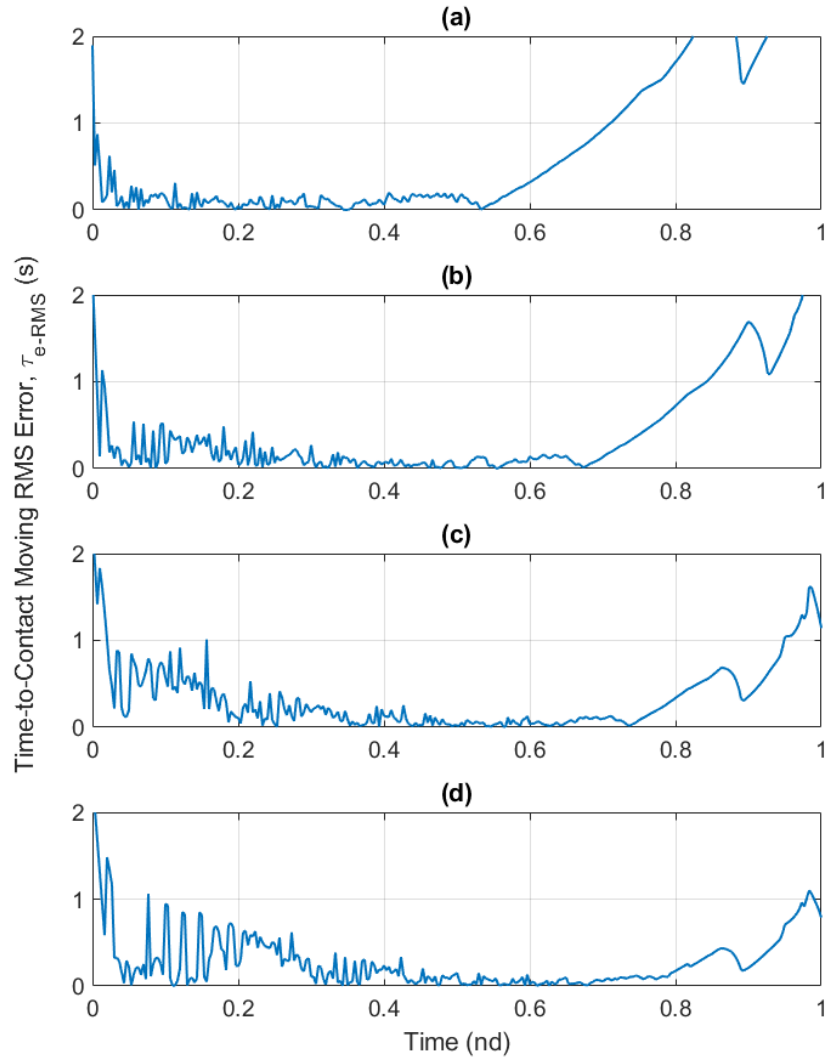


Figure 4-21: Tau measurement error using the direct gradient method for a constant velocity descent in the helipad landing scene with varying camera resolutions: (a) 1080p (b) 720p (c) VGA (d) HVGA

The higher HD resolutions exhibit less measurement noise early in the manoeuvre but experience a divergence in accuracy towards in the latter half. Conversely, low resolutions have high noise early but maintain accuracy for longer. These results mirror those of the approach speed variation and frame rate experiments closely. Higher resolution images also dramatically increase the computing resources required to perform the Tau calculations in real time. Figure 4-22 shows the mean computation for one time step at each of the resolutions above. The computation time is broken down into the three main phase of the process: acquiring the frame from the video source and computing the image gradients, averaging the gradients across the frame buffer, and performing the Tau computation.

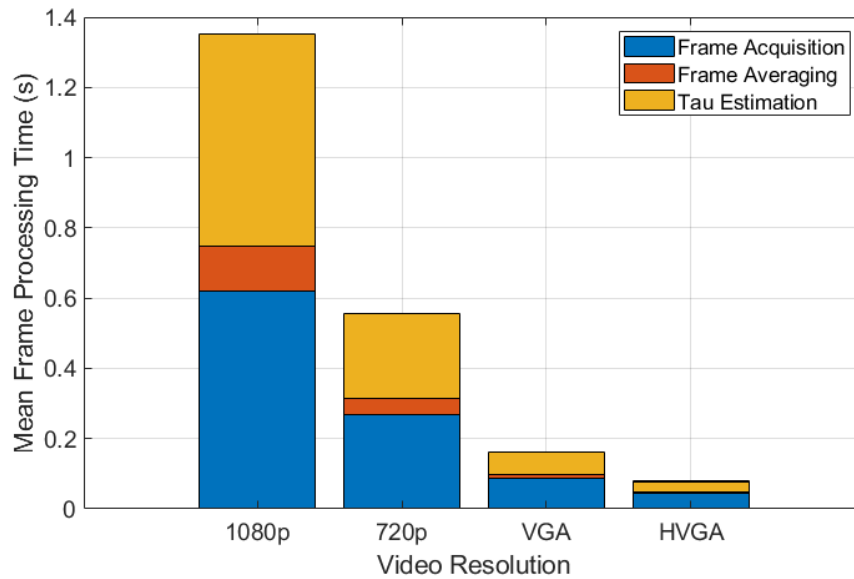


Figure 4-22: Mean frame processing time for varying resolution in direct gradient method Tau estimation, broken down by the three main operations performed during the calculation

For this implementation, computation time at 1080p resolution is over 1.3 seconds per frame, clearly not practical for real-time use. Frame processing time reduces in proportion to the number of elements in the image matrix. At HVGA resolution the processing time was approximately 0.065 seconds, which would allow for a frame rate of 15 frames per second. However, low resolution and low frame rate is likely to cause very large amounts of noise early when time-to-contact is still high. The Matlab programming language is useful for prototyping and analysis but not well optimised for high execution speed, so faster implementations will be possible.

Several conclusions can be drawn from these three investigations:

1. Tau estimates from the direct gradient method will be noisy where optical flow components are consistently below 1 pixel per frame
2. Tau estimates from the direct gradient method become unusable once the residual of the constant brightness equation become non-trivial. Initial indications suggest that Tau measurement error will exceed 0.5 s if the mean value of the residual across the image exceeds 10.
3. Varying camera frame rate and pixel sizing allows some control over these issues. Manipulating the two parameters throughout a manoeuvre should allow for increased accuracy in Tau measurements.



Though the effects of varying pixel size and frame rate are now clearer, in reality, both are relatively difficult to alter in real time during a live video stream. Altering frame rate would require very careful scheduling to allow all necessary operations to be performed in a timely and efficient manner. Changing the camera resolution may cause the field of view of the camera to change unexpectedly and introduce new sources of error. It may be possible to record multiple video streams at once, but this would multiply both required memory use and processing requirement for each video stream, as well increasing sensor weight for the multiple camera units. The simplest solution to exert control over the Tau measurement accuracy is to maintain a constant frame rate and use image processing to change the effective resolution of captured images. This type of image processing is known as subsampling.

#### 4.3.3.4 Image Subsampling

Image subsampling is the process of reducing the effective resolution of an image by averaging pixels across the frame in blocks. Subsampling an image once will combine 4 pixels (a 2 by 2 block) into 1, with a value that is the mean of the original 4 values. Subsampling again will combine 4 of the new pixels into 1, equal to 16 pixels of the original image pixels in a 4x4 block. This concept is illustrated in Figure 4-23.

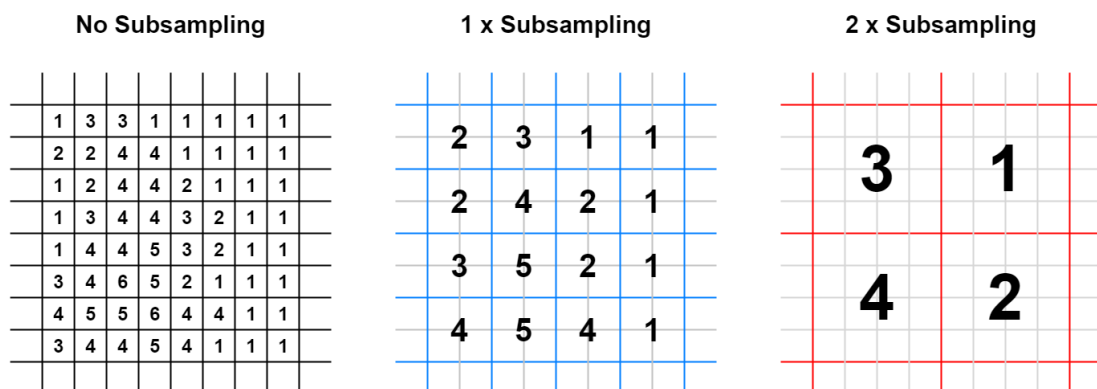


Figure 4-23: A numeric representation of an image plotted over a pixel grid, where the numbers represent the intensity of the image at each pixel location. The base image is shown on the left with corresponding 1- and 2-times subsampled equivalents in the middle on the right respectively

Subsampling was implemented using the virtual helipad scene by creating multiple instances of the frame buffer object used to hold and process frames and gradients. A block-processing method was added to the object class to average regions of input frames and populate a frame of reduced size with the corresponding values. The

block-processing method divides the image matrix into rectangular blocks and applies a custom function to each block. The results of that function in each block are then assembled into a new, smaller matrix. The block processing method used in this case is available in the Matlab image processing toolbox. The results of varying the level of subsampling on a 720p resolution video on the RMS error of the Tau measurement are shown in Figure 4-24. The effect is similar to that of changing the resolution of the camera but does not require any change in camera settings. Subsampling the image three times produces very accurate results within 0.2 seconds of the true value through the final third of the manoeuvre, and almost eliminates the large error divergence; it does begin to increase in the final moments of the descent but the magnitude is small compared to other levels of subsampling.




Figure 4-24: Tau measurement error using the direct gradient method for a constant velocity descent for the helipad landing scene with varying levels of subsampling: (a) no subsampling (b) 1 x subsampling (c) 2 x subsampling (d) 3 x subsampling

This subsampling method provides the ability to maintain an accurate Tau estimate across a manoeuvre if it can be implemented and scheduled efficiently for real-time use.

#### 4.3.3.5 Iteration Count

The more complex extensions to the direct gradient methods (those detailed in Appendix B.2.4 and Appendix B.2.5) require multiple systems of equations to be solved iteratively. The number of iterations does not affect simple cases with motion in one-dimension, but more iterations will be required for complex scenarios. Figure 4-25 shows the RMS error in the Tau measurement using the direct gradient method

for varying numbers of iterations on that method. The visual scene is the helicopter landing with motion in six degrees-of-freedom, captured at a 720p resolution at 30 frames per second.





Figure 4-25: Tau RMS error of direct gradient method estimation for varying iteration number in method.

More iterations on the direct gradient method do decrease the error in Tau a small amount, but this improvement is only seen in a short phase of the manoeuvre after the Tau estimate has already begun to diverge. Figure 4-26 shows the corresponding frame processing times for the varying number of iterations in the direct gradient method. Varying the iteration count does not affect the frame acquisition and averaging processes, but multiplies the time taken for the Tau estimation process as one might expect. Greater numbers of iterations were tested but showed diminishing returns. As a result, it is best to minimise the number of iterations used in the direct gradient method. Spare processing overhead can be directed to this task to improve accuracy slightly.



*Figure 4-26: Mean frame processing time for varying number of iteration in direct gradient method Tau estimation, broken down by the three main operations performed during the calculation*

The length of time to process a frame using this configuration is clearly far too long to allow for useful implementation in real time as discussed in section 4.3.3.3, but the effect of iteration count is still clear.

#### 4.3.4 Simulink Implementation

It was necessary to create an implementation of the direct gradient Tau estimation method in Matlab Simulink to enable integration of the flight simulation models described in section 3.1. Memory for every variable must be pre-allocated for any Simulink simulation, and their size cannot change during the run. As a result, if multiple levels of subsampling are desired for Tau estimation then each must be implemented separately; the size of the image changes when subsampled so different instances are required for each level. It would be inefficient to perform the Tau calculation for every level of subsampling at every time step, so the system is set up to enable different subsystems for different levels of subsampling when desired. The basic structure and flow of the Tau estimator is shown in Figure 4-27.

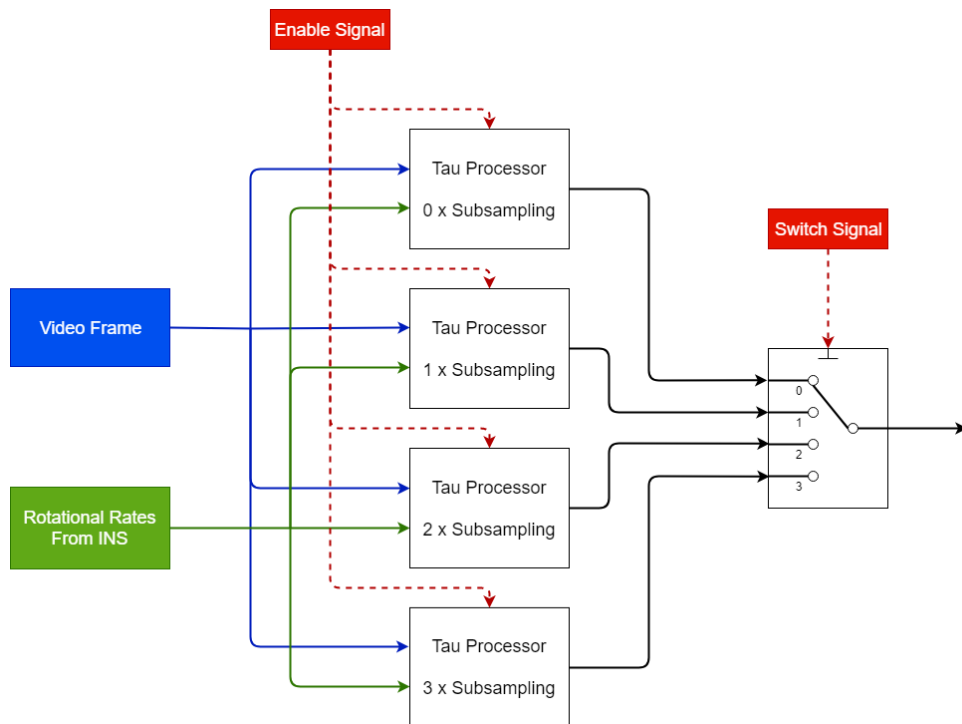


Figure 4-27: Subsampling switching logic for the Simulink implementation of the direct gradient method with multiple levels of subsampling

The enabling of different subsystems is complicated by the temporal averaging required for accurate image gradients. All image parameters are averaged over a buffer, so the full length of the frame buffer must be filled before the Tau estimate will be accurate. Therefore, each system must be enabled  $n$  frames before it can be used, where  $n$  is the length of the frame buffer. Hence, two signals are needed to enable and switch to the different subsampling systems independently.

All four systems can also be enabled for the whole simulation to assess where the best switching points are. Figure 4-28 shows the Tau estimates for four different subsampling conditions, all constantly enabled through the manoeuvre. The video source is the helipad landing scene with only vertical motion, using an HVGA resolution camera with a frame rate of 30. The point of minimum Tau RMS error is marked with a vertical dotted line.




Figure 4-28: Tau RMS measurement error for the Simulink implementation of the direct gradient method for a constant velocity descent in the helipad landing scenario with varying levels of subsampling: (a) no subsampling, (b) 1x subsampling, (c) 2x subsampling, (d) 3x subsampling. The point of minimum error for each configuration is shown by a vertical red dotted line.

The points of minimum error can then be used to define the switching logic, by switching at the midpoint between each minimum. Since the resolution is low to enable a faster execution rate, accuracy is low through the first third of the manoeuvre and the minimum points are clustered quite closely. The resulting switching logic is shown in Figure 4-29. This example of switching logic will give the best results for this camera setup and manoeuvre but may not optimal for all scenarios.




Figure 4-29: Switching (bottom) and enabling (top) logic signals for Simulink implementation of direct gradient method. Subsystems configured to run different subsampling configurations are turned on and off by the enabling signals, and selection of the subsystem that connects to the Tau output of the whole system is selected by the switching signal.

To prevent any step changes in estimate, a length of time equal to half the buffer length was used before and after the switch time to average the two estimates into one. The resulting time-to-contact estimate from this subsampling switching logic is shown in Figure 4-30.





Figure 4-30: Measured time-to-contact and true time-to-contact for the Simulink implementation of direct gradient method. Virtual camera used an HVGA resolution and frame rate of 30, descending towards helipad in one dimension.



The estimate through the second half of the manoeuvre tracks the true value very well, but the signal is very noisy in the first half. This is due to the low camera resolution used to reduce frame processing time. For comparison, the same scene was analysed with a virtual camera using a 720p resolution. The same process for defining subsampling logic was used, which produced a wider spread of switching points. The resulting time-to-contact estimate is shown in Figure 4-31. There is still a significant amount of noise in the first third of the manoeuvre, but it is lower than for the HVGA resolution above. Correspondingly, the accuracy is slightly lower as  $\tau$  approaches zero.



*Figure 4-31: Measured time-to-contact and true time-to-contact for the Simulink implementation of direct gradient method. Virtual camera used a 720p resolution and frame rate of 30, descending towards helipad in one dimension.*

A higher level of subsampling could potentially be used for the terminal phase to address this problem. It is also worth noting that these results are for a one-dimensional descent; scenarios with multi-axis motion tend to reduce measurement noise slightly. This is evidenced by the  $\tau$  estimate in Figure 4-32, where an HVGA camera descends towards the helipad with multi-axis linear and rotational motion. There is still some early noise, but both the amplitude and frequency are lower.




Figure 4-32: Measured time-to-contact and true time-to-contact for the Simulink implementation of direct gradient method. Virtual camera used an HVGA resolution and frame rate of 30, descending towards helipad with multi-axis motion.

From the results in the above figures it is clear that accurate Tau measurements are reliably attainable with this camera configuration through the time interval investigated. However, estimates can become too noisy to be usefully applied above 8 – 10 seconds time-to-contact, and diverge away from the true value below 1 – 2 seconds time-to-contact. Using high resolution cameras should increase this range above 10 seconds, and using high frame rate cameras could extend the range closer to the point of contact.

#### 4.3.5 Fixed-Wing Platform Analysis

The gradient method was also used to analyse videos that represent fixed-wing landing scenarios. These differ from the rotary-wing scenarios examined due to the high speed translational motion that the camera undergoes while descending toward the landing plane; fixed-wing landings are very definitely two-dimensional movements, while the examined rotary-wing landings are mainly one-dimensional. Both the simple runway visual scene described in section 3.3.4 and the VR runway of section 3.4.2 were examined.

##### 4.3.5.1 Simple Runway

Analysis of the simple runway scene was performed first, using two virtual camera recording at 30 frames per second, one at HVGA resolution and one at 720p resolution. Both videos depict a camera descending towards the virtual runway at 5

m/s from an initial height of 50 m, while moving along the runway at 15 m/s. This speed is approximately equivalent to 30 knots, the approach speed of the 3DR Aero UAS platform. The videos were processed using the Simulink implementation of the direct gradient method with multiple levels of subsampling applied. The resulting time-to-contact estimates are shown in Figure 4-33.

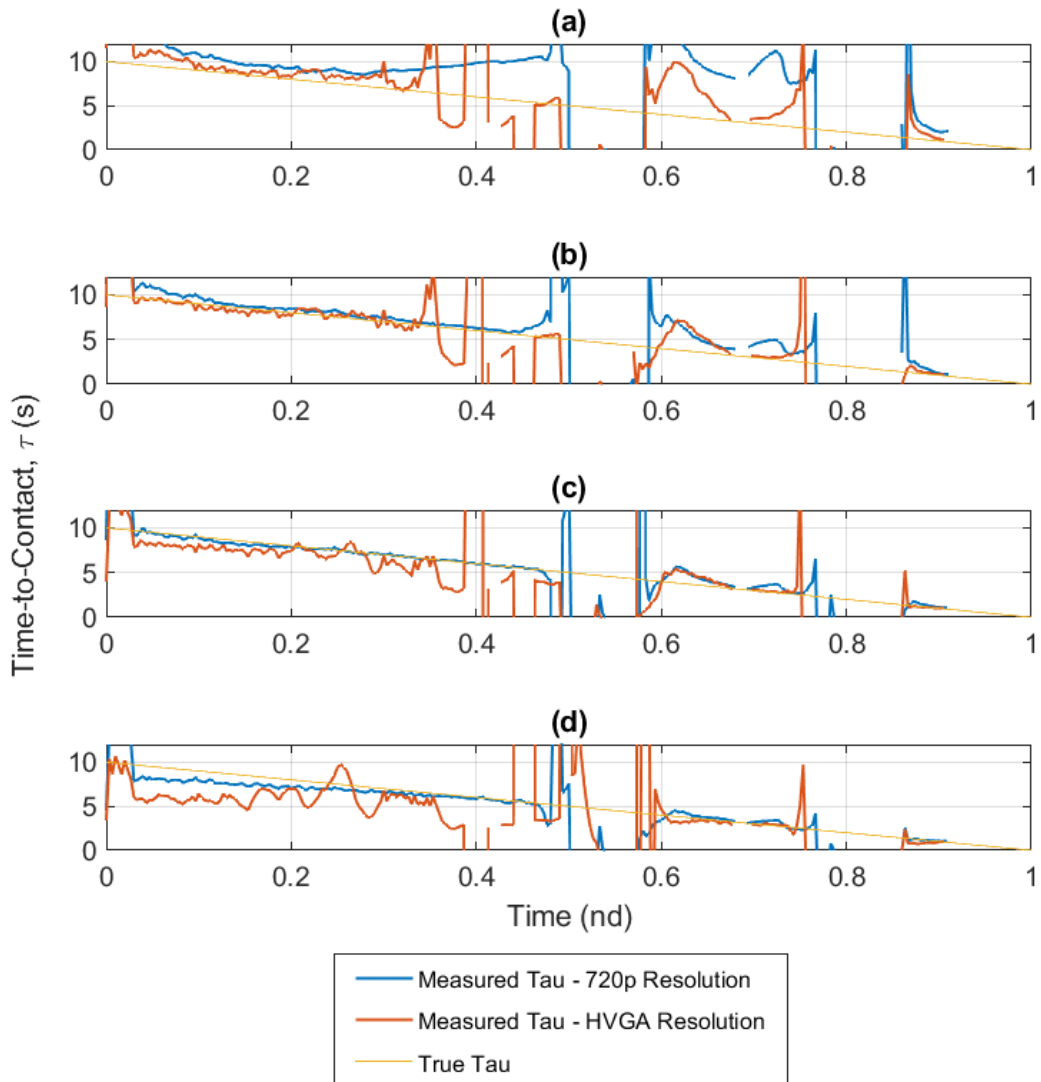


Figure 4-33: Tau estimates for the simple runway landing scene for two different camera resolutions and varying levels of subsampling: (a) no subsampling (b) 1 x subsampling (c) 2 x subsampling (d) 3 x subsampling

The Tau estimate tracks the true value well through the first third of the manoeuvre but begins to diverge heavily from this point. This is caused by the same problem

manifested in the dimension tracking and optical flow methods, where features disappear from the visual scene; shortly before 4 seconds into the manoeuvre the runway threshold markings disappear from view and the amount of information in the image is greatly decreased. As new line features appear in shot, the estimate improves again for short phases, between 0.6 and 0.75, and for a brief period at 0.9 of the total time duration.

As analysis of rotary-wing descents suggested translational motion does seem to decrease noise in the Tau estimate, especially at the start of the manoeuvre and at higher resolutions. As before, estimates from video at higher resolutions and low subsampling levels diverge away from the true value as the manoeuvre progresses. A similar relationship seems to hold true for subsampling; the performance of subsampled estimates is poor early but improves as time-to-contact decreases. It differs slightly here for the fixed-wing case in that the early error for heavily subsampled cases is not due to noise but seems to be a constant offset, though the reason for this behaviour is unclear.

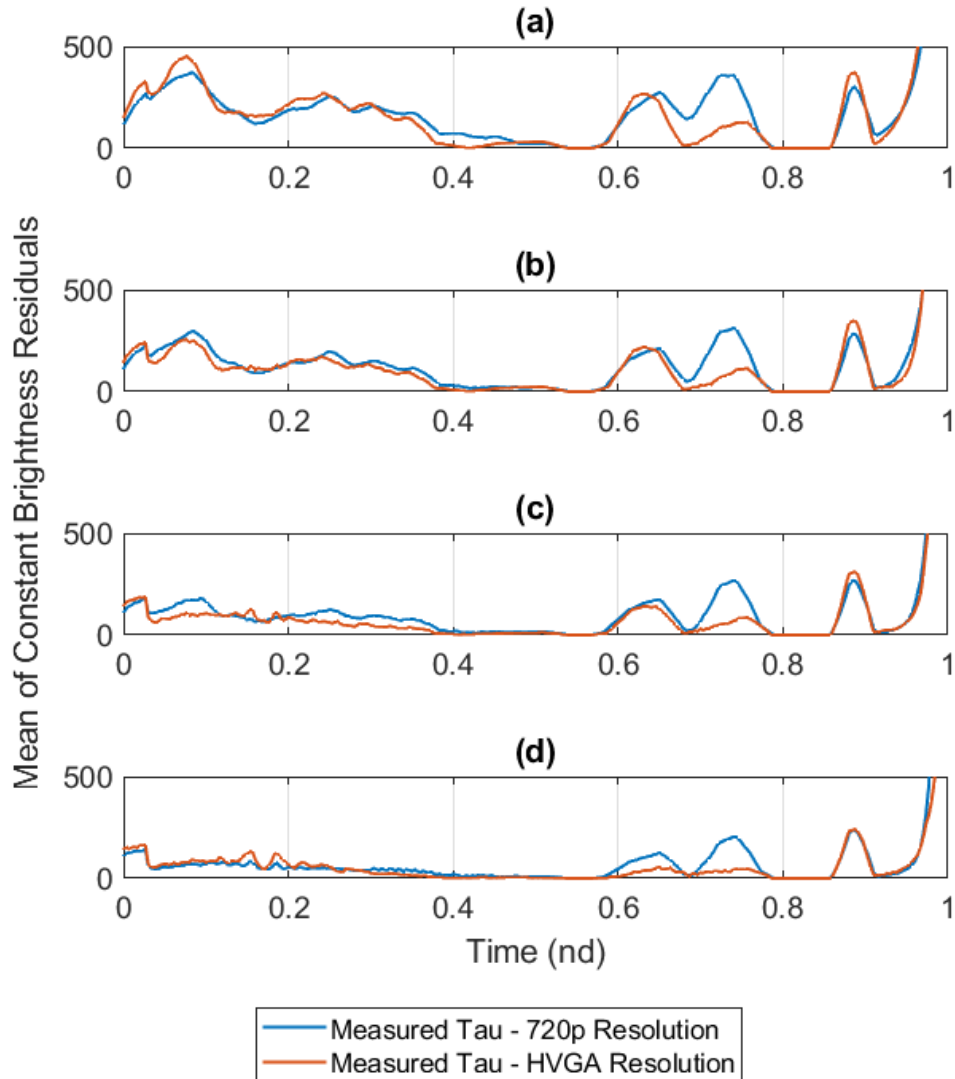


Figure 4-34: Mean of constant brightness assumption residuals for the simple runway landing scene for two different camera resolutions and varying levels of subsampling: (a) no subsampling (b) 1 x subsampling (c) 2 x subsampling (d) 3 x subsampling

It is again useful to analyse the residuals of the constant brightness equation to see where divergences occur again which are shown in Figure 4-34. Interestingly, the mean residual of the constant brightness residual does not seem to be proportional to the error in this situation; the residual in plot (a) is much higher than it is in (d) through the first 20% of the descent, but the Tau estimate in Figure 4-33 is much better for the no subsampling condition (a). Additionally, the residual reduces to zero through the video sections where there are no features in view, and the magnitudes are completely different that those seen in the rotary wing case. These factors

effectively eliminate the possibility of using the constant brightness residual as a metric for prediction of Tau estimate accuracy in cases with translational motion.

#### *4.3.5.2 VR Runway*

Since lack of image content is clearly an issue in analysis of the simple runway videos, the VR runway scene was also used for analysis of fixed-wing landings. An issue with the camera view clipping into the terrain is present in the VR runway videos, which was unavoidable. As a result, no time-to-contact estimate is possible through the last two to three seconds of the descent. This demonstrates the necessity of the simple runway analysis tool where greater camera control is possible. There is still a section of the VR videos where no runway markings are visible, and the image information is insufficient for a Tau estimate between 0.35 and 0.45 of the manoeuvre length. The aircraft is identical to the previous case, with the camera descending from 50 m at  $m/s$ , moving down the runway at 15  $m/s$ . An HVGA resolution camera was used and varying levels of subsampling were applied.

The Tau estimate for this VR video scene is shown in Figure 4-35 for four different levels of subsampling. The extra visual content of the VR scene seems to improve the Tau estimate significantly throughout the manoeuvre, as error is generally much lower, and estimates for all subsampling levels are less noisy. As before, estimates from low subsampling levels are closer to the true value early in the manoeuvre, while higher subsampling is more effective as time-to-contact approaches zero.





Figure 4-35: Tau estimates from the direct gradient method for the VR runway landing scene for varying levels of subsampling: (a) no subsampling (b) 1 x subsampling (c) 2 x subsampling (d) 3 x subsampling

Combining the four levels of subsampling into a single signal through the Simulink switching logic, a single Tau estimate was obtained and is shown in Figure 4-36. In contrast to the results of the rotary-wing analysis, the estimate contains more noise close to the ground, rather than far away. The magnitude of optical flow in each scenario (early in a two-dimensional descent, versus late in a one-dimensional descent) seems to be similar, though rotary-wing analysis suggested that flow magnitude was not the driving factor in estimate error. However, higher subsampling close to the ground does still provide superior performance. It does appear a good Tau estimate is feasible through most of the fixed-wing landing if sufficient visual


information is available throughout the descent, but performance is poorer close to the ground.



*Figure 4-36: Combined Tau estimate from the direct gradient method for the VR runway landing scene. Aircraft descends from 50m at 5 m/s, while travelling down the runway at 15 m/s. Scene was imaged at HVGA resolution.*


The approach speed of the 3DR Aero is very low, and the majority of fixed-wing aircraft will have a much higher stall speed. Consequently, tests were also carried with a camera moving down the runway at higher speeds. Figure 4-37 shows a camera moving down the runway at 40 m/s while descending, more similar to the approach of a larger, general aviation aircraft.





*Figure 4-37: Combined Tau estimate from the direct gradient method for the VR runway landing scene. Aircraft descends from 50m at 5 m/s, while travelling down the runway at 40 m/s. Scene was imaged at HVGA resolution.*

The Tau estimate from this higher speed landing supports the idea that higher optical flow magnitude increases the noise in the Tau estimate; it is higher than the slower equivalent in Figure 4-36, and increases as the camera descends and flow increases further. A possible solution to this problem is to increase the frame capture rate of the camera to effectively reduce the optical flow per frame magnitude. Figure 4-38 shows the RMS error in the Tau estimate for both a 30 and 60 frame-per-second video of the above scenario.



*Figure 4-38: Tau RMS measurement error from the direct gradient method for the VR fixed wing landing at 40 m/s, for two different cameras capturing at 30 fps and 60 fps respectively*

As predicted, the Tau RMS error is reduced for the higher frame rate video, suggesting that this is a helpful approach for fixed-wing cases at high speed. However, using a frame rate of 60 instead of 30 obviously doubles the computational load, with twice as many estimates required per second. As a result, extra computational resources maybe required for fixed-wing implementations of Tau that must run in real time.

## 4.4 Direct Gradient Method Extension for Brightness Changes

It has been demonstrated in the previous subsection that motion by a camera can violate the constant brightness assumption that forms the basis of both optical flow and direct gradient methods for Tau estimation, and that these violations can lead to reduced estimate accuracy. However, there are other factors that can also break this fundamental principle. The most obvious is changing image brightness; the illumination conditions in the visual scene can change captured images significantly. The visual scenes used for analysis so far have all used constant global illumination so this has not been a problem, but in reality, illumination of the visual scene can be variable. The amount of light entering the camera may change as clouds move or shadows develop from vehicle movement.

Modifications to the constant brightness assumption have been previously proposed to extend optical flow field estimation methods. Section 2.4.2.2.4 discussed one approach to dealing with changing illumination conditions by modelling brightness changes as a multiplicative field and two-dimensional offset to the image, originally detailed in reference [86]. This approach has not been previously applied to the direct gradient method for Tau estimation, but the concept is theoretically applicable in a similar way. An extension of the direct gradient method to account for global illumination changes is proposed in this section.

### 4.4.1 Overview

The proposed extension is based on a modification of the constant brightness equation, which for clarity is given again in equation (4-38):

$$I_x u + I_y v + I_t = 0 \quad (4-37)$$

The modified version, here termed the variable brightness equation, is the sum of the image gradients and optical flow components equalled to the raw image,  $I$ , multiplied by a scalar multiplier,  $m$ , instead of set equal to zero as in the constant brightness equation:

$$I_x u + I_y v + I_t = I m \quad (4-38)$$

This approach is slightly different from the similar extensions for optical flow methods, as it is restricted to using a single scalar multiplier, rather than a multiplicative field and offset. Optical flow is by definition a vector field, so computing extra fields in that process is achievable. The central motive of the direct gradient method is to estimate Tau without recovering the optical flow field, so it would be counterproductive to introduce more fields that need to be recovered. To use a scalar field of brightness change in the gradient method, some physical model of how brightness changes across the image would have to be implemented. This is analogous to how the pinhole camera model is used to create expressions for the optical flow vectors in the gradient method. It would necessitate knowledge of the environmental geometry, vehicle geometry and dynamics, and parameters of the illumination sources. Therefore, a single brightness change multiplier is used to account for illumination changes instead. Since all parameters are summed across the image in the direct gradient method, this single multiplier can be considered an average of a field of multipliers.

#### 4.4.2 Derivation

The same approach to expanding the constant brightness equation for the gradient method is applied to the variable brightness equation, by substituting in expressions for the optical flow components. The expressions for the optical flow components are the same as those used for the direct gradient method described in Appendix B.2.5:

$$u = \left( f \frac{U}{Z_0} - x \frac{W}{Z_0} \right) \left( 1 - x \frac{p_s}{f} - y \frac{q_s}{f} \right) + \frac{xy}{f} \omega_x - \frac{x^2 + f^2}{f} \omega_y + y \omega_z \quad (4-39)$$

$$v = \left( f \frac{V}{Z_0} - y \frac{W}{Z_0} \right) \left( 1 - x \frac{p_s}{f} - y \frac{q_s}{f} \right) + \frac{y^2 + f^2}{f} \omega_x - \frac{xy}{f} \omega_y - x \omega_z \quad (4-40)$$

To simplify these expressions, seven parameters are defined as combinations of the variables in equations (4-39) and (4-40), which are all detailed in section 4.3.1. Equations (4-39) and (4-40) are substituted into (4-38) to create the fundamental equation of the extension. Equations (4-20) through (4-26) are used to simplify this expression to:

$$\left(I_x A + I_y B + C(I_x x + I_y y)\right) \left(1 + x \frac{P_s}{C} + y \frac{Q_s}{C}\right) + I_x J + I_y K + I_t = I_m \quad (4-41)$$

This expression only contains six unknown variables, while the rest are all parameters of the camera or image. Several more parameters are then used to reduce the equation further, detailed in section 4.3.1. Substituting equations (4-27) through (4-30) into equation (4-41) yields the much simpler expression:

$$CFD + I_\omega = I_m \quad (4-42)$$

To find the unknowns in equation (4-42) it is best to formulate it as a least squares minimisation in the form:

$$\sum (CFD + I_\omega - I_m)^2 \quad (4-43)$$

Equation (4-43) can then be differentiated with respect to each of the six unknown quantities to find the minimum points. However, the resulting equations are nonlinear, so to find a solution the hierarchical solution proposed by Horn in reference [119] must be used. This approach sets the terms of  $D$  and  $F$  as constants in turn and solving for the other in an iterative process. Assuming  $P_s/C$  and  $Q_s/C$  are constant and therefore  $F$  is constant, then substituting back in the expression for  $D$  gives the equation:

$$\sum \left( CF \left( G + I_x \frac{A}{C} + I_y \frac{B}{C} \right) + I_\omega - I_m \right)^2 \quad (4-44)$$

Differentiating equation (4-44) with respect to  $A$ ,  $B$ ,  $C$  and  $m$  respectively:

$$\sum (F(CG + I_x A + I_y B) + I_\omega - I_m) F I_x = 0 \quad (4-45)$$

$$\sum (F(CG + I_x A + I_y B) + I_\omega - I_m) F I_y = 0 \quad (4-46)$$

$$\sum (F(CG + I_x A + I_y B) + I_\omega - I_m) F G = 0 \quad (4-47)$$

$$\sum -(F(CG + I_x A + I_y B) + I_\omega - I_m) I = 0 \quad (4-48)$$

Then, assuming  $A$  and  $B$  are constants and therefore  $D$  is constant, substituting back in the expression for  $F$  gives the equation:

$$\sum \left( CD \left( 1 + x \frac{P_s}{C} + y \frac{Q_s}{C} \right) + I_\omega - Im \right)^2 \quad (4-49)$$

Differentiating equation (4-52) with respect to  $P_s$ ,  $Q_s$ ,  $C$  and  $m$ :

$$\sum (D(C + xP_s + yQ_s) + I_\omega - Im)Dx = 0 \quad (4-50)$$

$$\sum (D(C + xP_s + yQ_s) + I_\omega - Im)Dy = 0 \quad (4-51)$$

$$\sum (D(C + xP_s + yQ_s) + I_\omega - Im)D = 0 \quad (4-52)$$

$$\sum -(D(C + xP_s + yQ_s) + I_\omega - Im)I = 0 \quad (4-53)$$

If you multiply out equations (4-45) through (4-53) and re-arrange them then they can then be formulated into two linear matrix equations of the form  $Ax = b$ , that can be solved iteratively to yield estimates of the six unknown parameters, including time-to-contact:

$$\begin{bmatrix} \sum F^2 I_x^2 & \sum F^2 I_x I_y & \sum F^2 G I_x & -\sum F I_x I \\ \sum F^2 I_x I_y & \sum F^2 I_y^2 & \sum F^2 G I_y & -\sum F I_y I \\ \sum F^2 I_x G & \sum F^2 I_y G & \sum F^2 G^2 & -\sum F G I \\ -\sum F I_x I & -\sum F I_y I & -\sum F G I & \sum I^2 \end{bmatrix} \begin{bmatrix} A \\ B \\ C \\ m \end{bmatrix} = - \begin{bmatrix} \sum F I_x I_\omega \\ \sum F I_y I_\omega \\ \sum F G I_\omega \\ -\sum I I_\omega \end{bmatrix} \quad (4-54)$$

$$\begin{bmatrix} \sum D^2 x^2 & \sum D^2 xy & \sum D^2 x & -\sum DxI \\ \sum D^2 xy & \sum D^2 y^2 & \sum D^2 y & -\sum DyI \\ \sum D^2 x & \sum D^2 y & \sum D^2 & -\sum DI \\ -\sum DI & -\sum DIx & -\sum DIy & \sum I^2 \end{bmatrix} \begin{bmatrix} P_s \\ Q_s \\ C \\ m \end{bmatrix} = - \begin{bmatrix} \sum DxI_\omega \\ \sum DyI_\omega \\ \sum DI_\omega \\ -\sum II_\omega \end{bmatrix} \quad (4-55)$$

It is important that the only parameter needed to solve these equations (above those needed for the base direct gradient method) is the image itself. Therefore, only an image, its derivatives, and a couple of base camera parameters are needed in total to produce a Tau estimate, and this is still possible with changes of illumination between video frames with the proposed changes.

Although it does not seem immediately clear that the small initial change to the constant brightness assumption of adding the  $Im$  term in equation (4-38) will have this effect; the extra image term effectively provides a multiplier for the temporal gradient that modulates it if there is a global change in image brightness. This is best illustrated in the 4<sup>th</sup> element of the vector on the right-hand side of equations (4-54) and (4-55),  $-\sum II_{\omega}$ . If the global image brightness,  $I$ , decreases, then the change will register in the temporal gradient,  $I_{\omega}$ , and increase the magnitude of the gradient. Without this correction, this change in temporal gradient would register as camera motion and disrupt the time-to-contact estimate. However, because the gradient is multiplied by the image frame itself, the decrease in brightness (captured as a reduced magnitude of the elements in  $I$ ) will effectively reduce the value of the gradient to cancel out the erroneous initial measurement. The full process of the implemented corrected gradient method is as follows, with the extra conceptual step highlighted in bold italic text:

1. Capture a video frame
  - 1.1. Store the video frame in a frame buffer
2. Calculate the image gradients in the X & Y directions of the captured video frame,
  - 2.1. Store the X & Y gradients in a 2nd buffer
3. Calculate the image gradient in the time direction (between frames) using the current frame and previous frame in the frame buffer
  - 3.1. Store the temporal gradient in a 3rd buffer
4. Calculate the radial gradient of the frame using the X & Y gradients and the pixel positions in the image frame of reference (equation (4-28))
  - 4.1. Store the radial gradient in a 4th buffer
5. Average the spatial, temporal and radial gradients over the length of the 4 respective buffers
6. Assemble parameter groups (equations (4-27) to (4-30))
7. Calculate the expected movement between the current video frame and previous frame caused by rotational motion of the camera; rotational rates

from a gyroscope are combined with the pixel positions in the image frame of reference (equations

8. Combine the expected movement from rotational motion of the camera with the averaged image temporal gradient to find a corrected temporal gradient
9. ***Use the raw video frame to modulate the corrected temporal gradient, which corrects for changes in the image caused by illumination changes, rather than relative camera motion***
10. Estimate translational motion of the camera from ratios of the spatial and radial gradients to the corrected temporal gradient, as well as an initial guess of relative slope between camera and target (equation (4-54))
11. Use the estimated translational motion of the camera from step 8 and the ratios of spatial, radial and temporal gradients to estimate relative slope between camera and target (equation (4-55))
12. Iterate through steps 10 and 11 multiple times using updated estimates of translational motion and relative slope to sequentially improve estimate accuracy of each quantity
13. Calculate time-to-contact from the estimated translational motion parameters

A flow diagram of this process with some example image and gradient frames is provided in Appendix B.3.2 for further clarity.

#### 4.4.3 Testing

This brightness change direct gradient method was implemented using the same frame buffer Matlab class of objects described in section 4.3.1. A small amendment was made to the Matlab class and associated Simulink implementation to also log and average the actual image frames themselves over the length of the buffer. This was necessary to maintain temporal alignment with the image gradients used in the process; the time averaging must be applied to all variables used in the calculation or they will represent different points in time.

The virtual helipad landing scene was used to test the extension to the direct gradient method. The basic video generator only produces grayscale images on an unsigned



8-bit integer scale; the colour of the helipad can vary between 0 (black) and 255 (white). To simulate changes in global illumination, a simple multiplier was applied to each video frame that could be varied with time.

Three test scenarios are shown below, with three different brightness change profiles, all plotted in Figure 4-39. The three scenarios are:

1. The first scenario is a gradual decrease of brightness to 40% of the initial over two seconds.
2. The second test is another gradual decrease of brightness over two seconds, but brightness reduces to 20% of its initial value.
3. The third test examines a step change in brightness to 40% of initial value.

These brightness profiles are intended to demonstrate the effectiveness of the method and are not necessarily representative of any particular physical scenario. The camera trajectory is identical for each of the three scenarios; the camera descends from 50m at a constant speed of 5m/s in the vertical axis.





Figure 4-39: Global image brightness level change over time for the three scenarios used to test the brightness change extension to the direct gradient method

All three videos are captured at 30 frames per second at an HVGA resolution, and all depict the same descent towards the helipad. The camera descends from a height of

50 m at a vertical speed of 5 m/s. Small amplitude sinusoidal translational motion occurs in the X axis, in addition to small sinusoidal rotations in the roll and yaw axes. Four example frames from the test two video are shown in Figure 4-40 to illustrate how the brightness changes over time.



*Figure 4-40: Example frames from the helipad scene with global brightness changes from brightness change test scenario two at varying altitudes: (a) 27.5 m (b) 22.5 m (c) 17.5 m (d) 12.5 m*

Tau estimates for test scenario one are plotted in Figure 4-41. Results from both the standard direct gradient method assessed in section 4.3, and the proposed brightness change extension are shown, in addition to the true value of time-to-contact.




Figure 4-41: Tau estimate from the direct gradient method for test scenario one with and without correction from brightness change extension

The Tau estimates are identical while global image brightness is constant, but as the brightness change begins halfway through the manoeuvre, the error of the uncorrected estimate increases. This increased error occurs due to the violation of the constant brightness assumption that underpins the method. The estimate using the brightness change correction continues to track the true time-to-contact value well through the brightness change phase, before the two estimates converge again when brightness becomes constant again. This result is encouraging and suggests that the proposed extension to account for brightness change works as intended. Figure 4-42 shows the Tau estimates produced from test scenario two, where brightness changes gradually but by a larger amount.

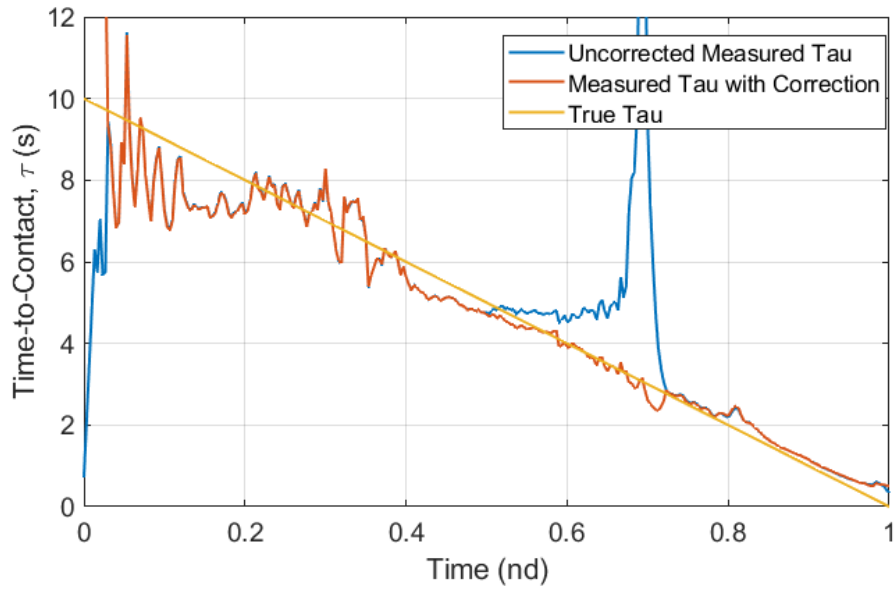


Figure 4-42: Tau estimate from the direct gradient method for test scenario two with and without correction from brightness change extension

The error of the uncorrected estimate increases dramatically during the brightness change in the test two scenario. However, the estimate corrected for brightness change still tracks the true value of Tau well throughout the event, with minimal performance degradation. The results from the final test scenario are shown in Figure 4-43.




Figure 4-43: Tau estimate from the direct gradient method for test scenario three with and without correction from brightness change extension

Test three assesses a step change in the brightness of the image. This is an extreme case that is unlikely to be encountered in the real world. The movement of a shadow over the visual scene could produce a similar effect. This step change produces a spike in the error of both corrected and uncorrected estimates of Tau, but the amplitude is approximately 25% of the uncorrected value with the brightness change correction applied. Interestingly, the correction produces a doublet pattern in the estimate, with error switching from positive to negative just after the step change, before returning to track the true value. The reason for this behaviour is unclear, but the errors are still smaller than for the uncorrected method.

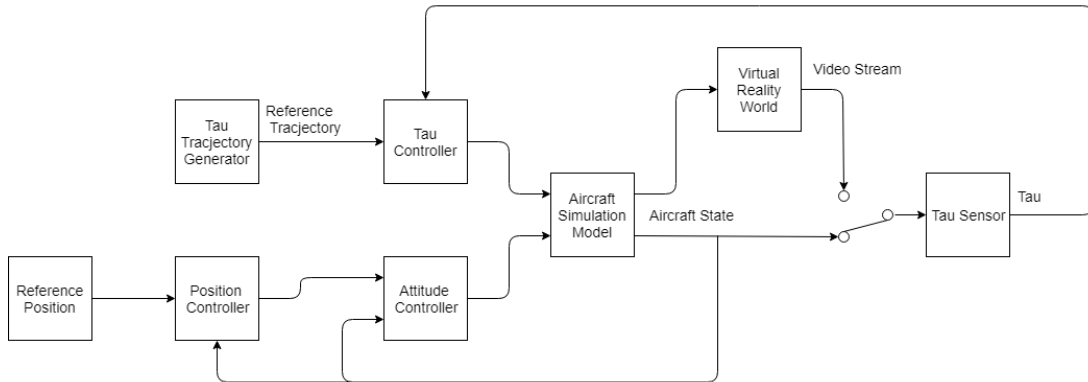
These results support the hypothesis that the proposed extension to the direct gradient method can indeed mitigate the effects of brightness change in the visual scene while perceiving Tau.

## 5 Chapter 5 –Tau Guided UAS Landings

### 5.1 Rotary Wing Tau Guidance System Implementation

A Tau guidance system has been implemented in Simulink and connected to a framework for simulating linear aircraft models. The architecture is independent of any particular aircraft model so that multiple models can be simulated by loading different configuration parameters via a number of setup scripts. Configurations for all of the linearised aircraft models described in section 3.2.3 were created. A separate model was used to integrate the SimMechanics-based 3DR X8 model described in section 3.2.1.

The Simulink Virtual Reality toolbox was also integrated into this architecture to allow generation of virtual camera images for changing aircraft position and orientation in the simulation loop. A simple block diagram of the system architecture is shown in Figure 5-1.



*Figure 5-1: High-level block diagram of simulation architecture that is implemented in Simulink for testing of Tau guidance systems*

A Tau trajectory generator passes reference values of time-to-contact to the Tau controller. The trajectory generator can generate first, second or third order intrinsic Tau guides, each selected by a simple switch. The Tau controller only has control authority over the collective control channel of the simulation model. The other channels are controlled by an inner attitude feedback control loop, and an outer position control loop. These systems use PID controllers to keep the aircraft hovering

over the centre of the ship deck, while holding a constant heading. Tau estimates can be generated from spatial parameters of the aircraft state, or estimated from the VR video stream using the direct gradient method with brightness gradient extension detailed in section 4.4. Whichever Tau estimate is desired can be selected with a switch and is then fed back to the Tau controller.

A detailed block diagram of the full flight controller architecture used for the following testing is given in Appendix C-1. The same controller architecture was used for all four aircraft models that were tested, but the control gains used were adjusted between models. A table containing all of the relevant control gains for all four rotorcraft is also given in Appendix C-2 to allow for full reconstruction of the system when combined with the linear models in Appendix A. The control gains were coarsely tuned manually, first for basic stability, and then to eliminate any undesirable oscillations. The gains should be tuned more rigorously in any future work to further reduce any oscillations present, reduce settling time of the flight control system when changes are commanded, and improve disturbance rejection capabilities, which will be vital for real operations in a maritime environment where atmospheric conditions can be turbulent. Simulink does include tools for automatic gain tuning which could be used for this task, but it may be more effective to redesign the controllers using Bode diagram, root locus, or Nichols plot design methods [120].

### 5.1.1 Tau Control Law Analysis

Several control laws have been previously proposed for implementing Tau guidance systems, as discussed in section 2.5 and detailed in reference [30], but it is not clear which is most effective and how the control laws interact with the various Tau guidance laws. This section gives a brief analysis of three candidate control laws and how perform when applied to the first-, second- and third-order intrinsic Tau guidance laws detailed in section 2.2.5.4. The Tau guidance laws will necessarily generate singularities and zeros at various points, so any guidance law must deal with these problems. The simplest is a proportional feedback control law, where the control input is just the difference between the reference Tau guide value and the measured value, multiplied by a gain,  $K_p$ :

$$u_c = K_p [\tau(t) - \tau_{ref}(t)] \quad (5-1)$$

This law is simple but will pass any infinite values of Tau straight on as a control input unless some saturation functions are applied. The second candidate control law is an inverse proportional feedback law. The inverse of measured Tau and the reference guide are used instead of the raw values, so that large values of Tau cause the error signal to tend towards zero, instead of infinite:

$$u_c = K_p \left[ \frac{1}{\tau(t)} - \frac{1}{\tau_{ref}(t)} \right] \quad (5-2)$$

The third and final law assessed is a nonlinear ratio control, where the ratio of reference and measured Tau is used as the error signal:

$$u_c = K_p \left[ 1 - \frac{\tau_{ref}(t)}{\tau(t)} \right] \quad (5-3)$$

These three control laws have been implemented separately to drive the collective signal of the 3DR X8 simulation model and applied to reference Tau trajectories generated by each of the three Tau guidance laws defined. Since values of infinity can cause errors in the Simulink solver, a saturation function has been applied to limit the maximum value of both the measured and reference Tau values to a magnitude of 100 s. In this case, the quadcopter is landing on a flat, static surface. Tau is calculated from measurements of the aircrafts position and velocity for these tests so that Tau is ‘perfectly’ known, and therefore no problems are encountered due to Tau measurement accuracy.

#### 5.1.1.1 First-Order Intrinsic Tau Guidance Law

The tracking performance of the three control laws are shown for a first-order intrinsic Tau guide manoeuvre in Figure 5-2. All laws are able to track the reference Tau trajectory well, but the proportional guidance law encounters issues at landing, when it commands an upward acceleration. This leads to the aircraft touching down on the runway and then lifting back off and accelerating away. This is caused by the infinite value of Tau measured as the velocity of the aircraft reaches zero as it touches the ground, leading to unpredictable control inputs.






Figure 5-2: Trajectory and Tau tracking RMS error of three candidate control laws applied a first order intrinsic Tau guide

This divergence and subsequent take-off is also a risk for the inverse controller as the error signal will become infinite as  $\tau_{ref}$  tends to 0. The acceleration should be towards the deck, keeping it grounded, but singularities can behave unpredictably in software.

#### 5.1.1.2 Second-Order Intrinsic Tau Guidance Law

For a second order Tau guide, plotted in Figure 5-3, the inverse control law showed poor tracking performance through the first two seconds of the manoeuvre, and the proportional control law also exhibited some noise at the start and end of the manoeuvre. Having already observed the divergence at the end of the manoeuvre for the proportional control law, the proportional controller was disabled at the end of the Tau manoeuvre to prevent the undesired take-off. The nonlinear ratio controller had the lowest tracking error of the three control laws throughout the manoeuvre.




Figure 5-3: Trajectory and Tau tracking RMS error of three candidate control laws applied a second order intrinsic Tau guide

### 5.1.1.3 Third-Order Intrinsic Tau Guidance Law

A third-order guide was also assessed for the three control laws and the results are shown in Figure 5-4. Similar to the second-order guide, the inverse controller shows poor performance during manoeuvre initiation, this time extending to almost 4 seconds. The proportional control law also creates a noisy error signal for the first 2 seconds of the manoeuvre. Again, the nonlinear ratio controller exhibits the best tracking performance, so this law was chosen as the method for implementing a Tau controller.

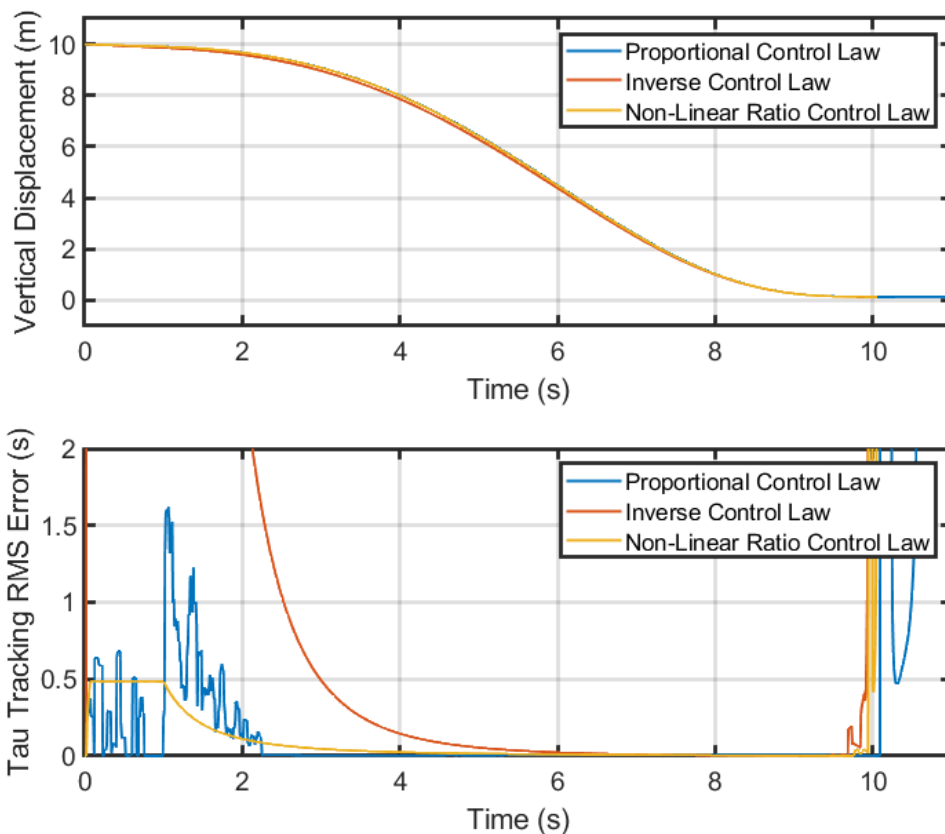



Figure 5-4: Trajectory and Tau tracking RMS error of three candidate control laws applied a third order intrinsic Tau guide

## 5.2 Frigate Deck Tau Guided Landings

The Tau guidance system was evaluated in the context of the ship deck landing MTE detailed in section 3.5.1. Since the Fire Scout is used for operational maritime landings, the MQ-8B linear model was used for the majority of analysis. The frigate deck VR world was used as the visual scene, and the deck was actuated with motion from varying sea states. A simple model for collision mechanics was implemented to detect when the landing gear of the aircraft contacted the deck of the frigate, which triggered a switch to update the position of the aircraft in relation to the motion of the ship deck, rather than with the output of the linear model.

Figure 5-5 shows the position of the aircraft, and relative velocity between aircraft and ship deck, for two landing manoeuvres performed by the MQ-8B. The aircraft begins in a hover 10 m above the deck and descends vertically, following an intrinsic Tau guide reference trajectory using spatially computed Tau estimates. Both second- and third-order intrinsic Tau guides have been tested here with coupling constants of  $k = 0.4$ , but a first-order guide is omitted since the rotorcraft begins at rest, contrary to the assumed initial conditions of a first-order guide. The deck is actuated with motion calculated for sea state 4. Note the heave velocity of the deck has some low amplitude noise that also manifests in any relative velocity measurements, though this does not seem to significantly affect overall performance of the system.



*Figure 5-5: MQ-8B Fire Scout heave dynamics through a sea state 4 ship deck landing manoeuvre using both second- and third-order intrinsic Tau guides*

The descents are noticeably different to those on a static helipad examined previously, since the aircraft seems to track the motion of the ship deck throughout the manoeuvre. The relative velocity profiles are of the expected form given by the

intrinsic Tau guide equations defined in sections 2.2.5.1 and 2.2.5.3, demonstrating that Tau guidance is reactive to its targeted end point, in this case the ship deck. This is evident when examining the terminal phase of the second-order guide Tau guide flight; the aircraft descends for the first 9 seconds of the manoeuvre but begins to climb again as the deck of the ship passes through a trough and begins to rise. The distance between the aircraft and deck is still closing, but the aircraft is climbing in the inertial frame to attempt to track reference value of time-to-contact provided to the controller. It is also clear in the  $k = 0.5$  Tau guide that the aircraft decelerates as the deck reaches its peak between 5 and 6 seconds, and then accelerates again to catch up as the deck drops. This behaviour can be made more or less obvious by tuning the coupling constant,  $k$ , as illustrated in Figure 5-6.




Figure 5-6: MQ-8B Fire Scout heave displacement through a sea state 4 ship deck landing manoeuvre using a second-order intrinsic Tau guide with varying coupling constants

For lower values of coupling constant, the aircraft will descend quickly in the first half of the manoeuvre and decelerate in relation to its target early, descending more slowly when close to the deck. This leads to the aircraft tracking the motion of the deck a short distance above it, and in this case nearly 2 full seconds where the rotorcraft is climbing to maintain separation with the deck, but still only makes contact as relative velocity closes to zero. This may present a problem in higher sea states where the amplitude of deck motion is higher, and the aircraft may not be able

to meet the demands required to track the deck. A higher value of  $k$ , such as the 0.5 case illustrated, will decelerate more slowly in relation to the deck at first, then brake more aggressively as distance closes. This will apply to a greater extent for cases where  $k > 0.5$ , and will lead to positive contact with the deck at a speed above 0. As a result, it is recommended that for autonomous deck landings using Tau guidance a coupling constant of  $0.4 \leq k \leq 0.5$  should be used.

### 5.2.1 Benefits of Tau Guidance for Deck Landings

The inherent ability of Tau guided systems to react to changing environmental dynamics is extremely useful, especially for situations such as ship-deck landings. Figure 5-7 shows this ability in comparison to a common alternative approach, where the aircraft descends at a constant rate. This rate reduces by half as the aircraft approaches the deck to reduce touchdown velocity.

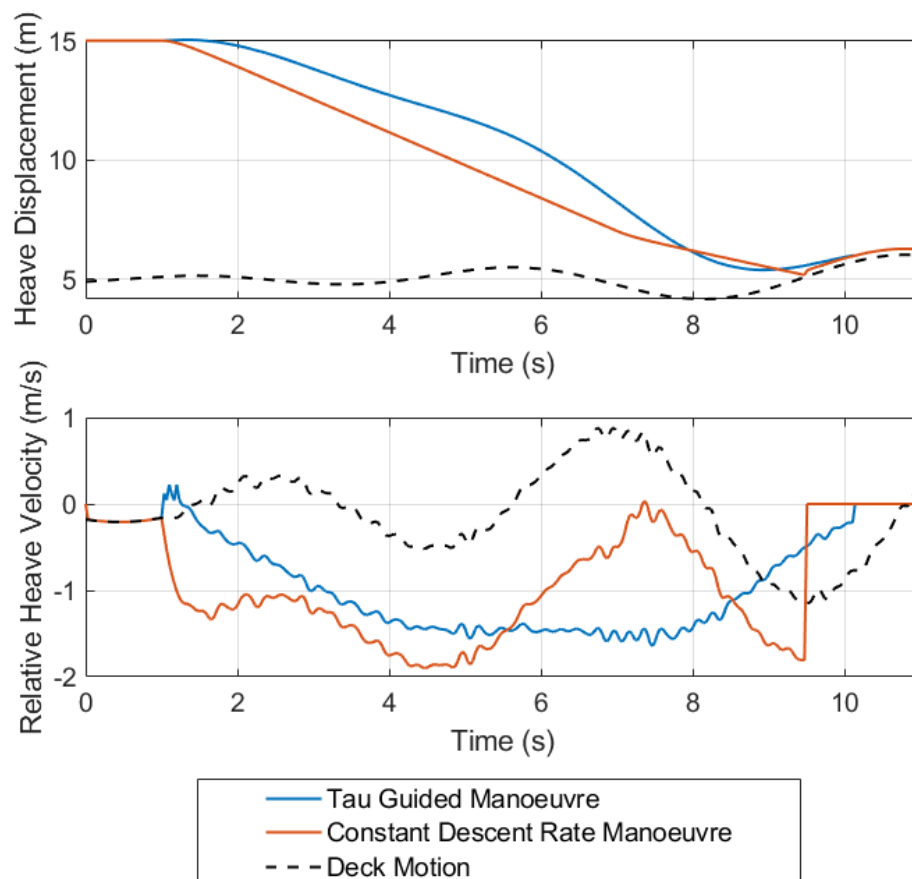


Figure 5-7: MQ-8B Fire Scout heave dynamics through a sea state 4 ship deck landing manoeuvre using both second-order intrinsic Tau guide, and a constant descent rate manoeuvre

The descent rate is measured in the inertial frame and fed back into a PID controller along with the reference descent rate. The resulting error signal is used to drive the collective channel of the helicopter. In hardware the descent rate of the aircraft would be found from fused GPS-INS data and would not be reactive to any environmental dynamics. This is clear from the relative velocity trace for the constant rate descent, as it varies with the velocity of the deck itself, with an offset equal to the reference descent rate. Since this method does not take into account the motion of the deck, soft contact with the deck cannot be assured and the aircraft hits the ship at a velocity of 1.8 m/s.




Figure 5-8: MQ-8B Fire Scout heave dynamics through a sea state 4 ship deck landing manoeuvre using both second-order intrinsic Tau guide, and a constant relative descent rate manoeuvre

A variant of this system was also used for comparison, shown in Figure 5-8, where the aircraft descends at a constant rate relative to the deck. This is similar to the UCARS system used for the maritime recovery of the MQ-8B Fire Scout by the US Navy and mentioned in section 2.1.1; a beacon on the landing deck measures ship

motion and relays it to the aircraft so it may account for the motion in its approach. In this case, it is possible to attain lower and more predictable touchdown velocities but relies heavily on the measure of relative position and velocity. This reliance was illustrated by the Fire Scout accident reported in reference [61], where communications issues with the beacon used to measure the relative velocity were disrupted.

This behaviour is also present while following a Tau guide that is not reactive to the motion of the deck. For example, if using spatial parameter estimates that do not include measurement of the deck motion. Instead, Tau is measured relative to a fixed datum, such as the expected height of the deck in sea state 1, instead of relative to the actual position of the deck. Figure 5-9 shows that this approach yields a similar result to the constant rate example and the aircraft hits the deck at 1.7 m/s. These results support the idea that Tau should be directly measured to properly utilise the benefits of Tau guidance.



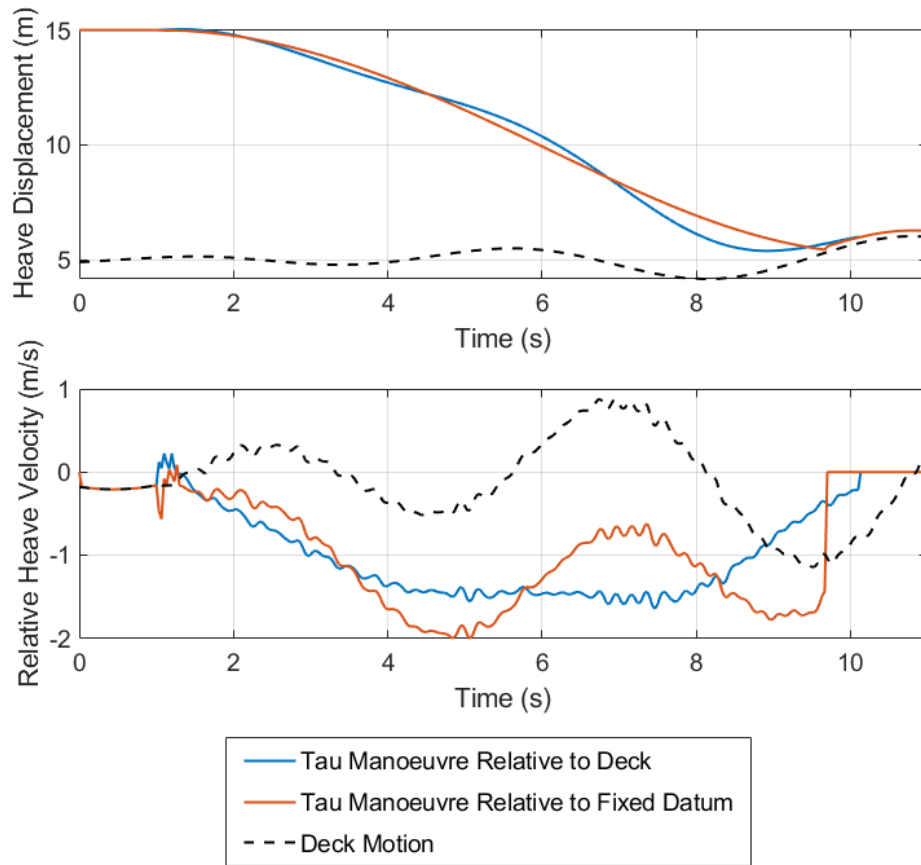


Figure 5-9: MQ-8B Fire Scout heave dynamics through a sea state 4 ship deck landing manoeuvre using a second-order intrinsic Tau guide, measured relative to the deck and a fixed datum

### 5.2.2 Start Point Variation

It is conceivable that the Tau guide used in the examples so far is taking advantage of some quiescent period during the deck motion and will not be as effective in slightly different scenarios. To address this, a further experiment was conducted that varied the start point of the manoeuvre in time relative to the 30 second period of recorded ship motion data. To achieve this, the start and end points of the ship motion were matched and smoothed to allow the data to be looped, and 20 equally spaced points in time were used as start points for the Tau landing manoeuvre. For the sake of comparison, the process was repeated with a Tau manoeuvre where Tau was measured relative to a fixed datum, i.e. not reactive to the deck.




Figure 5-10: MQ-8B Fire Scout touchdown velocity for sea state 4 deck landings using a second-order intrinsic Tau guide. Touchdown velocity plotted for Tau manoeuvres reactive to the deck motion, and without reactivity to the deck motion are shown.

Figure 5-10 shows the results of this experiment, with the touchdown velocity of both Tau guides reactive and not reactive to the deck shown. Also plotted is the heave displacement of the deck at the expected time of contact. Manoeuvre duration was set to ten seconds for all of these test simulations, and all use a coupling constant of 0.4. The touchdown velocity for a Tau guide not reactive to the deck varies between 0.5 and 2 m/s with the motion of the deck, and a number of points did not record any result, since the aircraft did not make contact with the deck during the expected manoeuvre duration. The average touchdown velocity of the constant descent rate controller was 1.3 m/s. Conversely, following the Tau guide relative to the deck the aircraft consistently touched down with a velocity lower than 0.5 m/s, with many results between 0.1 and 0.2 m/s and an average of 0.2 m/s. The quiescent period of ship motion between 15 and 25 seconds naturally produces lower touchdown velocities as one might expect. Since another advantage of Tau guidance is that the manoeuvre duration is a defined parameter, it is possible to schedule the touchdown point to coincide with this quiescent period if it can be predicted. The points of higher

touchdown velocity clearly coincide with the points of highest positive heave velocity for the deck. This is likely due to the reversal in heave required by the aircraft to match the deck motion in these situations seen previously. If this is indeed an issue, then it raises the question about how higher sea states may affect performance.

### 5.2.3 Sea State Variation

To examine how sea state affected performance of the Tau guidance system the same start point variation experiment was performed using higher sea state deck motion. Heave motion of the deck is significantly amplified above sea state 4, with maximum displacement increasing from 1 m to nearly 3 m for sea state 5. Figure 5-11 shows the touchdown velocity for twenty test points dispersed through a 30 second sample of sea state 5 deck motion. Again, a Tau descent relative to a fixed datum with no reactivity is also plotted for the sake of comparison.

The non-reactive descents exhibit substantially higher touchdown velocities, above 3 m/s over at least half the points, but are reduced during the quiescent period. The reactive Tau descents touch down result in substantially reduced velocities throughout the deck motion. Velocity is consistently below 1 m/s, though it just passes this mark when the touchdown point coincides with a large heave of the deck. However, contact velocity does not seem to increase significantly during phases where the deck is descending, or at a peak or trough; it only causes issue when the velocity is at a positive maximum. This suggests that the heave control power of the aircraft is the limiting factor in the effectiveness of the Tau guidance system.




Figure 5-11: MQ-8B Fire Scout touchdown velocity for sea state 5 deck landings using a second-order intrinsic Tau guide. Touchdown velocity plotted for Tau manoeuvres reactive to the deck motion, and without reactivity to the deck motion are shown.

This perspective is reinforced by the results from the same experiment repeated at sea state 6, shown in Figure 5-12. Maximum deck displacement is 5 m for this condition, and this leads to contact velocities of between 4 and 7 m/s without reactivity to the deck. This high velocity would likely cause damage to the aircraft. Using the reactive Tau guidance system, the aircraft is still able to achieve touchdown velocities of less than 1 m/s for 16 of the 20 points. On upswings of the deck the contact velocity increases to 2 m/s, again suggesting that the aircraft has reached the limit of its ability to track the deck motion.




Figure 5-12: MQ-8B Fire Scout touchdown velocity for sea state 6 deck landings using a second-order intrinsic Tau guide. Touchdown velocity for Tau manoeuvres reactive to the deck motion, and without reactivity to the deck motion are shown.

To test the hypothesis that the regions of increased touchdown velocity were caused by limitations in the handling qualities of the aircraft, the control power of the MQ-8B was augmented and the same experiment was performed. The control power was artificially increased by doubling the stability derivative for heave force due to collective control,  $Z_{\delta C}$ , in the state-space models input matrix. Figure 5-13 shows the results of this experiment. The peaks in touchdown velocity are reduced significantly for the aircraft with augmented control power. It can therefore be concluded that the manoeuvre should be performable for any sea state if the aircraft has sufficient heave control power.




Figure 5-13: MQ-8B Fire Scout touchdown velocity for sea state 6 deck landings using a second-order intrinsic Tau guide, with and without augmented control power.

Another useful characteristic is the defined temporal trajectory of Tau guides. This guarantees the manoeuvre will take place within a defined duration if tracking performance of the Tau controller is adequate. Figure 5-14 shows the average manoeuvre lengths for all of the simulations covered in Figure 5-10, Figure 5-11 and Figure 5-12, with error bars signifying the range of recorded values. Both data series occur at integer sea states but are separated slightly along the x axis to show the data more clearly. The manoeuvre duration specified to the Tau trajectory generator was 10 seconds for all simulations. For sea state 4, the majority of durations fell within a second of the specified 10 second duration with an average of 9.5 seconds, while the duration for a Tau descent not reactive to the deck took between 10 and 7.5 seconds, with a mean value of 9 seconds. Both mean value and range increase with sea state number. Although a Tau manoeuvre should deliver an exact manoeuvre duration of 10 seconds that will never be the case in reality, due to the physical configuration of the aircraft. Tau is measured from the downward looking camera that measures Tau, so the time-to-contact of the landing gear that will actually hit the deck is actually slightly smaller than the measured value. As a result, true time-to-contact of the landing gear will always slightly lead the measured value, and touchdown will occur

slightly earlier than the specified manoeuvre duration suggests. Similar to touchdown velocity, this manoeuvre duration is influenced by the ability of the aircraft to follow the Tau guide, so mean value will be closer to specified duration for aircraft with greater control power.




Figure 5-14: Manoeuvre durations for 60 simulated MQ-8B deck Tau guided landings across sea states four to six. Note that results are discrete to sea state 4, 5 and 6: lateral displacement of the two data sets is purely for clarity of results.

These degradations in performance due to lack of control power beg the question, is it possible to predict how much control is actually need for a specific manoeuvre?

### 5.3 Tau Manoeuvre Limitations

It is possible to provide some insight on where an aircraft will be unable to perform a particular manoeuvre by leveraging a simple model of rotorcraft heave dynamics from reference [121]. This model relates the acceleration of the aircraft to the heave force stability derivatives from a state-space model of a rotorcraft:

$$\dot{w} - Z_w w = Z_{\delta_C} \delta_C + Z_w w_g \quad (5-4)$$

Where  $w$  and  $\dot{w}$  are the heave velocity of the aircraft and its time derivative,  $Z_w$  is the heave damping stability derivative,  $Z_{\delta_C}$  is the heave force due to collective deflection,  $\delta_C$  is collective deflection and  $w_g$  is vertical wing gust velocity. In this case the wind gust can be neglected for simplicity but would be useful to include in further analysis that includes the effect of ship airwake.

The heave velocity,  $w$ , and acceleration,  $\dot{w}$ , required at any time,  $t$ , in a Tau guided manoeuvre can be found analytically using the equations presented in section 2.2.5. For a second order intrinsic Tau guide applied to the heave axis, they are:

$$w = -D_\tau \frac{3}{k} t^2 (T_G^3 - t^3)^{\frac{1}{k}-1} \quad (5-5)$$

$$\dot{w} = D_\tau \frac{3}{k} t \left[ \left( \frac{3}{k} - 1 \right) t^3 - 2T_G^3 \right] (T_G^3 - t^3)^{\frac{1}{k}-2} \quad (5-6)$$

Where:

$$D_\tau = \frac{x_0}{\frac{2}{T_G^k}} \quad (5-7)$$

Therefore, for the aircraft to be able to follow the guide, the control power must exceed the demands of the manoeuvre:

$$\begin{aligned} D_\tau \frac{3}{k} \left( \left( \frac{3}{k} - 1 \right) t^3 - 2T_G^3 \right) (T_G^3 - t^3)^{\frac{1}{k}-2} \\ + Z_w (t^2 (T_G^3 - t^3)^{\frac{1}{k}-1}) \leq Z_{\delta_c} \delta_{c_{max}} \end{aligned} \quad (5-8)$$

For this investigation, the SH60B Seahawk state-space model described in section 3.2.3 was examined. The model was generated from the nonlinear FLIGHTLAB model trimmed in a hover, so the stability derivatives are only strictly valid for this condition. The stability derivatives will change with heave velocity in reality but for this investigation the assumption was made that this change is small, so the heave stability derivatives are considered to remain constant throughout. The relevant stability derivatives for the SH60B heave model are  $Z_w = -0.0816$ ,  $Z_{\delta_c} = 2.2193$  and the maximum collective deflection from the trim in the hover is  $\delta_{c_{max}} = 5.515$ . Only positive acceleration is examined here, since for negative decelerations the ultimate limiting factor is acceleration due to gravity. For the standard Tau guide used throughout this chapter,  $T_G = 10$ ,  $k = 0.4$  and  $x_0 = 10 \text{ m}$ , and these values are used again here. For this Tau manoeuvre the left and right sides of equation (5-8) are evaluated separately, with the collective control power terms on the right-hand side treated as a threshold that, if breached, will lead to tracking failure during the



manoeuvre. This threshold and accelerative demands of the Tau guide are plotted in Figure 5-15.





Figure 5-15: Heave acceleration required for an aircraft to perform an ideal second-order Tau guide vertical landing, and the maximum available heave acceleration for the SH60B Seahawk in a hover condition.

Clearly for this manoeuvre, the aircraft does not approach the limits of its capability at any point and should easily be able to complete it. However, if the same manoeuvre is performed over a heaving ship deck, the motion of the ship must also be taken into account in the model as follows:


$$(\dot{w} + \dot{w}_D) - Z_w(w + w_D) \leq Z_{\delta_C} \delta_C \quad (5-9)$$

Using the ship deck motion profiles detailed in section 3.5.1.1, the velocity and acceleration profiles were found by taking the first- and second-order gradients of the heave displacement, and added to the Tau guide acceleration demands for several sea states. Since it isn't immediately clear where the maximums of this function lie, they were calculated across the Tau manoeuvre duration and the whole envelope of the ship deck motion profile. The accelerative demands for a sea state four landing are plotted in Figure 5-16.



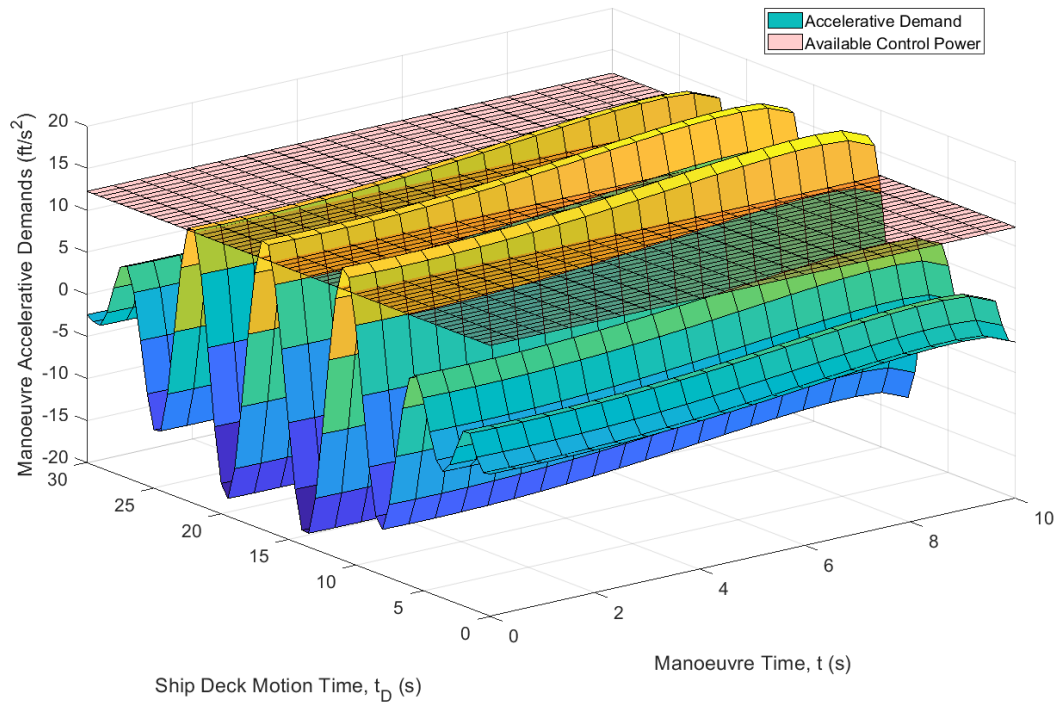
*Figure 5-16: Heave acceleration required for an aircraft to perform an ideal second-order Tau guide vertical landing on heaving ship deck in sea state 4, and the maximum available heave acceleration for the SH60B Seahawk in a hover condition.*

The acceleration still does not approach the available control power and there is a significant margin in this sea state 4 example, but the motion of the deck is clearly reflected in the acceleration required throughout the manoeuvre.



*Figure 5-17: Heave acceleration required for an aircraft to perform an ideal second-order Tau guide vertical landing on heaving ship deck in sea state 5, and the maximum available heave acceleration for the SH60B Seahawk in a hover condition*

However, for a sea state 5 landing as shown in Figure 5-17, the accelerative demand reaches 11 ft/s<sup>2</sup>, very close to breaking the threshold of maximum available control power of 12.2 ft/s<sup>2</sup>. One step further to sea state 6 and this barrier is broken, as shown in Figure 5-18.



*Figure 5-18: Heave acceleration required for an aircraft to perform an ideal second-order Tau guide vertical landing on heaving ship deck in sea state 6, and the maximum available heave acceleration for the SH60B Seahawk in a hover condition*

In this scenario, the Seahawk helicopter is not able to produce enough heave acceleration to match the demand of the Tau guide and deck motion combined in certain regions. Therefore, the aircraft is not be able to track the Tau guide at these points. If these regions occur during the terminal phase of the manoeuvre as the aircraft nears the deck, it will likely hit the deck with a higher velocity than desired and potentially damage the aircraft. On the other hand, if the aircraft can touch down between these points, there is still a large enough margin to complete the manoeuvre with a low touchdown velocity. These findings correlate well with the results of the sea state variation experiments performed in the previous section, where touchdown velocity was increased above the expected value when the touchdown point coincided with the deck heaving upwards in high sea state.

To test the hypothesis that this rotorcraft dynamic models can be used to predict the performance limits of an aircraft, the SH60B was investigated using the same start point variation experiment applied to the MQ-8B earlier. The two other aircraft models, the R-Max and T-Rex, were also used to demonstrate a range of aircraft using

Tau guidance. The Fire Scout is included again for the sake of comparison. Figure 5-19 shows each of the four aircraft undertaking the same second-order Tau guide landing on a ship deck in sea state 4, with ten different start points evaluated. The deck motion at the expected end point is also plotted.




Figure 5-19: Touchdown velocity for sea state 4 deck landings using a second-order intrinsic Tau guide for four different aircraft

All four aircraft are able to follow the Tau guide and land with a touchdown velocity of less than 1 m/s at all points through the deck motion profile. No parameters were changed in the Tau controller, which demonstrates that this control method is essentially platform independent; it is equally applicable to all four aircraft and can deliver similar performance if the handling qualities of the aircraft are sufficient.

The R-Max, the aircraft with the worst heave control power, exhibits slightly higher touchdown velocities throughout the deck motion profile, even at this relatively low sea state. On the other hand, the SH60B exhibits low touchdown velocities throughout, but it does increase slightly through points of deck upswing. Figure 5-21 shows the same analysis for sea state 6 deck motion.




Figure 5-20: Touchdown velocity for sea state 6 deck landings using a second-order intrinsic Tau guide for four different aircraft

Again, the R-Max experienced high touchdown velocities due to its low vertical control power. As predicted during the previous analysis, the SH60B is also not able to track the deck motion through areas of high deck upswing, and lands heavily. As a result, it can be concluded that the modelling technique proposed above can be used for indication of the performance limits of rotorcraft performing Tau deck landings.

The order of the aircraft from highest to lowest touchdown velocity is not consistent at all test points, especially when the point of touchdown does not coincide with a deck upswing in the time history. For example, the best performer out of the MQ8 and T-Rex switches at each of the first 4 test points, depending on how the deck is moving. This suggests that there are more factors in the determination of Tau landing performance than purely heave control power. This variability may be due to a range of handling qualities considerations or control system tuning problems and warrants further investigation. However, since the touchdown velocities are low at these points, these factors do not seem to be affecting touchdown velocity performance significantly, and as a result are not further investigated in this thesis.

## 5.4 Optical Tau Guided Frigate Deck Landings Results

The Simulink 3D Animation toolbox was used to implement the VR frigate deck landing scene detailed in section 3.4.1 into the Simulink aircraft model, and the resulting video stream was processed using the direct gradient method for Tau estimation. The brightness change extension proposed in section 4.4 was used. The resulting Tau estimate was fed back into the Tau guidance system to evaluate how well this Tau estimate can be utilised for the landing manoeuvre.

### 5.4.1 Sea State 1

The system was initially evaluated in sea state 1, which corresponds to a flat, calm ocean where no heave displacement of the deck occurs at all. The Tau reference trajectory, measured Tau value, and the ground truth value are shown in Figure 5-21 for the MQ-8B model following a second-order intrinsic Tau guide.

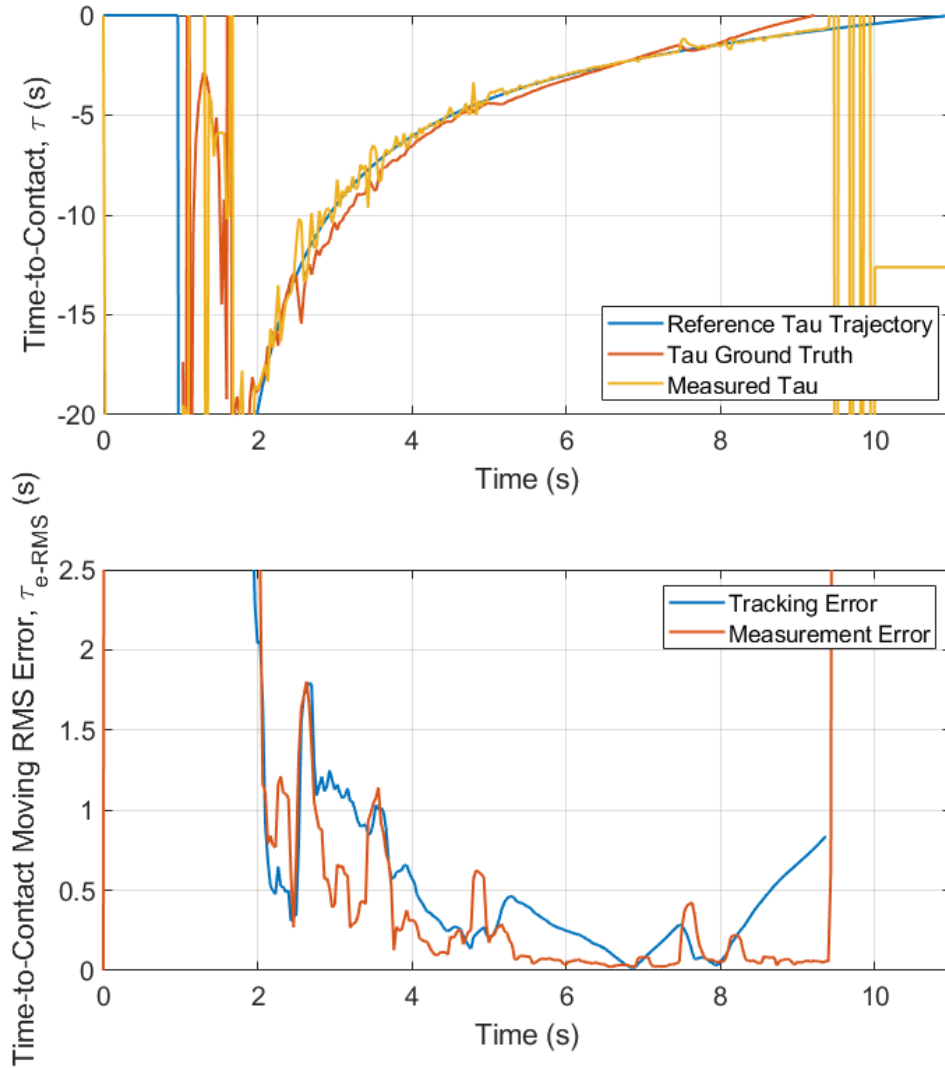


Figure 5-21: Tau quantities through an optical Tau guided ship deck landing by the MQ-8B in sea state 1. Time-to-contact with the deck measured by the direct gradient method, along with the ground truth time history of Tau and the reference trajectory that the controller is attempting to track are shown in the top plot. The RMS tracking error of the controller between the ground truth and reference trajectory, and measurement error between the measured and ground truth values of Tau.

The measured value of Tau is the one used by the controller to generate the error signal for control inputs, equivalent to the tracking error plotted above. The tracking performance is very good between 4 and 9 seconds, with error below 0.5 seconds consistently, excepting small discontinuities at the points where the subsampling level switches just before the 5 and 8 second marks. There are larger errors in the first 3 seconds of the manoeuvre, but this is unsurprising since the value of the



reference trajectory is high and changes quickly at this point. It is also not problematic since the aircraft is a long way from the deck at this point. The error between the measured value of Tau and the ground true value is also low between 4 and 8 seconds but begins to increase rapidly after 8 seconds as the estimated value becomes larger than the true value. This should result in an acceleration toward the deck, and this is observed by the heave displacement and relative velocity plotted in Figure 5-22; instead of decelerating at the end of the manoeuvre, the measurement error causes the aircraft to remain at a high velocity and hit the deck at 1.3 m/s. The high error at the beginning of the manoeuvre also seems to cause a large spike in velocity as the aircraft begins to descend.





Figure 5-22: Heave dynamics for MQ-8B for a sea state 1, second-order intrinsic Tau guide landing using optical Tau estimates

Analysis in chapter 4 suggested that there would be small errors as the camera approached the deck, it was not clear that it would have this great an effect, and it does not seem possible to mitigate for these problems simply with the available

techniques. As a result, it must be concluded that optical Tau perception is not suitable for the terminal phase of this deck landing manoeuvre. The author proposes that in this section, performance could be augmented with an active ranging sensor for improved performance. Since this divergence in measurement occurs in the last 2 m of descent, an ultrasonic sensor could be used effectively for this situation. This solution was implemented in simulation, with Tau measured from the optical camera feed through the first 7 seconds of the Tau manoeuvre, and then switched to a spatial estimate calculated using the aircraft height above the deck and relative velocity. Figure 5-23 shows the same descent carried out with this switch to a spatial Tau estimate for the final 3 seconds of the descent.




*Figure 5-23: Heave dynamics for MQ-8B for a sea state 1, second-order intrinsic Tau guide landing using optical Tau estimates with spatial augmentation. The point where spatial augmentation of the Tau estimate is activated is marked with a red dotted line at 8 seconds*

This approach prevents the divergence in Tau measurement error from commanding a high velocity toward the deck in the terminal phase and allows for smooth touchdown at low velocity.

#### 5.4.2 Higher Sea States

Optical Tau guided landings were also performed with higher sea state motion. Figure 5-24 shows the same second-order Tau guided descent by the Fire Scout onto a deck moving with sea state 3 motion. Again, the augmentation with spatial data is used for the final 3 seconds of the manoeuvre. The aircraft is again able to make a soft landing on the deck with minimal relative velocity.



*Figure 5-24: Heave dynamics for MQ-8B for a sea state 3, second-order intrinsic Tau guide landing using optical Tau estimates with spatial augmentation*

The deck motion for sea state 3 is still relatively benign, so the results for the same experiment performed in sea state 5 are plotted in Figure 5-25. This descent gives a much better of how optically perceived Tau can still be a powerful tool in this scenario, as the aircraft clearly slows its descent to avoid the peak in deck position 5

seconds into the manoeuvre. This is accomplished using only optical Tau as a control variable, and occurs at a range where an ultrasonic ranging sensor would not be accurate, though it is still necessary to augment the optical Tau estimate with a spatial one through the final moments of the landing.




Figure 5-25: Heave dynamics for MQ-8B for a sea state 5, second-order intrinsic Tau guide landing using optical Tau estimates with spatial augmentation

This descent again exhibits the large transient velocity after manoeuvre initiation, which is undesirable. Since third-order intrinsic Tau guides should start with zero acceleration and develop velocity more slowly, they may offer a mechanism for reducing this velocity spike. Figure 5-26 shows the heave dynamics for an optically guided third-order Tau descent in sea state 5. Spatial Tau augmentation is again used over the final 3 seconds of the manoeuvre. Unfortunately, following a third-order intrinsic Tau guide does not reduce the velocity spike that occurs during manoeuvre initiation. This seems to be because the reference Tau trajectory is similar to a second-order guide in that they both produce a reference Tau of negative infinity at the start of the reference trajectory, which despite saturation functions being

included, is still partially transmitted to the controller. If the deck was stationary then the measured value of Tau would also be infinite and this would not be a problem, but since the deck heaves relative to the aircraft constantly, a non-infinite value of Tau is measured and this leads to a large error signal passed to the controller. The velocity spike does not result in a large displacement, so should not pose a large risk to the safety of the aircraft.




Figure 5-26: Heave dynamics for MQ-8B for a sea state 5, third-order intrinsic Tau guide landing using optical Tau estimates with spatial augmentation

## 5.5 Fixed-Wing Tau Guided Landings

### 5.5.1 Tau Guidance System Implementation

The Tau guidance system detailed in chapter 5 has also been applied to the simulation model of the 3DR Aero, detailed in section 3.2.2. The technical review in chapter 2 only yielded one example of Tau guidance used for the autonomous control of fixed-wing aircraft, which was Ridgeway's Tau pilot model applied to a simulation of a Grob Tutor aircraft in reference [21]. This example calculated Tau from spatial variable to complete a range of manoeuvres. No examples at all have been found of fixed-wing aircraft using optical Tau as a guidance variable. It was demonstrated in chapter 4 that it is possible to generate Tau estimates for a camera attached to a fixed-wing aircraft, so it should be possible to implement this.

A Tau guidance system using the nonlinear ratio control law was implemented on the linear state-space model of the 3DR Aero in Matlab Simulink for testing purposes. The control loops differ from the ones applied to the rotorcraft models described in section 5.1, since translational rate command is not a suitable command type for a fixed-wing aircraft. Instead, control loops were added for heading- and airspeed-hold functions. The full flight controller architecture is given in Appendix C-3, and the matrix of control gains used in the flight controller is also given in Appendix C-4. The gain values were tuned manually for simplicity and allowed the aircraft to complete the manoeuvre without significant excursions from the desired trajectory. As for the rotorcraft flight controller, the system should be better tuned using more formal methods in future to improve performance.

As before, a constant heave rate controller was added in parallel to the Tau guidance system so the user can switch between the two systems. In this case, the output of the Tau controller is connected directly to the elevator control channel to allow direct alteration of the flight path. Both a first- and second-order Tau guide reference trajectory were included. The first-order system is set up to descend at a constant rate while monitoring time-to-contact and activate a constant rate of change of Tau deceleration when  $\tau < 5$  s. The desert runway VR world detailed in section 3.4.2 was

also integrated into the simulation model to provide a video stream to generate optical Tau estimates.

### 5.5.2 Flat Runway Tau Landings

The results for a first-order Tau guided landing are shown in Figure 5-27. The Aero descends at a constant rate of 2.5 m/s from a starting altitude of 50 m. Tau is measured spatially in this case. When time-to-contact with the ground plane breaks the 5 s threshold, the Tau controller is activated, and the aircraft decelerates to touch down with low heave velocity.




Figure 5-27: Fixed-wing runway landing of the 3DR Aero following first-order intrinsic Tau guidance using spatially computed time-to-contact estimates. Aircraft descends from 50m at a constant rate and initiates a Tau flare manoeuvre when time-to-contact passes a threshold value of 5 seconds

For this landing Tau was estimated using the direct gradient method applied to the video stream of the VR runway, though the estimate was not used in the guidance system. The estimate through the manoeuvre is plotted in Figure 5-28. Clearly the measurement is highly noisy and tracking error is poor. No estimate at all is available during the actual deceleration phase of the manoeuvre because the virtual camera begins to clip through the terrain and displays no image, despite the camera is still being 5 to 10 m above the ground plane. It is not clear why this problem occurred and unfortunately troubleshooting failed to identify the problem. New tools for modelling visual scenes may need to be integrated into the system to provide a better environment for testing of optical Tau guided fixed-wing aircraft. However, the benefits of Tau guidance in this context can still be evaluated using spatial estimates of time-to-contact.




Figure 5-28: Time-to-contact estimate from the 3DR Aero runway landing described above, where the aircraft follows a first-order Tau guide triggered at a threshold time-to-contact of 5 seconds. Tau is calculated using the extended direct gradient method.

### 5.5.3 Landings at Unprepared Sites

The benefits of Tau guidance for fixed-wing aircraft landings are similar to those discussed for the ship deck landings in chapter 5; the aircraft is able to respond to



changing environmental dynamics with ease, but it is admittedly far less likely that a fixed-wing UAS will be required to land on a moving object. However, the high forward speed of the aircraft can produce the effect of a moving ground plane. If the elevation of the terrain changes as the aircraft moves over it, this appears as if the ground is moving towards or away from a downward looking camera. Such terrain variations may be present if an aircraft must operate from an unprepared runway, as is the case with many SUAS. Several experiments have been performed to assess how a Tau guided fixed-wing aircraft with perform landing on uneven terrain. Figure 5-29 shows the 3DR Aero descending towards a runway that constantly slopes upward toward the aircraft at a 5-degree angle.

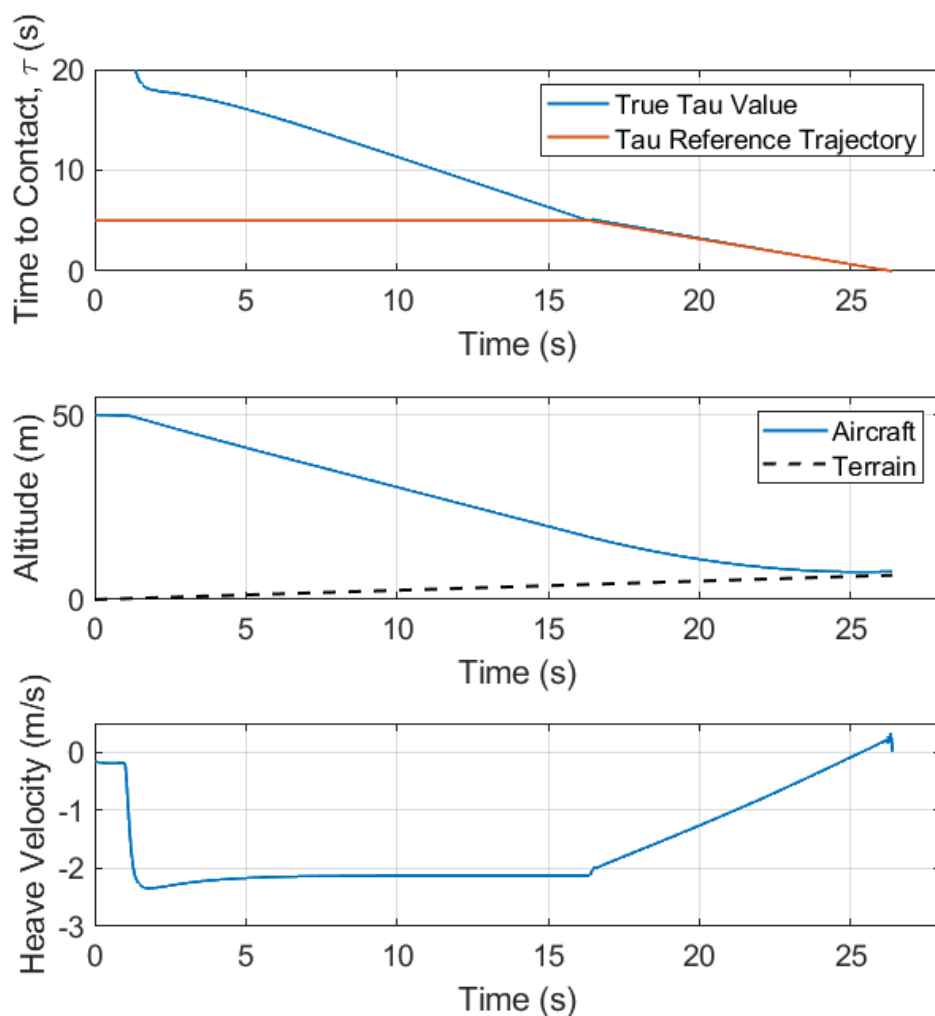


Figure 5-29: Fixed-wing landing of the 3DR Aero following first-order intrinsic Tau guidance using spatially computed time-to-contact estimates. Aircraft descends over terrain that slopes upward at 5

degrees in the direction of travel. Aircraft descends from 50m at a constant rate and initiates a Tau flare manoeuvre when time-to-contact passes a threshold value of 5 seconds

As observed in the ship deck landing analysis, the aircraft is able to respond to the changing demands of the environment and still perform a soft landing with no modifications to the time-to-contact trajectory necessary. The same is true if the terrain does not vary linearly. Figure 5-30 shows a 3DR Aero descent towards terrain that varies over a 10 m range in a sinusoidal pattern. Variations in the spatial estimate of time-to-contact are visible while the aircraft descends at a constant rate, but the controller tracks the reference trajectory well once it is activated.

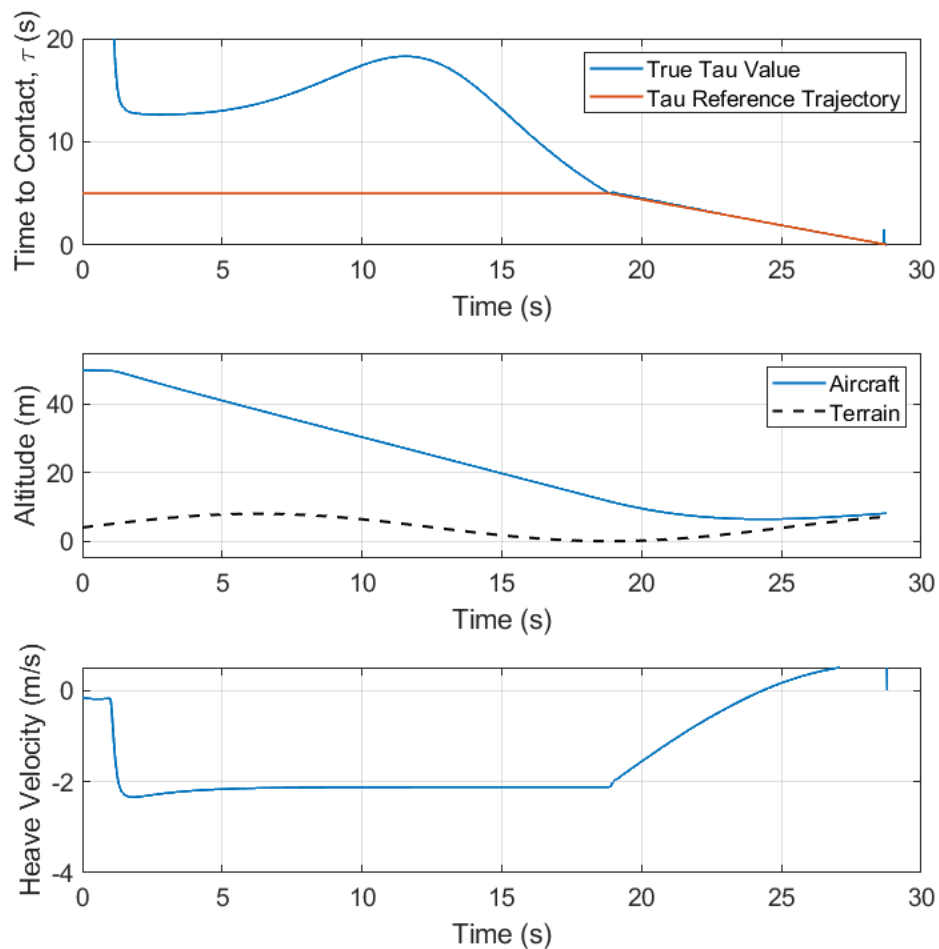


Figure 5-30: Fixed-wing landing of the 3DR Aero following first-order intrinsic Tau guidance using spatially computed time-to-contact estimates. Aircraft descends over unprepared ground where terrain elevation varies sinusoidally. Aircraft descends from 50m at a constant rate and initiates a Tau flare manoeuvre when time-to-contact passes a threshold value of 5 seconds

These simple experiments demonstrate that optical Tau guidance is also useful when applied to fixed-wing landings, especially for landings at unprepared sites with potential terrain variations. However, further work is still required to successfully implement Tau perception methods on fixed-wing platforms.

## 6 Chapter 6 – Conclusions

The research presented in this thesis has sought to better illuminate how the guidance methods that stem from ecological Tau theory can be applied to unmanned aerial systems. Tau theory provides powerful mechanisms to explain the natural behaviour of animals and humans, both when controlling their own bodies and machines, by leveraging the time-to-contact of an observer with an obstacle or target. Tau has been shown to follow defined patterns during many natural actions and is theorised to be directly perceived from an observer's environment.

As the market for UAS continues to grow, technological solutions must be found to answer the many problems this fast-growing industry faces. Analysis of accident data has shown that unreliable communication links, pilot errors and a limited ability to detect and avoid hazards are primary causes of many UAS accidents. The author proposes that the application of Tau theory to guide UAS may aid in tackling some of these problems. In an effort to establish how useful Tau guidance can be, and what obstacles must be overcome for its successful use, several distinct areas have been examined. This chapter will summarise the findings of each element of this thesis.

### 6.1 Conclusions of the Research

#### 6.1.1 Tau Perception

The perception of time-to-contact is not a new problem, and many solutions already exist for this task, but it is still a fundamental problem to the application of Tau guidance with no universally accepted answer. Many systems that have implemented Tau guidance thus far have used spatial information from GPS-INS systems to do so, which compromises the key benefit of Tau theory, namely the implicit situational awareness gained from optical sensing of Tau. Active ranging techniques can be simply applied for Tau perception, but literature shows that the effective ranges and characteristics of available sensors are not currently adequate in a great enough range of conditions to be useful. Reliable ranging sensors also tend to be expensive and heavy for use on small unmanned aircraft. Optical sensors are cheap, viable at many ranges, and offer an obvious analogue to the natural

mechanism for Tau perception. Chapter 4 of this thesis examined several optical techniques for Tau perception to ascertain which best suit the needs of UAS.

Dimension tracking methods are simple, intuitive and offer accurate estimates of Tau at varied ranges. However, it relies heavily on the ability to detect and track features within an image. Feature content can be poor in some images, and translational motion parallel to the ground plane leads to loss of features very quickly. Without sufficient robust feature content, no Tau estimate can be made at all, which heavily limits the usefulness of dimension tracking methods. They also do not deal well with some of the expected complexities of an aerial observer, namely six degree of freedom motion, and landing on inclined planes, since both distort the relationship between apparent size of an object on the image plane and its true size.

Optical flow divergence methods should in theory be able to provide highly accurate estimates of Tau, since the optical flow field characterises relative motion very well. However, the obstacles involved in estimating the optical flow field itself make Tau perception through this method very difficult. The most popular methods of flow field estimation again depend on robust feature detection. Though widely implemented in common image processing packages, optical flow implementations are somewhat less than intuitive, and all experiments have produced heavily flawed optical flow fields even with simple camera dynamics.

As a result, direct gradient methods have been identified as the most effective method for optically sensing time-to-contact at current. Gradient methods leverage elements of both dimension tracking and optical flow to calculate Tau directly. The method has been implemented through both a custom Matlab class of objects, and a Simulink system. This method only requires knowledge of some intrinsic camera parameters, and the gradients of the frames in a video sequence. However, due to the reliance of the method on the constant brightness assumption, the effectiveness of the method is heavily dictated by the relationship between the parameters of the camera, and the dynamics of the camera.

The frame rate, resolution (and hence effective pixel size) of a video camera used for gradient method estimation of Tau will affect the accuracy of results heavily. It has

been found that for cases where translational motion parallel to the ground plane is small, the residual of the constant brightness assumption when populated with a simple approximation of an optical flow field is a good indicator of how accurate a Tau measurement will be. The error in Tau measurement was found to increase linearly with the mean of the constant brightness equation residuals, with 1 s of error in Tau corresponding to a residual of 30. To maintain residuals below this threshold, the properties of the camera can be manipulated using subsampling, a method of reducing effective resolution through block processing each video frame. However, for cases with high translational motion (such as a camera attached to a fixed-wing aircraft), this relationship seems to break down, and the residuals of the constant brightness assumption did not provide an indicative measure of Tau measurement accuracy. Subsampling is still useful for maintaining effectiveness and scheduling increased levels of subsampling with a reliable scaler can produce accurate results.

Another drawback of the gradient method is the relatively large computational loads for processing each frame. This factor can be mitigated by using lower resolution images or frame rates, though this must also be balanced with the most effective image sizing for Tau perception as well. It was found that an HVGA (320x480 pixels) camera resolution recording at 30 frames per second presented a good compromise between all factors, and through subsampling could be effective throughout a range of manoeuvres with low translational motion. For cases with high relative motion, increasing frame capture rate to 60 fps was found to significantly improve Tau perception performance.

The most powerful element of the direct gradient method is its extensibility. Several extensions to the method have already been proposed for the method, that allow it to account for observer motion in 6 degrees of freedom, with relative slope between observer and target. A further extension is proposed by the author that augments the underlying constant brightness assumption with the ability to account for global brightness changes in the image and maintain accurate Tau estimates. This new method for Tau perception was tested using several image brightness profiles that change over time, and was found to mitigate deviations caused by the brightness change in the base gradient method well.

### 6.1.2 Tau Guidance of Rotary UAS

Tau guidance of rotary UAS was tested extensively in the context of frigate deck landings in various sea states. It was found that the nonlinear ratio control law originally proposed by Kendoul is most effective for Tau control when tested against proportional and inverse control laws. The frigate deck landing illustrates the benefits of optical Tau guidance in an obvious way. Measurement of Tau naturally captures the motion of the deck and the aircraft adapts to it seamlessly due to the temporal nature of the reference Tau trajectory. The ability to fit manoeuvres into temporal gaps that Tau guidance allows is also very useful, as it allows manoeuvres to be scheduled for quiescent periods in deck motion if they can be predicted. In comparison to other commonly implemented techniques for UAS landings, the Tau guidance system consistently delivered very low touchdown velocities across the whole envelope of deck motion in multiple sea states.

Performance of the Tau guidance system was found to only be limited by the control power of the aircraft when Tau estimates are sufficiently accurate. A simple model for prediction of the performance limits with respect to sea state and Tau manoeuvre aggressiveness was developed to predict when aircraft would hit the deck with a higher than expected velocity. The model was applied to the SH60B Seahawk, and it correctly predicted that the touchdown velocity of the aircraft would increase above nominal in sea state 6 for the Tau trajectory defined by the MTE. Several other aircraft were also tested with the same Tau guidance system to demonstrate that it was equally applicable to any platform.

Optical measurements of Tau were then used in the loop with the simulation model to control the aircraft solely with optical Tau. Tracking performance through the first 4 seconds of the manoeuvre was poor due to a large velocity spike brought about the singularity in a second-order intrinsic Tau guide at  $t = 0$ . Between 4 and 8 seconds the aircraft tracked the desired profile closely, with an RMS error of less than 0.5 s. However, in the last two seconds, the measured value of Tau diverged from the true value and caused the aircraft to accelerate toward the deck and make contact at 1.3 m/s. While this is not excessive, it would be amplified in higher sea states and is substantially larger than the results obtained with better Tau estimates. It may be

possible to improve optical Tau measurements in the last 2 m above the deck where this issue occurred, but it is recommended that Tau measures are augmented with an ultrasonic sensor in this area. Ultrasonic sensors are effective at this range, are lightweight and cheap, so would not negatively impact a UAS in this situation.

### 6.1.3 Tau Guidance of Fixed-Wing UAS

A similar Tau guidance system was applied to a simulation of a small fixed-wing UAS, the 3DR Aero. Tau guidance provides similar benefits to fixed wing aircraft as it does to RUAS; it enables reactivity to environmental dynamics and guarantees soft landing if Tau measurements are reasonably accurate. The level of analysis of fixed-wing aircraft fell short of the detail applied to rotary aircraft, due in part to technical failings of the analysis tools. However, it was demonstrated that fixed wing aircraft following Tau guides are able to react to changing terrain at unprepared landing sites. It was also demonstrated that it is possible to estimate Tau from fixed-wing platforms, but accuracy was not sufficient to enable control of the UAS from it. Since fixed-wing aircraft are more likely to land in static environments, spatial GPS-INS state information could potentially be used to implement Tau guidance without much performance degradation. Further work will be required to tune and improve optical Tau perception from fixed-wing platform.

### 6.1.4 Overall Conclusions

Referring back to the aims and objectives set out in section 1.5, the above conclusions can be assessed against them. Project objectives 1, 2 and 3 have been completed, though implementation of optical Tau guidance on fixed-wing platforms remains elusive, preventing true completion of objective 4. Addressing the three project aims:

1. It has been demonstrated that Tau has the potential to address some of the problems facing UAS: Tau guidance improves reactivity to external obstacles and hazards, reduces reliance on external signals such as GPS, and improves performance in landing manoeuvres, even challenging ship deck landings.
2. The direct gradient method has been identified as an effective method for Tau perception, and an improvement to it has been proposed that further expands capability.



3. Optical Tau guidance extends the same benefits to fixed-wing flight as it does to rotary-wing flight, but the implementation of the former has proved more difficult and warrants further investigation.

## 6.2 Potential Further Work

1. Further investigation into the effects that camera properties and camera motion have on Tau estimate accuracy is needed to allow the direct gradient method to run at maximum accuracy consistently. Currently it is not always clear whether the algorithm is providing the best possible measurement in all situations, so better characterisation of how all factors affect the output is needed. This is especially true for the fixed-wing implementation of the Tau sensor where performance is still poor.
2. More rigorous testing of the Tau sensor in a wider range of environments is also needed to build confidence in its effectiveness. Further testing of the new extension to the direct gradient method that accounts for brightness change is also needed with a great variety of brightness change models, such as partial occlusion of the image.
3. More simulation testing of the Tau guidance system is needed with higher fidelity simulation models to identify potential problems ahead of hardware testing. Several of the linear state-space models used for the majority of testing were produced from nonlinear FLIGHTLAB models which could be used for higher fidelity testing. These models would capture more physics that may impact the effectiveness of the Tau guidance system, especially close to the ground. More detailed simulations of the maritime environment that include the atmospheric disturbances caused by the air wake of the ship superstructure should also be investigated for progression of deck landing operations.
4. Better tuning of control systems is required to improve both on-axis Tau control and other off-axis controllers, since all control gains were tuned simply through trial and error to this point. In future all gains should be tuned more rigorously using more formal techniques to reduce unwanted oscillations, reduce settling time, and improve disturbance rejection ability.

5. Hardware testing of the Tau sensor and Tau guidance systems are needed to confirm both systems are viable for full implementation. Optimisation of the extended direct gradient method is required for this step to reduce computation time and allow for a faster execution rate.

## Appendix A – Aircraft Linear State-Space Models

This appendix details the state and control matrices for the linear state-space models described in section 3.2. All of the models are time-invariant state-space models of the form in equations A-1 and A-2, where  $x$  is the state vector,  $A$  is the state matrix,  $B$  is the control matrix,  $u$  is the control vector,  $y$  is the output vector,  $C$  is the output matrix which is a 9-by-9 identity matrix,  $D$  is the feedforward matrix which is a 9-by-4 zero matrix and  $x_{init}$  is the initial state of the aircraft.

$$\dot{x}(t) = Ax(t) + Bu(t) \quad (A-1)$$

$$y(t) = Cx(t) + Du(t) + x_{init} \quad (A-2)$$

This model is shown as a block diagram in Figure 6-1 and was implemented in Simulink in this way. A saturation is applied to the control input that limits its range,  $u_{range}$ , to the one specified by the original FLIGHTLAB model each state-space model was derived from. The initial control positions,  $u_{init}$ , are also included here by altering the saturation range accordingly.

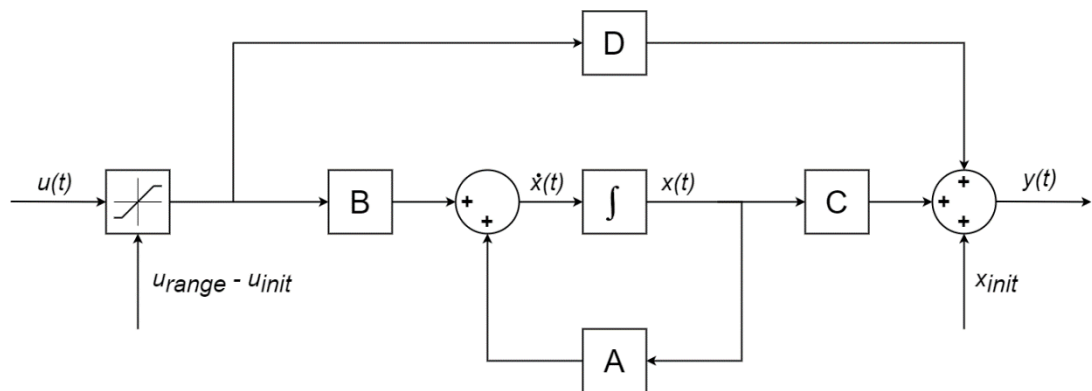


Figure 6-1: Block diagram of state-space model implementation

## Appendix A.1 – Sikorsky SH-60B Seahawk Model

The FLIGHTLAB model that the state-space model was derived from was trimmed in a hover at 100 feet altitude with no ambient wind. The state matrix, state vector, control matrix, control vector, initial state vector, control range matrix and initial control vector for the Sikorsky SH-60B Seahawk model are:

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 & -0.0024 & 0.0504 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9989 & 0.0467 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.0467 & 1.0002 \\ 0 & -28.6125 & 0 & -0.0170 & 0.0054 & 0.0033 & -1.5513 & 3.0874 & -0.0771 \\ 32.2 & -0.3086 & 0 & 0.0011 & -0.0301 & 0.0020 & -1.9423 & -1.7114 & 0.3063 \\ 1.501 & -1.3408 & 0 & 0.0052 & -0.1004 & -0.0816 & -1.3164 & 2.2791 & -0.9121 \\ 0 & -0.4390 & 0 & 0.0010 & -0.0239 & 0.0009 & -4.5792 & -1.5790 & -0.0513 \\ 0 & -0.6856 & 0 & 0.0028 & 0.0024 & 0.0005 & 0.1961 & -0.9972 & -0.0262 \\ 0 & -0.0229 & 0 & 0 & 0.0037 & -0.0010 & -0.1963 & -0.1387 & -0.1726 \end{bmatrix} \quad (A-3)$$

$$x = [\phi \ \theta \ \psi \ u \ v \ w \ p \ q \ r]^T \quad (A-4)$$

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0.0067 & -0.1651 & -0.0937 & 0.0943 \\ 0.0918 & 0.0179 & 0.0292 & -0.0723 \\ 0.0044 & -0.0130 & 2.2193 & -0.7718 \\ 0.1155 & 0.0204 & -0.0030 & -0.0319 \\ 0.0005 & 0.0319 & -0.0263 & 0.0048 \\ 0.0060 & 0.0011 & -0.0032 & 0.0228 \end{bmatrix} \quad (A-5)$$

$$u = \begin{bmatrix} \delta_{lat} \\ \delta_{lon} \\ \delta_{col} \\ \delta_{ped} \end{bmatrix} \quad (A-6)$$

$$x_{init} = [-0.0466 \ 0.0503 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]^T \quad (A-7)$$

$$u_{range} = \begin{bmatrix} 0 & 10 \\ 0 & 10 \\ 0 & 10 \\ 0 & 5.38 \end{bmatrix} \quad (A-8)$$

$$u_{init} = \begin{bmatrix} 5.0886 \\ 5.6640 \\ 4.4850 \\ 1.5693 \end{bmatrix} \quad (A-9)$$

All angles are measured in radians, angular rates are measured in radians per second and linear rates are measured in feet per second.

## Appendix A.2 – Northrop Grumman MQ-8B Fire Scout Model

The FLIGHTLAB model that the state-space model was derived from was trimmed in a hover at 100 feet altitude with no ambient wind. The state matrix, state vector, control matrix, control vector, initial state vector, control range matrix and initial control vector for the Northrop Grumman MQ-8B Fire Scout are:

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & -32.2 & 0 & -0.0200 & 0.0037 & 0.0089 & 0.0516 & 2.3968 & -0.0629 \\ 32.2 & 0 & 0 & -0.0014 & -0.0377 & 0.0066 & -2.2261 & -0.3187 & 0.2369 \\ 0 & 0.0024 & 0 & 0.0160 & 0.0031 & -0.3982 & -0.0327 & 0.2366 & 1.4921 \\ 0 & 0 & 0 & -0.0150 & -0.0999 & 0.0040 & -23.8176 & 2.2019 & -0.2229 \\ 0 & -0.0001 & 0 & 0.0102 & 0.0006 & 0.0065 & -0.8831 & -3.2247 & -0.0934 \\ 0 & 0 & 0 & -0.0006 & 0.0118 & -0.0016 & -0.1714 & -0.1178 & -0.2659 \end{bmatrix} \quad (A-10)$$

$$x = [\phi \ \theta \ \psi \ u \ v \ w \ p \ q \ r]^T \quad (A-11)$$

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -0.0335 & -2.0765 & 0.4312 & 0.0503 \\ 1.0939 & 0.1288 & -0.0400 & -1.8699 \\ -0.0050 & -0.0679 & -14.3644 & 0.6881 \\ 5.4903 & -0.2828 & -0.1798 & -1.9850 \\ 0.1627 & 1.3543 & 0.2273 & 0.4957 \\ 0.0543 & -0.0334 & 0.7158 & 1.2408 \end{bmatrix} \quad (A-12)$$

$$u = \begin{bmatrix} \delta_{lat} \\ \delta_{lon} \\ \delta_{col} \\ \delta_{ped} \end{bmatrix} \quad (A-13)$$

$$x_{init} = [-0.0496 \ 0.0071 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]^T \quad (A-14)$$

$$u_{range} = \begin{bmatrix} -5 & 5 \\ -5 & 5 \\ 0 & 10 \\ -2.5 & 2.5 \end{bmatrix} \quad (A-15)$$

$$u_{init} = \begin{bmatrix} -0.3790 \\ 0.4960 \\ 5.1910 \\ 0.5660 \end{bmatrix} \quad (A-16)$$

All angles are measured in radians, angular rates are measured in radians per second and linear rates are measured in feet per second.

### Appendix A.3 – 3DR Aero Model

The FLIGHTLAB model that the state-space model was derived from was trimmed in straight and level flight at 30 knots airspeed at 100 feet altitude with no ambient wind. The state matrix, state vector, control matrix, control vector, initial state vector, control range matrix and initial control vector for the 3DR Aero are:

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & -0.0053 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0.001 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.001 & 1 \\ 0 & -32.2 & 0 & -0.2533 & -0.1386 & 0.3835 & 0.0025 & 0.2301 & 0.0081 \\ 32.2 & -0.0002 & 0 & 0.0014 & -1.4159 & 0.0015 & -2.2714 & -0.0088 & -47.017 \\ 0.032 & 0.1695 & 0 & -1.2111 & 0.0028 & -16.831 & 0.1066 & 44.5761 & -0.0187 \\ 0 & 0 & 0 & 0.0220 & -1.7098 & 0.0937 & -34.234 & -0.0024 & 3.7364 \\ 0 & 0 & 0 & 0.0201 & -0.0091 & -1.8986 & -0.0057 & -3.5441 & 0.0035 \\ 0 & 0 & 0 & 0 & 0.7385 & 0.0045 & -0.0118 & 0.0014 & -2.0435 \end{bmatrix} \quad (A-17)$$

$$x = [\phi \ \theta \ \psi \ u \ v \ w \ p \ q \ r]^T \quad (A-18)$$

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0.0017 & 0.0057 & 0.0503 & 0.0023 \\ 0 & 0.0260 & 0 & -0.3862 \\ 0.4259 & -1.0662 & 0 & 0 \\ 0 & 0.4238 & -0.0022 & -0.0723 \\ 0.3144 & -0.0521 & -0.0053 & 0.0004 \\ 0 & 0.0040 & 0 & 0.2500 \end{bmatrix} \quad (A-19)$$

$$u = \begin{bmatrix} \delta_{lon} \\ \delta_{lat} \\ \delta_{thr} \\ \delta_{ped} \end{bmatrix} \quad (A-20)$$

$$x_{init} = [-0.001 \ -0.0053 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]^T \quad (A-21)$$

$$u_{range} = \begin{bmatrix} 0 & 100 \\ 0 & 100 \\ 0 & 100 \\ 0 & 100 \end{bmatrix} \quad (A-22)$$

$$u_{init} = \begin{bmatrix} 57.1586 \\ 50.9737 \\ 39.0769 \\ 49.9805 \end{bmatrix} \quad (A-23)$$

All angles are measured in radians, angular rates are measured in radians per second and linear rates are measured in feet per second.

## Appendix B – Tau Perception Derivations

This appendix details the derivations for the optical flow field estimation techniques and direct gradient method variations discussed in section 4.

### Appendix B.1 – Optical Flow Field Estimation

#### Appendix B.1.1 – Horn and Schunck Optical Flow Estimation

Any image can be regarded as an intensity or brightness pattern, and expressed as a function in these terms:

$$I(x, y, t) \tag{B-1}$$

This image brightness or intensity,  $I$  is a function of position on the image plane and time. The constant brightness assumption states that the intensity of a particular point in the brightness pattern does not change significantly with time; in other words, the information contained within the images of a video sequence does not change between video frames:

$$\frac{d}{dt}I(x, y, t) = 0 \tag{B-2}$$

The customary constant brightness equation is derived by applying the chain rule for differentiation:

$$\frac{dI}{dx} \frac{dx}{dt} + \frac{dI}{dy} \frac{dy}{dt} + \frac{dI}{dt} = I_x u + I_y v + I_t = 0 \tag{B-3}$$

Where  $u$  and  $v$  are the optical flow components in the  $x$  and  $y$  directions, and the  $I_x$ ,  $I_y$  and  $I_t$  terms are the image gradients with respect to  $x$ ,  $y$  and time directions. A single equation is insufficient for calculating the two optical flow components, so a further constraint equation must be introduced for a solution. This second equation is typically where optical flow methods differ most.

Horn and Schunck's second constraint equation assumes that optical flow varies smoothly across the image. The constant brightness assumption states that the brightness of a point remains constant between frames but allows for its motion within the image. If all points in an image moved independently, it would be nearly

impossible to recover the optical flow velocities. However, in the majority of cases, a video will depict rigid objects undergoing motion so that neighbouring points will move with similar velocities; the optic flow field will have some degree of smoothness. This inference can be used to derive a second constraint equation to use in the optical flow calculations by using the Laplacian of the vector flow fields as a measure of smoothness:

$$\nabla^2 u = \frac{\delta^2 u}{\delta x^2} + \frac{\delta^2 u}{\delta y^2} \quad (B-4)$$

$$\nabla^2 v = \frac{\delta^2 v}{\delta x^2} + \frac{\delta^2 v}{\delta y^2} \quad (B-5)$$

For a case where the observers move parallel to the scene, all flow vectors will be aligned and therefore perfectly smooth so the above quantities will be equal to zero. Since it is unlikely that motion is exclusively parallel, it is beneficial to formulate the problem as a minimisation for a solution.

For the solution, it is necessary to calculate image derivatives for each frame in the sequence. However, it is important that the derivatives with respect to each dimension are consistent; they should all be calculated for the same point in time and space. This is straightforward for the spatial dimensions as a simple gradient calculation method can be applied using convolution. A central difference method is typically used for this purpose:

$$\frac{df}{dx} = \frac{f\left(x + \frac{1}{2}h\right) - f\left(x - \frac{1}{2}h\right)}{h} \quad (B-6)$$

However, this becomes more complex in the time dimension, as a central difference method can be applied in two ways: across three frames or across two. Both will introduce a time lag to the system. If the time gradient is calculated across three frames it introduces a full frame lag to the system, as an extra image at frame  $n + 1$  is required to compute the gradient for frame  $n$ . Calculating the frame across two frames reduces the size of this frame lag by half, but means the gradient is assessed at a point between frames. Since it is only possible to assess the spatial image gradients at times when frames are captured, these gradients are no longer



consistent with the time gradient. One solution to this problem is to interpolate the gradients between frames using multiple measurements to ensure consistency. Horn and Schunk used this method, illustrated in Figure 6-2 below, where a formation of cubes represents pixels position in time and space:




Figure 6-2: Pixel structure in  $x$ ,  $y$  and time directions used for averaging

By using the indices  $i$ ,  $j$  and  $k$  to represent values on the  $x$ ,  $y$  and  $t$  axes, an estimation for the gradient values at the centre of the cube can be formulated:

$$I_x \approx \frac{1}{4x} \left( I_{i,j+1,k} - I_{i,j,k} + I_{i+1,j+1,k} - I_{i+1,j,k} + I_{i,j+1,k+1} - I_{i,j,k+1} + I_{i+1,j+1,k+1} - I_{i+1,j,k+1} \right) \quad (B-7)$$

$$I_y \approx \frac{1}{4y} \left( I_{i+1,j,k} - I_{i,j,k} + I_{i+1,j+1,k} - I_{i,j+1,k} + I_{i+1,j,k+1} - I_{i,j,k+1} + I_{i+1,j+1,k+1} - I_{i,j+1,k+1} \right) \quad (B-8)$$

$$I_t \approx \frac{1}{4t} \left( I_{i,j,k+1} - I_{i,j,k} + I_{i+1,j,k+1} - I_{i+1,j,k} + I_{i,j+1,k+1} - I_{i,j+1,k} + I_{i+1,j+1,k+1} - I_{i+1,j+1,k} \right) \quad (B-9)$$

These equations can be applied over a whole image simply by convolving the image with simple masks for each direction in every frame, summing with a consecutive frame and dividing by the relevant interval:

$$I_x = \left( I(x, y, t - 1) * \begin{bmatrix} -\frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} \end{bmatrix} + I(x, y, t) * \begin{bmatrix} -\frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} \end{bmatrix} \right) \frac{1}{x_{px}} \quad (B-10)$$

$$I_y = \left( I(x, y, t - 1) * \begin{bmatrix} \frac{1}{4} & \frac{1}{4} \\ -\frac{1}{4} & -\frac{1}{4} \end{bmatrix} + I(x, y, t) * \begin{bmatrix} \frac{1}{4} & \frac{1}{4} \\ -\frac{1}{4} & -\frac{1}{4} \end{bmatrix} \right) \frac{1}{y_{py}} \quad (B-11)$$

$$I_t = \left( I(x, y, t - 1) * \begin{bmatrix} -\frac{1}{4} & -\frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} \end{bmatrix} + I(x, y, t) * \begin{bmatrix} \frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} \end{bmatrix} \right) \frac{1}{dt} \quad (B-12)$$

Where  $x_{px}$  and  $y_{py}$  represent the size of each pixel in the  $x$  and  $y$  directions, and  $dt$  represents the time period between each frame.

The Laplacians of the flow field to apply the smoothness constraint must also be found. A convenient approximation is:

$$\nabla^2 u \approx K(\bar{u}_{i,j,k} - u_{i,j,k}) \quad (B-13)$$

$$\nabla^2 v \approx K(\bar{v}_{i,j,k} - v_{i,j,k}) \quad (B-14)$$

The local average velocities  $\bar{u}_{i,j,k}$  and  $\bar{v}_{i,j,k}$  are defined by a weighted average of the neighbouring points. The problem is then to minimise the sum of the errors of the two constraint equations. A minimisation can be constructed using the above approximations and a weighting factor  $\alpha$  through the calculus of variations to define the two following equations:

$$(u - \bar{u}) = -I_x \frac{(I_x \bar{u} + I_y \bar{v} + I_t)}{\alpha^2 + I_x^2 + I_y^2} \quad (B-15)$$

$$(v - \bar{v}) = -I_y \frac{(I_x \bar{u} + I_y \bar{v} + I_t)}{\alpha^2 + I_x^2 + I_y^2} \quad (B-16)$$

These equations can then be solved iteratively to produce estimates for the flow field across successive frames.

### Appendix B.1.2 – Lucas-Kanade Optical Flow Estimation

For this method the constant brightness equation is defined for the  $n$  pixels in a window around the point of interest, each represented by  $p_i$ . The constant brightness equations for each pixel can be represented as linear matrix equation:

$$\begin{bmatrix} I_x(p_1) & I_y(p_1) \\ I_x(p_2) & I_y(p_2) \\ \vdots & \vdots \\ I_x(p_n) & I_y(p_n) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} -I_t(p_1) \\ -I_t(p_2) \\ \vdots \\ -I_t(p_n) \end{bmatrix} = A_{xy} \bar{v} = b \quad (B-17)$$

The Lucas-Kanade method assumes that optical flow is constant in this region, so there are more equations than unknowns and is therefore over-determined, so is best solved with a least squares method. The matrix constant brightness equation can be expressed in a more useful form by multiplying by the transverse of  $A_{xy}$  and rearranging for the flow vector:

$$A_{xy}^T A_{xy} \bar{v} = A_{xy}^T b \quad (B-18)$$

$$\bar{v} = (A_{xy}^T A_{xy})^{-1} A_{xy}^T b \quad (B-19)$$

Which expands to:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \sum_i I_x(p_i)^2 & \sum_i I_x(p_i)I_y(p_i) \\ \sum_i I_y(p_i)I_x(p_i) & \sum_i I_y(p_i)^2 \end{bmatrix}^{-1} \begin{bmatrix} -\sum_i I_x(p_i)I_t(p_i) \\ -\sum_i I_y(p_i)I_t(p_i) \end{bmatrix} \quad (B-20)$$

## Appendix B.2 – Derivations for Direct Gradient Method for Tau Perception

The direct gradient method relies on the constant brightness equation (B-3) and the two expressions for optical flow components derived from the weak perspective projection equations:

$$u = f \left( \frac{U}{Z} - \frac{xW}{fZ} \right) \quad (B-21)$$

$$v = f \left( \frac{V}{Z} - \frac{yW}{fZ} \right) \quad (B-22)$$

### Appendix B.2.1 – One Degree of Freedom Translational Motion

The direct gradient method uses the expressions for the optical flow components derived from the weak perspective projection equations, and substitutes them into the constant brightness equation directly:

$$I_x f \left( \frac{U}{Z} - \frac{xW}{fZ} \right) + I_y f \left( \frac{V}{Z} - \frac{yW}{fZ} \right) + I_t = 0 \quad (B-23)$$

Examining a simple case where motion is perpendicular to an approaching planar surface,  $U = V = 0$ , therefore:

$$-\frac{W}{Z} (x I_x + y I_y) + I_t = 0 \quad (B-24)$$

This equation is commonly simplified to:

$$CG + I_t = 0 \quad (B-25)$$

Where  $G$  is the radial gradient, defined as:

$$G = x I_x + y I_y \quad (B-26)$$

And  $C$  is the inverse of time-to-contact:

$$C = -\frac{W}{Z} = \frac{1}{\tau} \quad (B-27)$$

Since the constant brightness equation will never be exactly equal to zero due to quantization errors and noise in the image capture process, it is best to formulate the problem as a minimisation. Using a least-squares method the problem can be stated as a minimisation of the quantity:

$$\sum (CG + I_t)^2 \quad (B-28)$$

To find the minimum value of this expression it is simple to differentiate with respect to  $C$  and equate to zero to find the stationary point, yielding:

$$\sum (CG + I_t) G = 0 \quad (B-29)$$

$$C = -\frac{\sum GI_t}{\sum G^2} \quad (B-30)$$

Time-to-contact is then simply obtained by finding the inverse of  $C$ . This method only requires the three image gradients to function and is very computationally inexpensive once those gradients are found. The gradients can be summed over the whole image or a subset of it. The power of this method becomes clear when considering scenarios with increased complexity.

#### Appendix B.2.2 – Three Degree of Freedom Translational Motion

Horn also allowed for cases where translational motion occurs in all three axes instead of purely perpendicular to the ground plane. In this case,  $U$  and  $V$  are non-zero, so equation (4-17) can be used in its entirety. It is useful to separate the inverse of Tau out again, so the equation can be rearranged as:

$$f \frac{U}{Z} I_x + f \frac{V}{Z} I_y - \frac{W}{Z} (x I_x + y I_y) + I_t = 0 \quad (B-31)$$

Again, using  $C$  and  $G$  and setting:

$$A = f \frac{U}{Z} \quad (B-32)$$

$$B = f \frac{V}{Z} \quad (B-33)$$

The expression can be simplified to:

$$AI_x + BI_y + CG + I_t = 0 \quad (B-34)$$

There are three unknown variables in equation (B-34), but the problem can still be formulated as a least-squares problem:

$$\sum (AI_x + BI_y + CG + I_t)^2 \quad (B-35)$$

Taking the partial derivatives of equation (B-35) with respect to  $A$ ,  $B$  and  $C$  yields three equations to match the three unknowns:

$$\sum (AI_x + BI_y + CG + I_t)I_x = 0 \quad (B-36)$$

$$\sum (AI_x + BI_y + CG + I_t)I_y = 0 \quad (B-37)$$

$$\sum (AI_x + BI_y + CG + I_t)G = 0 \quad (B-38)$$

By expanding equations (B-36), (B-37), and (B-38) formulating them as a linear matrix equation:

$$\begin{bmatrix} \Sigma I_x^2 & \Sigma I_x I_y & \Sigma G I_x \\ \Sigma I_x I_y & \Sigma I_y^2 & \Sigma G I_y \\ \Sigma G I_x & \Sigma G I_y & \Sigma G^2 \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix} = - \begin{bmatrix} \Sigma I_x I_t \\ \Sigma I_y I_t \\ \Sigma G I_t \end{bmatrix} \quad (B-39)$$

Again, all of the required terms are products of the three image gradients which can be used to solve for time-to-contact, in addition to the  $A$  and  $B$  parameters. Solving for  $C$  is the primary focus, but these extra quantities allow for estimation of the image focus of expansion  $(x_0, y_0)$ , which can be useful for translational control:

$$\frac{A}{C} = \frac{f \frac{U}{Z}}{-\frac{W}{Z}} = -f \frac{U}{W} \quad (B-40)$$

$$\frac{B}{C} = \frac{f \frac{V}{Z}}{-\frac{W}{Z}} = -f \frac{V}{W} \quad (B-41)$$

Appendix B.2.3 – One Degree of Freedom Translational Motion with Relative Slope  
A further level of complexity can be added by considering a camera approaching an inclined plane with one degree of translational freedom along the optical axis. The slope of the plane can be defined by two parameters,  $p_s$  and  $q_s$ , to represent the amount of slope in the  $X$  and  $Y$  directions above a height reference,  $Z_0$ :

$$Z = Z_0 + p_s X + q_s Y \quad (B-42)$$

The weak perspective projection equations can be substituted into equation (B-42) and rearranged to:

$$Z = \frac{Z_0}{\left(1 - p_s \frac{x}{f} - q_s \frac{y}{f}\right)} \quad (B-43)$$

Equation (B-43) can be then substituted into the image rate equations ((B-29) and (B-30)) to obtain expression for the optical flow components when approaching an inclined surface. The horizontal translational rates are again set to zero, so  $U = V = 0$  and:

$$u = -x \frac{W}{Z_0} \left(1 - p_s \frac{x}{f} - q_s \frac{y}{f}\right) \quad (B-44)$$

$$v = -y \frac{W}{Z_0} \left(1 - p_s \frac{x}{f} - q_s \frac{y}{f}\right) \quad (B-45)$$

Substituting equations (B-44) and (B-45) into the constant brightness equation and rearranging:

$$-(x I_x + y I_y) \frac{W}{Z_0} \left(1 - p_s \frac{x}{f} - q_s \frac{y}{f}\right) + I_t = 0 \quad (B-46)$$

Setting:

$$P_s = \frac{p_s W}{f Z_0} \quad (B-47)$$

$$Q_s = \frac{q_s W}{f Z_0} \quad (B-48)$$

Then substituting in to equation (B-46) results in:

$$G(C + P_s x + Q_s y) + I_t = 0 \quad (B-49)$$

Since motion is along the optical axis the contact point with the plane will be at  $(0, 0, Z_0)$ , hence  $C$  will still be the inverse of time-to-contact with the ground plane. The problem can again be formulated as a least square minimisation in the same manner to produce the linear matrix equation:

$$\begin{bmatrix} \Sigma G^2 x^2 & \Sigma G^2 xy & \Sigma G^2 x \\ \Sigma G^2 xy & \Sigma G^2 y^2 & \Sigma G^2 y \\ \Sigma G^2 x & \Sigma G^2 y & \Sigma G^2 \end{bmatrix} \begin{bmatrix} P_s \\ Q_s \\ C \end{bmatrix} = - \begin{bmatrix} \Sigma GxI_t \\ \Sigma GyI_t \\ \Sigma GI_t \end{bmatrix} \quad (B-50)$$

Appendix B.2.4 – Three Degree of Freedom Translational Motion with Relative Slope  
It is possible to account for both three degree of freedom translational motion and relative slope between camera and ground plane. This can be achieved by substituting the full equations for the optical flow components (equation (B-21) and (B-22)) and the expression for altitude from slope (equation (B-43)) into the constant brightness assumption:

$$\left(1 - p_s \frac{x}{f} - q_s \frac{y}{f}\right) \left(f \left(\frac{U}{Z_0} - \frac{xW}{fZ_0}\right) I_x + f \left(\frac{V}{Z_0} - \frac{yW}{fZ_0}\right) I_y\right) + I_t = 0 \quad (B-51)$$

Using the previous variable groupings, this can be simplified to:

$$C \left(1 + x \frac{P_s}{C} + y \frac{Q_s}{C}\right) \left(G + I_x \frac{A}{C} + I_y \frac{B}{C}\right) + I_t = 0 \quad (B-52)$$

Setting:

$$F = 1 + x \frac{P_s}{C} + y \frac{Q_s}{C} \quad (B-53)$$

$$D = G + I_x \frac{A}{C} + I_y \frac{B}{C} \quad (B-54)$$

The problem can again be posed as a least-squares minimisation of the form:

$$\sum (CFD + I_t)^2 \quad (B-55)$$

Differentiating equation (B-55) with respect to the five variables  $A$ ,  $B$ ,  $C$ ,  $P_s$  and  $Q_s$  to find the minimum point yields five equations with five unknowns, but the equations are nonlinear so must be solved numerically. Horn proposed a hierarchical scheme to achieve this by sequentially assuming two of the variables are known and iteratively solving two sets of linear equations. By assuming that  $P_s/C$  and  $Q_s/C$  are constant and known, equation (B-55) can be differentiated with respect to  $A$ ,  $B$  and  $C$  to give the following equations:



$$\sum (F(CG + I_x A + I_y B) + I_t) F I_x = 0 \quad (B-56)$$

$$\sum (F(CG + I_x A + I_y B) + I_t) F I_y = 0 \quad (B-57)$$

$$\sum (F(CG + I_x A + I_y B) + I_t) F G = 0 \quad (B-58)$$

Assuming  $A$  and  $B$  are constants and differentiating with respect to  $P$ ,  $Q$  and  $C$ :

$$\sum (D(C + x P_s + y Q_s) + I_t) D x = 0 \quad (B-59)$$

$$\sum (D(C + x P_s + y Q_s) + I_t) D y = 0 \quad (B-60)$$

$$\sum (D(C + x P_s + y Q_s) + I_t) D = 0 \quad (B-61)$$

These equations can be arranged into two sets of linear matrix equations:

$$\begin{bmatrix} \Sigma F^2 I_x^2 & \Sigma F^2 I_x I_y & \Sigma F^2 G I_x \\ \Sigma F^2 I_x I_y & \Sigma F^2 I_y^2 & \Sigma F^2 G I_y \\ \Sigma F^2 G I_x & \Sigma F^2 G I_y & \Sigma F^2 G^2 \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix} = - \begin{bmatrix} \Sigma F I_x I_t \\ \Sigma F I_y I_t \\ \Sigma F G I_t \end{bmatrix} \quad (B-62)$$

$$\begin{bmatrix} \Sigma D^2 x^2 & \Sigma D^2 xy & \Sigma D^2 x \\ \Sigma D^2 xy & \Sigma D^2 y^2 & \Sigma D^2 y \\ \Sigma D^2 x & \Sigma D^2 y & \Sigma D^2 \end{bmatrix} \begin{bmatrix} P_s \\ Q_s \\ C_s \end{bmatrix} = - \begin{bmatrix} \Sigma D x I_t \\ \Sigma D y I_t \\ \Sigma D I_t \end{bmatrix} \quad (B-63)$$

Equation (B-62) and (B-63) can be solved iteratively to produce estimates for all five variables using only image gradients, camera properties and some initial estimates for  $P_s/C$  and  $Q_s/C$ .

Appendix B.2.5 – Six Degree of Freedom Translational and Rotational Motion with Relative Slope

To account for rotational motion, the fundamental weak perspective projection equations must be re-examined as per Horn in Ref. [122]. A point in the visual scene  $P$  creates an image  $p$  on a camera. The coordinates of these two points are:

$$P = [X, Y, Z] \quad (B-64)$$

$$p = [x, y, f] \quad (B-65)$$

And they are related to each other through the weak perspective projection equations:

$$p = [x, y, f] = f \left[ \frac{X}{Z}, \frac{Y}{Z}, \frac{Z}{Z} \right] = \frac{fP}{P \cdot \hat{z}} \quad (B-66)$$

Where  $\hat{z}$  is a unit vector in the  $Z$  direction. The velocity of point  $P$  relative to the camera frame will be:

$$\frac{dP}{dt} = -\bar{V} - \bar{\omega} P \quad (B-67)$$

Where the  $\bar{V}$  and  $\bar{\omega}$  are the translational and rotational velocities respectively, and their components are:

$$\bar{V} = [U, V, W] \quad (B-68)$$

$$\bar{\omega} = [\omega_x, \omega_y, \omega_z] \quad (B-69)$$

The corresponding motion of the image point  $p$  will be:

$$\frac{dp}{dt} = \frac{d}{dt} \left( \frac{fP}{P \cdot \hat{z}} \right) \quad (B-70)$$

$$\frac{dp}{dt} = \frac{f\dot{P} (P \cdot \hat{z}) - (P \cdot \hat{z})fP}{(P \cdot \hat{z})^2} \quad (B-71)$$

Using the vector triple product this can be rearranged to:

$$a \times (b \times c) = (c \cdot a)b - (a \cdot b)c \quad (B-72)$$

$$\frac{dp}{dt} = f \frac{\hat{z} \times (\dot{P} \times P)}{(P \cdot \hat{z})^2} \quad (B-73)$$

Substituting in equation (B-67):

$$\frac{dp}{dt} = \frac{f}{(P \cdot \hat{z})^2} \left[ \hat{z} \times ((-\bar{\omega} \times P - \bar{V}) \times P) \right] \quad (B-74)$$

Rearranging equation (B-66) and substituting in:

$$\frac{dp}{dt} = \frac{f}{(P \cdot \hat{z})^2} \left[ \hat{z} \times \left( \left( -\frac{P \cdot \hat{z}}{f} \bar{\omega} \times p - \bar{V} \right) \times \frac{P \cdot \hat{z}}{f} r \right) \right] \quad (B-75)$$

Since  $a \times b = -(b \times a)$  equation (B-75) can be re-ordered and terms cancelled to yield:

$$\frac{dp}{dt} = -\hat{z} \times \left[ p \times \left( p \times \frac{\bar{\omega}}{f} - \frac{\bar{V}}{P \cdot \hat{z}} \right) \right] \quad (B-76)$$

Substituting in the relevant vectors for  $p, P, \bar{V}$  and  $\bar{\omega}$ :

$$\frac{dp}{dt} = -\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \times \left[ \begin{bmatrix} x \\ y \\ f \end{bmatrix} \times \left( \begin{bmatrix} x \\ y \\ f \end{bmatrix} \times \begin{bmatrix} \omega_x/f \\ \omega_y/f \\ \omega_z/f \end{bmatrix} - \begin{bmatrix} U/Z \\ V/Z \\ W/Z \end{bmatrix} \right) \right] \quad (B-77)$$

Solving equation (B-77) yields:

$$\frac{dp}{dt} = \begin{bmatrix} \frac{-fU + xW}{Z} + \frac{xy}{f} \omega_x - \frac{x^2 + f^2}{f} \omega_y + y\omega_z \\ \frac{-fV + yW}{Z} + \frac{y^2 + f^2}{f} \omega_x - \frac{xy}{f} \omega_y - x\omega_z \\ 0 \end{bmatrix} \quad (B-78)$$

Since the elements of  $p$  are the  $x$  and  $y$  position in the image, the components of the time derivative of  $p$  will be the optical flow velocity components  $u$  and  $v$ . Therefore:

$$u = -f \frac{U}{Z} + x \frac{W}{Z} + \frac{xy}{f} \omega_x - \frac{x^2 + f^2}{f} \omega_y + y\omega_z \quad (B-79)$$

$$v = -f \frac{V}{Z} + y \frac{W}{Z} + \frac{y^2 + f^2}{f} \omega_x - \frac{xy}{f} \omega_y - x\omega_z \quad (B-80)$$

These expressions for the flow components can then be substituted into the constant brightness equation and the gradient methodology applied to produce time-to-contact estimates for cases with the observer in rotational motion if the angular velocities are known. These angular velocities can be measured with a gyroscopic inertial navigation system. Using the above expressions for surface slope and translational velocity, and the new rotational motion quantities:

$$J = \frac{xy}{f} \omega_x - \frac{x^2 + f^2}{f} \omega_y + y\omega_z \quad (B-81)$$

$$K = \frac{y^2 + f^2}{f} \omega_x - \frac{xy}{f} \omega_y - x\omega_z \quad (B-82)$$

The following equation can be constructed:

$$C \left( 1 + x \frac{P_s}{C} + y \frac{Q_s}{C} \right) \left( G + I_x \frac{A}{C} + I_y \frac{B}{C} \right) + JI_x + KI_y + I_t = 0 \quad (B-83)$$

Constructing a least-squares minimisation and further simplifying with  $F$  and  $D$  variable groupings:

$$\sum [CFD + JI_x + KI_y + I_t]^2 \quad (B-84)$$

Creating another variable grouping:

$$I_\omega = JI_x + KI_y + I_t \quad (B-85)$$

Allows the construction of two systems of equations which can again be solved iteratively to estimate time-to-contact given the image gradients, camera properties and rotational velocities:

$$\begin{bmatrix} \Sigma F^2 I_x^2 & \Sigma F^2 I_x I_y & \Sigma F^2 G I_x \\ \Sigma F^2 I_x I_y & \Sigma F^2 I_y^2 & \Sigma F^2 G I_y \\ \Sigma F^2 G I_x & \Sigma F^2 G I_y & \Sigma F^2 G \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix} = - \begin{bmatrix} \Sigma F^2 I_x I_\omega \\ \Sigma F^2 I_y I_\omega \\ \Sigma F^2 G I_\omega \end{bmatrix} \quad (B-86)$$

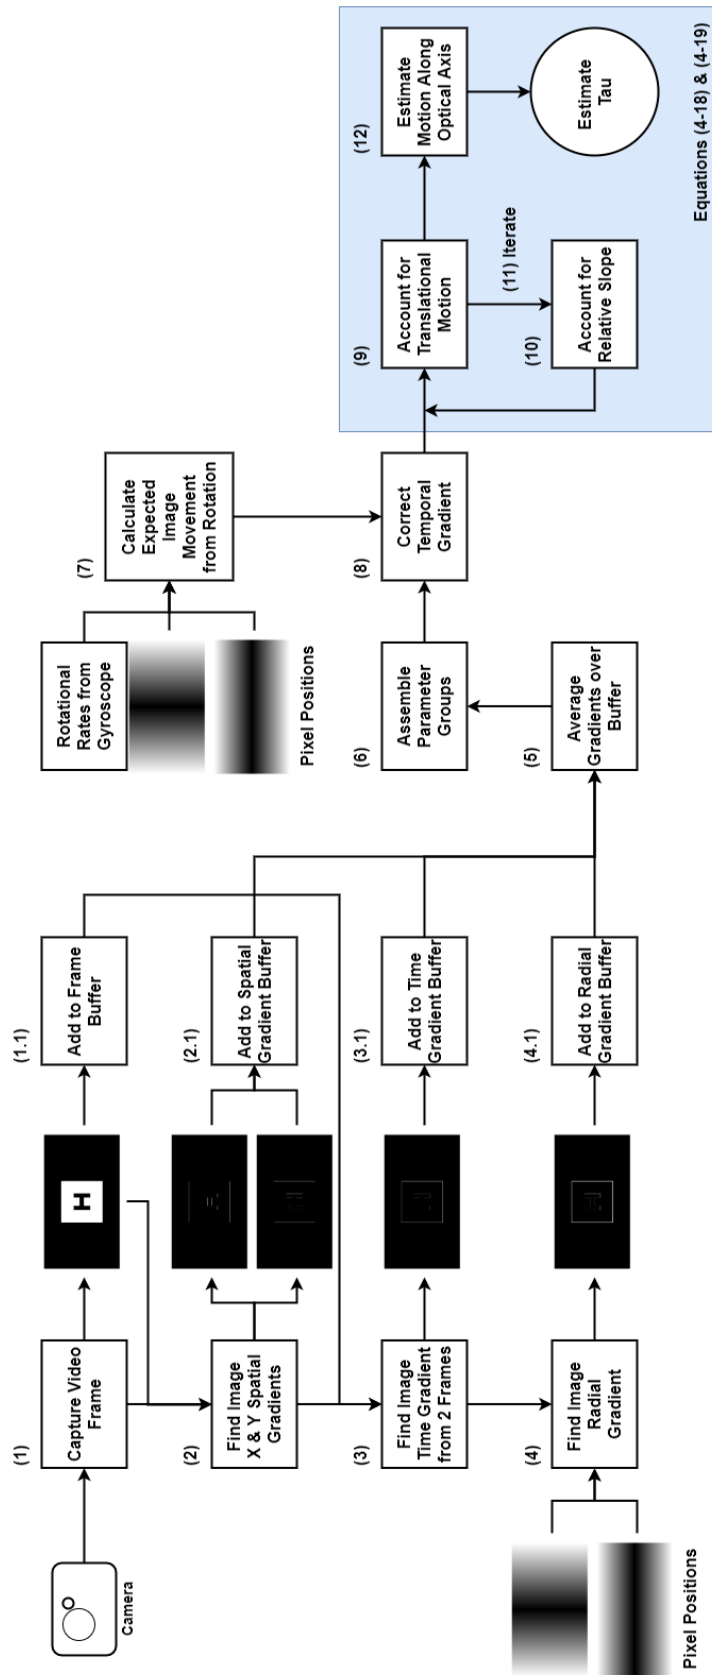
$$\begin{bmatrix} \Sigma D^2 x^2 & \Sigma D^2 xy & \Sigma D^2 x \\ \Sigma D^2 xy & \Sigma D^2 y^2 & \Sigma D^2 y \\ \Sigma D^2 x & \Sigma D^2 y & \Sigma D^2 \end{bmatrix} \begin{bmatrix} P_s \\ Q_s \\ C \end{bmatrix} = - \begin{bmatrix} \Sigma D x I_\omega \\ \Sigma D y I_\omega \\ \Sigma D I_\omega \end{bmatrix} \quad (B-87)$$

These equations allow for time-to-contact to be estimated for a camera moving in six degrees of freedom relative to a target inclined to the camera in two dimensions, which covers most possible scenarios.

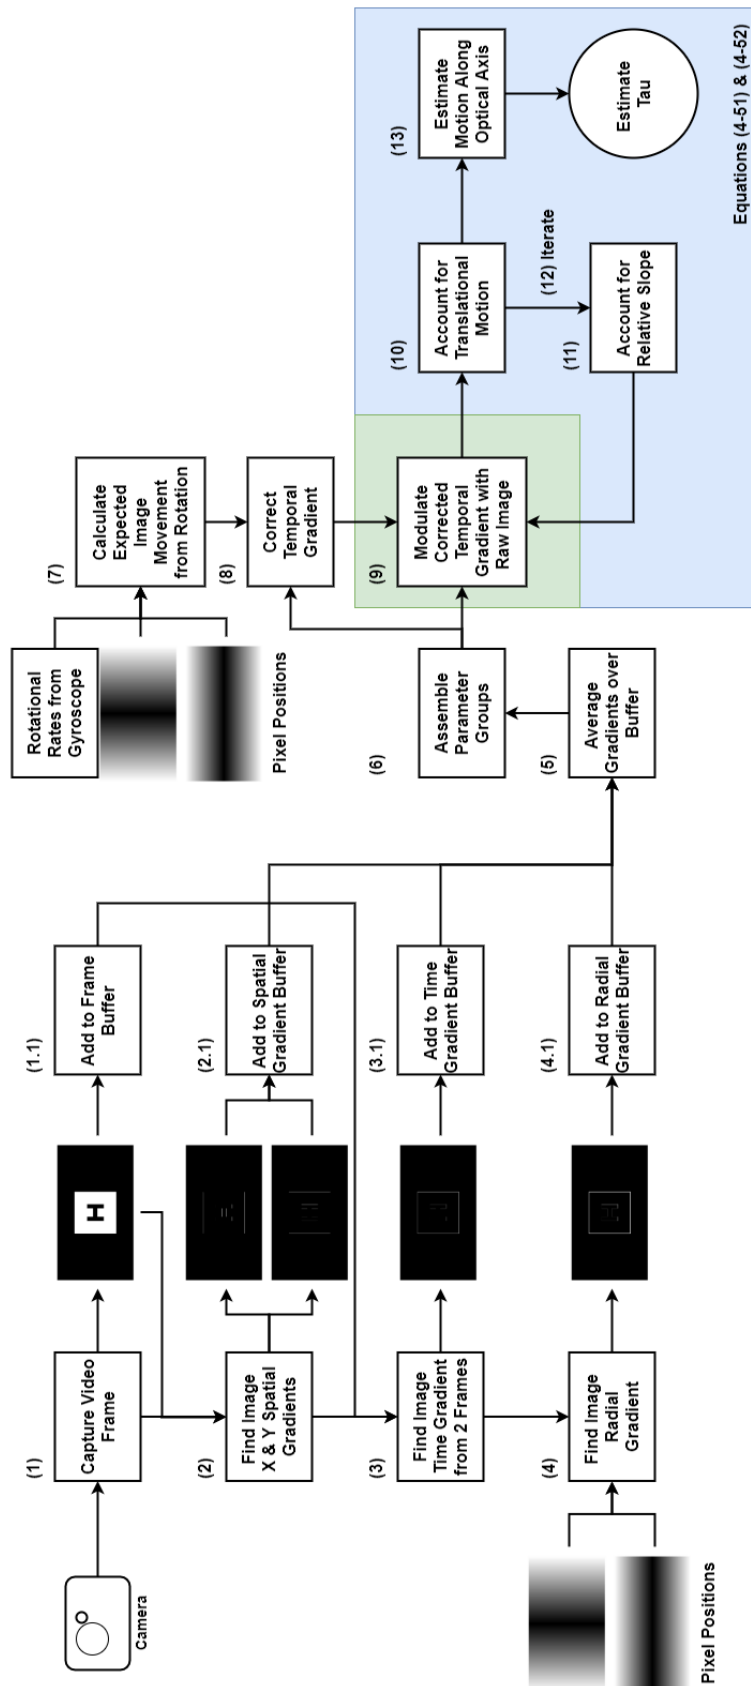
### Appendix B.3 – Direct Gradient Method Block Diagrams

This appendix contains block diagrams that detail the implementation of the direct gradient method in its base form, described in section 4.3, and with the extension proposed in section 4.4. The numbers by each block correspond to the numbered lists in these two respective sections that outline the process in words.

Appendix B.3.1 – Six Degree of Freedom Translational and Rotational Motion with Relative Slope Block Diagram

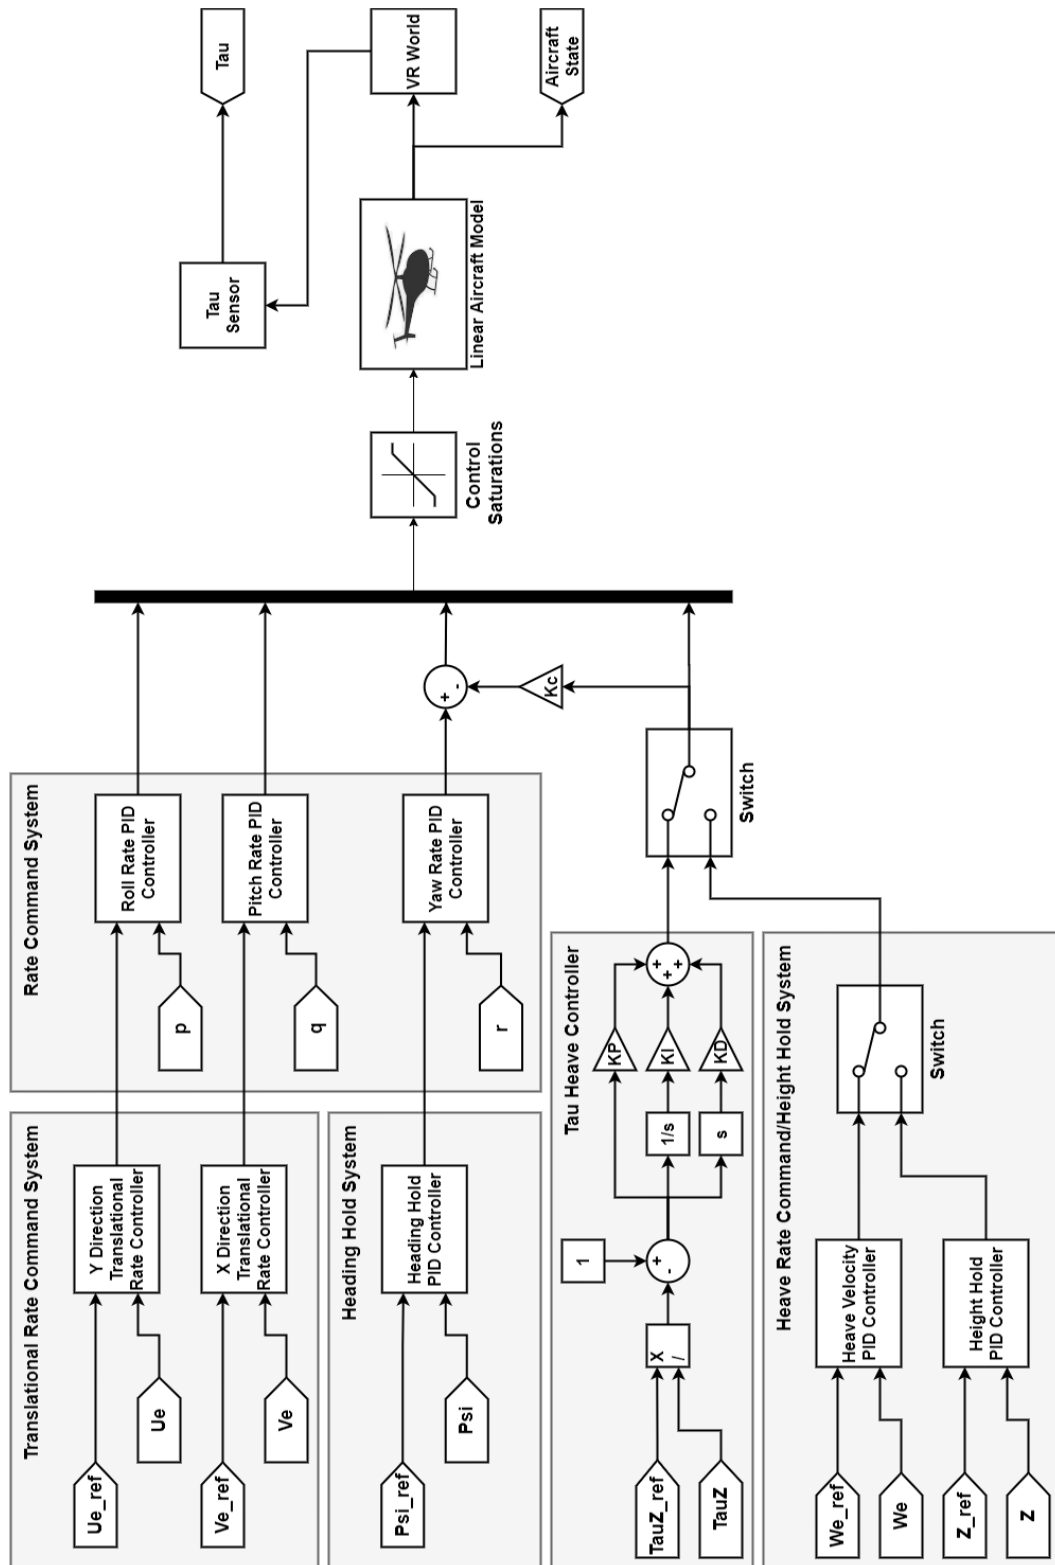


Appendix B.3.2 – Six Degree of Freedom Translational and Rotational Motion with Relative Slope with Brightness Change Correction Block Diagram



# Appendix C – Simulation Model Flight Controller Systems

## Appendix C.1 – Rotary-Wing Flight Controller Structure

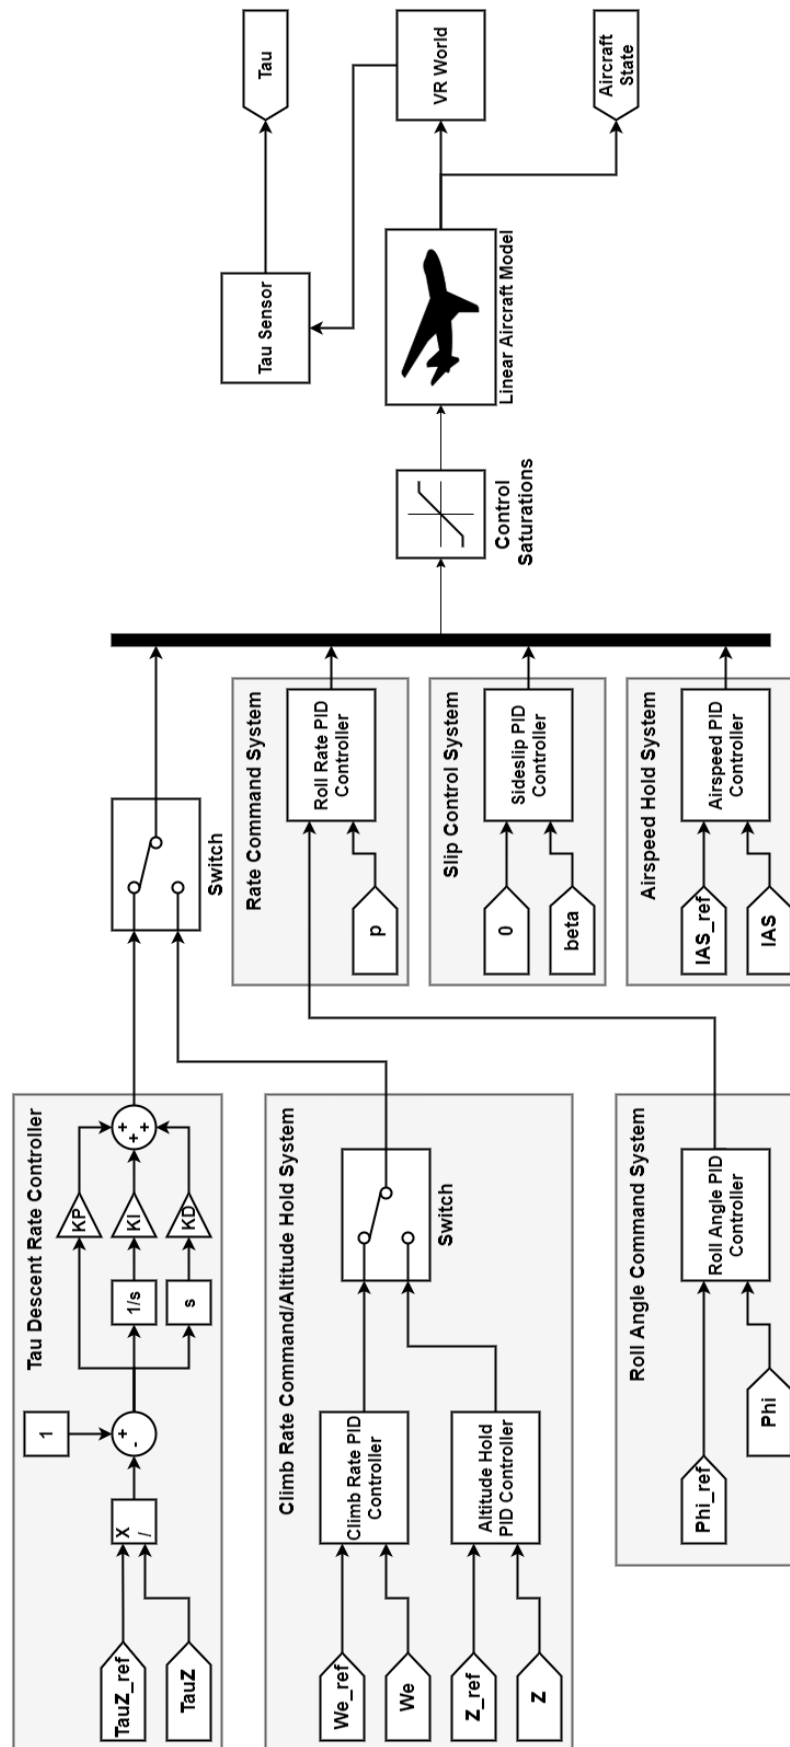




Appendix C.2 – Rotary-Wing Flight Controller Gains

Parameter		SH-60B	MQ-8B	R-Max	T-Rex
Tau Controller	$K_P$	-10	-2	-100	-10
	$K_I$	0	-0.2	0	-0.2
	$K_D$	0	0	0	0
Roll Rate Controller	$K_P$	-1	-1	-1	-1
	$K_I$	0	0	0	0
	$K_D$	0	0	0	0
Pitch Rate Controller	$K_P$	-2	-2	-2	-2
	$K_I$	0	0	0	0
	$K_D$	0	0	0	0
Yaw Rate Controller	$K_P$	-0.5	-0.5	-1	-0.5
	$K_I$	0	0	0	0
	$K_D$	0	0	0	0
X Direction Translation Rate Controller	$K_P$	-0.2	-0.2	-0.5	-0.2
	$K_I$	0	0	0	0
	$K_D$	0	0	0	0
Y Direction Translation Rate Controller	$K_P$	-0.2	-0.2	-0.5	-0.2
	$K_I$	0	0	0	0
	$K_D$	0	0	0	0
Heading Hold Controller	$K_P$	-2	-2	-2	-2
	$K_I$	0	0	0	0
	$K_D$	0	0	0	0
Heave Velocity Controller	$K_P$	-1	-1	-1	-1
	$K_I$	0	0	0	0
	$K_D$	0	0	0	0
Height Hold Controller	$K_P$	-10	-10	-10	-10
	$K_I$	0	0	0	0
	$K_D$	0	0	0	0
Crossfeed	$K_C$	0.5	0.5	0.5	0.5

# Appendix C.3 – Fixed-Wing Flight Controller Structure



## Appendix C.4 – Fixed-Wing Flight Controller Gains

Parameter		3DR Aero
Tau Controller	$K_P$	-100
	$K_I$	0
	$K_D$	0
Roll Rate Controller	$K_P$	-10
	$K_I$	0
	$K_D$	0
Roll Angle Controller	$K_P$	-10
	$K_I$	-10
	$K_D$	0
Airspeed Hold Controller	$K_P$	-50
	$K_I$	0
	$K_D$	0
Sideslip Controller	$K_P$	-1
	$K_I$	0
	$K_D$	0
Climb Rate Controller	$K_P$	-10
	$K_I$	0
	$K_D$	0
Height Hold Controller	$K_P$	-10
	$K_I$	0
	$K_D$	0

## References

- [1] Federal Aviation Authority, "FAA Aerospace Forecast: Fiscal Years 2017-2037," FAA, 2017.
- [2] Federal Aviation Authority, "FAA Aerospace Forecast: Fiscal Years 2019-2039," FAA, 2019.
- [3] Federal Aviation Authority, "U.S. Civil Airmen Statistics," 2019. [Online]. Available:  
[https://www.faa.gov/data\\_research/aviation\\_data\\_statistics/civil\\_airmen\\_statistics/](https://www.faa.gov/data_research/aviation_data_statistics/civil_airmen_statistics/). [Accessed 01 09 2019].
- [4] CBS News, "Amazon Unveils Futuristic Plan: Delivery by Drone," CBS News, 01 December 2013. [Online]. Available:  
<https://www.cbsnews.com/news/amazon-unveils-futuristic-plan-delivery-by-drone/>. [Accessed 01 September 2019].
- [5] Amazon, "Amazon Prime Air First Air Delivery," Amazon, 7 December 2017. [Online]. Available: <https://www.amazon.com/Amazon-Prime-Air/b?ie=UTF8&node=8037720011>. [Accessed 1 September 2019].
- [6] DHL, "DHL Launches its First Regular Fully Automated and Intelligent Urban Drone Delivery Service," DHL, 16 May 2019. [Online]. Available:  
<https://www.dpdhl.com/en/media-relations/press-releases/2019/dhl-launches-its-first-regular-fully-automated-and-intelligent-urban-drone-delivery-service.html>. [Accessed 01 September 2019].
- [7] The Guardian, "Domino's Planning Drone Pizza Delivery Service in New Zealand," The Guardian, 25 August 2016. [Online]. Available:  
<https://www.theguardian.com/business/2016/aug/25/dominos-planning-drone-pizza-delivery-service-new-zealand-auckland-trial>. [Accessed 01 September 2019].

- [8] Google, "X-Wing," Alphabet, 2019. [Online]. Available: <https://x.company/projects/wing/>. [Accessed 1 September 2019].
- [9] L. Brooke-Holland, "Overview of Military Drones used by the UK Armed Forces," House of Commons Library, London, 2015.
- [10] General Atomics Aeronautical, "Italian Air Force MQ-9," General Atomics Aeronautical, 2019. [Online]. Available: <http://www.ga-asi.com/predator-b>. [Accessed 01 09 2019].
- [11] DefPost, "British Army Watchkeeper Unmanned Aircraft Crashes Near Aberporth," DefPost, 18 06 2018. [Online]. Available: <https://defpost.com/british-army-watchkeeper-unmanned-aircraft-crashes-near-aberporth/>. [Accessed 01 09 2019].
- [12] Lockheed Martin, "Desert Hawk III," Lockheed Martin, [Online]. Available: <https://www.lockheedmartin.com/en-us/products/desert-hawk.html>. [Accessed 01 09 2019].
- [13] Boeing, "ScanEagle," Boeing, [Online]. Available: <https://www.boeing.com/defense/autonomous-systems/scaneagle/index.page>. [Accessed 01 09 2019].
- [14] Defence Imagery, "Black Hornet Nano," Defence Imagery, [Online]. Available: <http://www.defenceimagery.mod.uk/fotoweb/fwbin/download.dll/45153802.jpg>. [Accessed 01 09 2019].
- [15] BAE Systems, "Taranis," BAE Systems, [Online]. Available: <https://www.baesystems.com/en/product/taranis>. [Accessed 01 09 2019].
- [16] HM Government, "National Security Strategy and Strategic Defense and Security Review 2015," OGL, London, 2015.
- [17] Ministry of Defence, "Taranis Unveiled," *Desider: the magazine for defence equipment and support*, pp. 20-21, 10 August 2014.

- [18] Department of Defense, "Unmanned Systems integrated Roadmap 2017:2042," Department of Defense, Washington D.C., 2017.
- [19] Department for Transport, "Public dialogue on drone use in the UK: Moving Britain Ahead," OGL, London, 2016.
- [20] United States Postal Service: Office of Inspector General, "Public perception of Drone Delivery in the United States," United States postal Service, 2016.
- [21] G. Ridgway, *Optical Tau Theory: Current and Future in Fixed-Wing Operations*, Liverpool: University of Liverpool, 2012.
- [22] UAS Vision, "Swiss Post Suspends Drone Delivery Service Second Crash," 2 August 2019. [Online]. Available: <https://www.uasvision.com/2019/08/02/swiss-post-suspends-drone-delivery-service-after-second-crash/>.
- [23] C. Whitlock, "When drones fall from the sky," *Washington Post*, 20 June 2014.
- [24] A. Susini, "A Technocritical Review of Drones Crash Risk Probabilistic Consequences and its Societal Acceptance," in *Risk Information Management, Risk Models, and Applications*, Berlin, Germany, 2015.
- [25] K. W. Williams, "A Summary of Unmanned Aircraft Accident/Incident Data: Human Factors Implications," Federal Aviation Administration, Oklahoma City, 2004.
- [26] D. N. Lee, "A theory of visual control of braking based on information about time-to-collision," *Perception*, vol. 5, no. 4, pp. 437-459, 1976.
- [27] G. D. Padfield, "The Tau of Flight Control," *The Aeronautical Journal*, vol. 115, no. 1171, pp. 521-556, 2011.
- [28] M. Jump and G. D. Padfield, "Investigation of the Flare Manoeuvre Using Optical Tau," *Journal of Guidance, Control and Dynamics*, vol. 29, no. 5, pp. 1189-1200, 2006.

- [29] G. D. Padfield, D. N. Lee and R. Bradley, "How do helicopter pilots know when to stop, turn or pull up?," *Journal of American Helicopter Society*, vol. 48, no. 2, pp. 80-92, 2003.
- [30] F. Kendoul, "Four-Dimensional Guidance and Control of Movement Using Time-to-Contact: Application to Automated Docking and Landing of Unmanned Rotorcraft Systems," *International Journal of Robotics Research*, vol. 33, no. 2, pp. 237-267, 2013.
- [31] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen and G. Lachapelle, "GPS Vulnerability to Spoofing Threats and a Review of Antispoofing techniques," *International journal of Navigation and Observation*, vol. 2012, pp. 1-16, 2012.
- [32] M. T. Alkowalaty, V. M. Becerra and W. Holderbaum, "Bio-Inspired Autonomous Visual Vertical Control of a Quadrotor Unmanned Aerial Vehicle," *Journal of Guidance, Control and Dynamics*, vol. 38, no. 2, pp. 249-262, 2015.
- [33] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *IEEE International Conference on Robotics and Automation*, Shanghai, China, 2011.
- [34] University of Liverpool, "Postgraduate Code of Practice," 2019. [Online]. Available: <https://www.liverpool.ac.uk/media/livacuk/tqsd/code-of-practice-on-assessment/PGR-code-of-practice.pdf>.
- [35] Ministry of Defence, "Joint Doctrine Note 3/10, Unmanned Aircraft Systems Terminology, Definitions and Classification," Forms and Publications Section, DSDA Operations Centre, 2010.
- [36] Ministry of Defence, "Joint Doctrine Publication 0-30.2 Unmanned Aircraft Systems," Ministry of Defence, London, 2018.
- [37] C. Cole, "Accidents will happen: a review of military drone crash data as the UK considers allowing large military drone flights in its airspace," Drone Wars UK, Oxford, 2019.

- [38] F. Kendoul, "Survey of Advances in Guidance, Navigation, and Control of Unmanned Rotorcraft Systems," *Journal of Field Robotics*, vol. 22, no. 2, pp. 315-378, 2012.
- [39] H.-M. Huang, E. Messina and J. Albus, "Autonomy Levels for Unmanned Systems (ALFUS) Framework," NIST Special Publication, 2007.
- [40] L. Meier, P. Tanskanen, F. Fraundorfer and M. Pollefeys, "Pixhawk: A System for Autonomous Flight using Onboard Computer Vision," in *IEEE International Conference on Robotics and Automation*, Shanghai, China, 2011.
- [41] O. Amidi, T. Kanade and K. Fujita, "A visual odometer for autonomous helicopter flight," *Robotics and Autonomous System*, vol. 28, no. 2-3, pp. 185-193, 1999.
- [42] J. Kelly, S. Saripalli and G. & Sukhatme, "Combined visual and inertial navigation for an unmanned aerial vehicle.," in *6th International Conference on Field and Service Robotics*, Chamonix, France, 2007.
- [43] F. Caballero, L. Merino, J. Ferruz and A. Ollero, "A visual odometer without 3D reconstruction for aerial vehicles," in *Proceedings of the 2005 International Conference on Robotics and Automation*, Barcelona, Spain, 2005.
- [44] R. W. Madison, G. A. Andrews, D. Paul, S. A. Rasmussen and M. S. Bottkol, "Vision-aided navigation for small UAVs in GPS-challenged environments.," in *AIAA Infotech@Aerospace*, Rohnert Park, USA, 2007.
- [45] F. Kendoul, I. Fantoni and K. Nonami, "Optic flow-based vision system for autonomous 3D localization and control of small aerial vehicles," *Robotics and Autonomous Systems*, vol. 57, no. 6, pp. 591-602, 2009.
- [46] Ministry of Defence, "Joint Doctrine Note 2/11, The UK Approach to Unmanned Aircraft Systems," Forms and Publication Section, DSDA Operations Centre, 2011.



- [47] C. Theodore, D. Rowley, D. Hubbard, A. Ansar, L. Matthies, S. Goldberg and M. Whalley, "Flight trials of a rotorcraft unmanned aerial vehicle landing autonomously at unprepared sites," in *American Helicopter Society 62nd Annual Forum*, Phoenix, USA, 2006.
- [48] A. A. Proctor and E. N. Johnson, "Vision-only approach and landing," in *AIAA Guidance, Navigation and Control Conference and Exhibit*, San Francisco, California, 2005.
- [49] E. Johnson, A. J. Calise, Y. Watanabe, J. Ha and J. C. Neidhoefer, "Real time vision-based relative aircraft navigation," *Journal of Aerospace Computing, Information, and Communication*, vol. 4, pp. 707-738, 2007.
- [50] G. Conte and P. Doherty, "Vision-based unmanned aerial vehicle navigation using geo-referenced information," *EURASIP Journal on Advances in Signal Processing*, vol. 2009, no. 387308, pp. 1-18, 2009.
- [51] F. Lindsten, J. Callmer, H. Ohlsson, D. Tornqvist, T. B. Schon and F. Gustafsson, "Geo-referencing for UAV navigation using environmental classification," in *IEEE International Conference on Robotics and Automation*, Anchorage, USA, 2010.
- [52] A. J. Davison, I. D. Reid, N. D. Molton and O. Stasse, "MonoSLAM: Real-Time Single Camera SLAM," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 6, pp. 1-16, 2007.
- [53] J. J. Gibson, *The Perception of the Visual World*, Cambridge: The Riverside Press, 1950.
- [54] F. Ruffier and N. Franceschini, "Optic flow regulation: The key to aircraft automatic guidance," *Robotics and Autonomous Systems*, vol. 50, no. 4, pp. 177-194, 2005.

- [55] T. R. Fell, M. D. White, M. Jump and I. Owen, "Sensitivity study of a small maritime rotary UAS operating in a turbulent airwake," in *Proceedings of the 71st AHS Forum*, Virginia Beach, 2015.
- [56] B. Ferrier, A. Sehgal and R. Ernst, "Fire Scout UAV Launch and Recovery System Performance Improvement," NAVAIR, 2014.
- [57] Northrop Grumman, "Northrop Grumman-developed US Navy MQ-8B Fire Scout completes second test period onboard USS McInerney," May 2009. [Online]. Available: <https://news.northropgrumman.com/news/releases/northrop-grumman-developed-u-s-navy-mq-8b-fire-scout-completes-second-test-period-onboard-uss-mcinerney-ffg-8>.
- [58] EASA, "Regulation (EU) No 376/2014," 24 04 2014. [Online]. Available: <https://www.easa.europa.eu/document-library/regulations/regulation-eu-no-3762014>.
- [59] FAA, "Small Unmanned Aircraft Regulations (Part 107)," 21 June 2016. [Online]. Available: [https://www.faa.gov/news/fact\\_sheets/news\\_story.cfm?newsId=20516](https://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=20516).
- [60] G. Wild, J. Murray and G. Baxter, "Exploring civil drone accidents and incidents to help prevent potential air disasters," *Aerospace*, vol. 22, no. 3, pp. 1-11, 2016.
- [61] Z. Rosenberg, "Fire Scout flight suspended due to crashes," FlightGlobal, 9 April 2012. [Online]. Available: <https://www.flightglobal.com/news/articles/firescout-flights-suspended-due-to-crash-370520/>. [Accessed 1 September 2019].
- [62] D. Heeger, "Perception Lecture Notes: Visual Motion Perception," New York University, 2006. [Online]. Available:

<https://www.cns.nyu.edu/~david/courses/perception/lecturenotes/motion/motion.html>.

- [63] D. N. Lee, D. S. Young and D. Rewt, "How do somersaulters land on their feet?," *Journal of experimental psychology: human perception and performance*, vol. 18, no. 4, pp. 1195-1202, 1992.
- [64] D. N. Lee, J. R. Lishman and J. A. Thomson, "Regulation of gait in long jumping," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 8, no. 3, pp. 448-459, 1982.
- [65] W. H. Warren, D. S. Young and D. N. Lee, "Visual control of step length during running over irregular terrain," *Journal of experimental psychology: Human perception and performance*, vol. 12, no. 3, pp. 259-266, 1986.
- [66] D. N. Lee, D. S. Young, P. E. Reddish, S. Lough and T. Clayton, "Visual timing in hitting an accelerating ball," *Quarterly journal of experimental psychology*, vol. 35, no. A, pp. 333-346, 1983.
- [67] D. N. Lee and P. E. Reddish, "Plumetting gannets: a paradigm of ecological optics," *Nature*, vol. 293, pp. 293-294, 1981.
- [68] D. N. Lee, M. N. Davies, P. R. Green and F. R. v. d. Weel, "Visual control of velocity of approach by pigeons when landing," *Journal of experimental biology*, vol. 180, no. 1, pp. 85-104, 1993.
- [69] D. N. Lee, P. E. Reddish and D. T. Rand, "Aerial docking by hummingbirds," *Naturwissenschaften*, vol. 78, pp. 526-527, 1991.
- [70] D. N. Lee, F. R. v. d. Weel, E. Matejowski and J. D. Pettigrew, "Common principles of guidance by echolocation and vision," *Journal of comparative psychology*, vol. 171, no. A, pp. 563-571, 1992.
- [71] D. N. Lee, "Guiding movement by coupling of Taus," *Ecological Psychology*, vol. 10, no. 3, pp. 221-250, 1998.

- [72] N. Hogan, "An organizing principle for a class of voluntary movements," *Brain*, vol. 105, pp. 331-348, 1984.
- [73] H. Nagasaki, "Assymmetric velocity and acceleration profiles of human arm movement," *Experimental Brain Research*, vol. 74, pp. 319-326, 1989.
- [74] Z. Zhang, S. Zhang, P. Xie and O. Ma, "Bio-Inspired 4D Trajectory Generation for UAS Rapid Point-to-Point Movement," *Journal of Bionic Engineering*, vol. 11, no. 1, pp. 72-81, 2014.
- [75] J. Wann, "Anticipating arrival: Is the Tau margin a specious theory?," *Journal of experimental psychology: human perception and performance*, vol. 22, no. 4, pp. 1031-1048, 1996.
- [76] G. D. Padfield, G. Clark and A. Taghizad, "How long do pilots look forward? Prospective visual guidance in terrain hugging flight," *Journal of the American Helicopter Society*, vol. 52, no. 2, 2007.
- [77] H. A. Lockett, *The Role of Tau Guidance During Decelerative Helicopter Approaches*, Liverpool: University of Liverpool, 2010.
- [78] M. Jump and G. D. Padfield, "Progress in the development of guidance strategies for the landing flare manoeuvre using Tau-based parameters," *Aircraft Engineering and Aerospace Technology*, vol. 78, no. 1, pp. 44-52, 2006.
- [79] M. Jump, *Prospective Sky Guides: Developing Guidelines for Pilot Vision Aids*, Liverpool: University of Liverpool, 2007.
- [80] T. Kai, Y. Shimada and K. Ito, "Timing Control of the Mobile Robot Using Tau-Margin," *Artif Life Robots*, vol. 13, pp. 241-245, 2008.
- [81] R. H. Rasshofer and K. Gresser, "Automotive Radar and Lidar systems for Next Generation Driver Assistance Functions," *Advances in radio science*, vol. 3, pp. 205-209, 2005.
- [82] B. K. Horn, *Robot Vision*, Boston: MIT Press, 1986.

- [83] G. P. Stein, "System and method for detecting obstacles to vehicle motion and determining time to contact therewith using sequences of images". USA Patent 7,113,867 B1, 2006.
- [84] B. K. P. Horn and B. G. Schunck, "Determining Optical Flow," *Artificial Intelligence*, vol. 13, pp. 185-203, 1981.
- [85] B. D. Lucas and T. Kanade, "An Iterative Image Registration Technique with an Application to Stereo Vision," in *Proceeding of Imaging Understanding Workshop*, Pittsburgh, USA, 1981.
- [86] S. Negahdaripour and C.-H. Yu, "A generalized brightness change model for computing optical flow," in *International Conference on Computer Vision*, Berlin, Germany, 1993.
- [87] M. Dailey, "Apparatus and method for self calibrating visual time-to-contact sensor". USA Patent 5,559,695,24, 1996.
- [88] B. K. Horn, Y. Fang and I. Masaki, "Time to Contact Relative to a Planar Surface," in *IEEE Intelligent Vehicles Symposium*, Istanbul, Turkey, 2007.
- [89] H. Zhang and J. Zhao, "Biologically Inspired Vision Based Control Using Featureless Time-to-Contact Estimation," in *IEEE International Conference on Advanced Intelligent Mechatronics*, Banff, Canada, 2016.
- [90] M. Voskuijl, D. J. Walker, B. J. Manimala and A. W. Gubbels, "Simulation of Automatic Helicopter Deck Landings using Nature Inspired Flight Control and Flight Envelope Protection," *The Aeronautical Journal*, vol. 114, no. 1151, pp. 25-34, 2010.
- [91] G. Carico, R. Fang, R. Finch, G. J. W. H. Krijns and K. Long, "Ship/Helicopter Qualification Testing," Research and Technology Organisation of NATO, 2003.

- [92] Z. Yang, Z. Fang and P. Li, "Decentralized 4D Trajectory Generation for UAVs based on Improved Intrinsic Tau Guidance," *International Journal of Advanced Robotics Systems*, vol. 13, no. 3, pp. 1-13, 2016.
- [93] H. Zhang, B. Cheng and J. Zhao, "Extended Tau Theory for Robot Motion Control," in *IEEE International Conference Robotics and Automation (ICRA)*, Singapore, 2017.
- [94] H. Zhang and Z. Jianguo, "An Integrated Unmanned Aerial Vehicle System for Vision Based Control," in *Proceedings of the ASME 2017 Dynamics Systems and Control Conference*, Tysons, Virginia, USA, 2017.
- [95] S. Armendariz, V. M. Becerra and N. Bausch, "Bio-Inspired Autonomous Visual Vertical and Horizontal Control of a Quadrotor Unmanned Aerial Vehicle," *Electronics*, vol. 8, no. 184, pp. 1-20, 2019.
- [96] A. R. Vetrella, I. Sa, M. Popovic, R. Khanna, J. Nieto, G. Fasano, D. Accardo and R. Siegwart, "Improved Tau-Guidance and Vision-Aided Navigation for Robust Autonomous Landing of UAVs," in *11th Conference on Field and Service Robotics*, Zurich, Switzerland, 2017.
- [97] Advanced Rotocraft Technology, "FLIGHTLAB," Advanced Rotocraft Technology, [Online]. Available: <https://www.flightlab.com/flightlab.html>.
- [98] Mathworks, "Simulink," Mathworks, 2019. [Online]. Available: <https://uk.mathworks.com/products/simulink.html>.
- [99] OpenCV, "OpenCV About," OpenCV, 2020. [Online]. Available: <https://opencv.org/about/>. [Accessed 25 02 2020].
- [100] Python Software Foundation, "Python About," Python Software Foundation, 2020. [Online]. Available: <https://www.python.org/about/>. [Accessed 25 02 2020].

- [101 Drone User Manuals, "Review of 3D Robotics X8+," Drone User Manuals, [Online]. Available: <https://www.dronesusermanuals.com/3d-robotics-3dr/review-of-3d-robotics-x8-the-functional-oktokopter-with-great-potential/>. [Accessed 01 09 2019].
- [102 G. D. Wood and D. C. Kennedy, "Simulating Mechanical Systems in Simulink with SimMechanics," The MathWorks, 2003.
- [103 Mathworks, "Simulink 3D Animation," Mathworks, 2019. [Online]. Available: <https://uk.mathworks.com/products/3d-animation.html>.
- [104 J. N. S. Seddon, Basic Helicopter Aerodynamics, Oxford: Blackwell Science, 2002.
- [105 3D Robotics, "3DR Aero," 3D Robotics, [Online]. Available: <https://3dr.com/support/articles/aero/>. [Accessed 01 09 2019].
- [106 L. Lu, "Modelling of DA42 TWIN STAR and RTF Aero UAV in FLIGHTLAB," University of Liverpool Virtual Engineering Centre, 2015, 2015.
- [107 GoodFon, "SH60 Seahawk Wallpaper," GoodFon, [Online]. Available: <https://www.goodfon.com/wallpaper/sh60-seahawk-sh60-vertoliot-letit-poliot.html>. [Accessed 01 09 2019].
- [108 K. Schindler, "MQ-8B Fire Scout," US Navy, [Online]. Available: [https://www.navy.mil/view\\_image.asp?id=108121](https://www.navy.mil/view_image.asp?id=108121). [Accessed 01 09 2019].
- [109 Yamaha, "RMAX G1: An Unmanned Helicopter that Helped Unravel the Mystery of Continent Formation," Yamaha, 2015. [Online]. Available: <https://global.yamaha-motor.com/about/technology/electronic/009/>. [Accessed 01 09 2019].
- [110 RC-Aerobatics, "Align T-Rex 700 E," RC-Aerobatics, [Online]. Available: [http://www.rc-aerobatics.eu/reflex-t-rex-700-e\\_e.html](http://www.rc-aerobatics.eu/reflex-t-rex-700-e_e.html). [Accessed 01 09 2019].

- [111 R. Du Val and C. He, "Validation of the FLIGHTLAB Virtual Engineering Toolset,"  
] *The Aeronautical Journal*, vol. 122, no. 1250, pp. 519-555, 2018.
- [112 Raspberry Pi Foundation, "Raspberry Pi Camera Module V2 Product Page,"  
] Raspberry Pi Foundation, April 2016. [Online]. Available:  
<https://www.raspberrypi.org/products/camera-module-v2/>. [Accessed 13  
August 2019].
- [113 Raspberry Pi Foundation, "Raspberry Pi Camera Module Documentation,"  
] Raspberry Pi Foundation, April 2016. [Online]. Available:  
<https://www.raspberrypi.org/documentation/hardware/camera/>. [Accessed  
13 August 2019].
- [114 FAA, "Airport Markings: ICAO & US Requirements," FAA, 2012.  
]
- [115 DSA, "ShipMo3D," DSA, 2019. [Online]. Available: [https://dsa-  
\] ltd.ca/shipmo3d/overview/](https://dsa-ltd.ca/shipmo3d/overview/).
- [116 H. Bay, T. Tuytelaars and L. Van Gool, "SURF: Speeded Up Robust Features,"  
] *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346-359, 2008.
- [117 I. Culjak, D. Abram, T. Pribanic, H. Dzapo and M. Cifrek, "A brief introduction to  
] OpenCV," in *Proceedings of the 35th international conference MIPRO*, Opatija,  
2012.
- [118 T. Lindeberg, "Detecting Salient Blob-Like Image Structures and Their Scales  
] with a Scale-Space Primal Sketch: A Method for Focus-of-Attention,"  
*International Journal of Computer Vision*, vol. 11, no. 3, pp. 283-318, 1993.
- [119 B. K. Horn, Y. Fang and I. Masaki, "Hierarchical Framework for Direct Gradient-  
] Based Time-to-Contact Estimation," in *IEEE Intelligent Vehicles Symposium*,  
Xi'an, China, 2009.



- [120 G. Ellis, Control System Design Guide: Using Your Computer to Understand and  
] Diagnose Feedback Controllers, Burlington: Elsevier Science, 2012.
- [121 G. D. Padfield, Helicopter Flight Dynamics, Wiley-Blackwell, 1996.  
]
- [122 B. K. P. Horn and E. J. Weldon, "Direct Methods for Recovering Motion,"  
] *International Journal of Computer Vision*, vol. 2, no. 1, pp. 51-76, 1988.
- [123 Z. Zhang, P. Xie and O. Ma, "Bio-Inspired Trajectory Generation for UAV  
] Perching," in *IEEE/ASME International Conference of Advanced Intelligent Mechatronics*, Wollongong, Australia, 2013.
- [124 Mathworks, "Simulink," Mathworks, 2015. [Online]. Available:  
] <http://uk.mathworks.com/products/simulink/?refresh=true>. [Accessed 20 7 2015].
- [125 Department of Defense, "Unmanned Systems Integrated Roadmap,  
] FY2013,2038," 2013.
- [126 T. Merz, S. Duranti and G. Conte, "Autonomous Landing of an Unmanned  
] Helicopter based on Vision and Inertial Sensing," *Experimental Robotics IX*, vol. 21, no. 1, pp. 343-352, 2006.
- [127 H. Das, K. Sridhar and R. Padhi, "Bio-Inspired Landing of a Quadrotor Using  
] Improved State Estimation," *International Federation of Automatic Control*, vol. 51, no. 1, pp. 462-467, 2018.
- [128 D. W. King, A. Bertapelle and C. J. Moses, "UAV failure rate criteria for  
] equivalent level of safety," in *International Helicopter Safety Symposium*, Montreal, Canada, 2005.
- [129 L. Lu, M. Jump and G. D. Padfield, "Development of a Generic Time-To-Contact  
] Pilot Guidance Model," in *44th European Rotorcraft Forum*, Delft, Netherlands, 2017.

- [130 L. Meier, D. Honegger and M. Pollefeys, "PX4: A node-based multithreaded  
] open source robotics framework for deeply embedded platforms," in *IEEE International conference on robotics and automation*, Seattle, USA, 2015.
- [131 L. Meier, D. Honegger and M. Pollefeys, "PX4: A Node-Based Multithreaded  
] Open Source Robotics Framework for Deeply Embedded platforms," in *IEEE International conference on robotics and automation*, Seattle, USA, 2015.
- [132 D. Honegger, L. Meier, P. Tanskanen and M. Pollefeys, "An Open Source and  
] Open Hardware Embedded Metric Optical Flow CMOS Camera for Indoor and Outdoor Applications," in *2013 IEEE International Conference on Robotics and Automation*, Karlsruhe, Germany, 2013.