

Visualizing Rules, Regulations, and Procedures in Ecological Information Systems

Comans, Jan

DOI

[10.4233/uuid:9b3f9bb6-ef1b-41ed-803a-7e7976784b85](https://doi.org/10.4233/uuid:9b3f9bb6-ef1b-41ed-803a-7e7976784b85)

Publication date

2017

Document Version

Final published version

Citation (APA)

Comans, J. (2017). *Visualizing Rules, Regulations, and Procedures in Ecological Information Systems*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:9b3f9bb6-ef1b-41ed-803a-7e7976784b85>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Visualizing Rules, Regulations, and Procedures in Ecological Information Systems

Jan Comans

Visualizing Rules, Regulations, and Procedures in Ecological Information Systems

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben;
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op
donderdag 5 oktober 2017 om 15.00 uur

door

Jan COMANS

Ingenieur Luchtvaart en Ruimtevaart
Technische Universiteit Delft

geboren te Hasselt, België

Dit proefschrift is goedgekeurd door de

promotor: Prof.dr.ir. M. Mulder
copromotor: Dr.ir. M.M. van Paassen

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. ir. M. Mulder,	Technische Universiteit Delft
Dr. ir. M.M. van Paassen,	Technische Universiteit Delft

Onafhankelijke leden:

Prof. dr. R. Curran,	Technische Universiteit Delft
Prof. dr. J. Dankelman,	Technische Universiteit Delft
Prof. dr. J.M. Flach,	Wright State University, USA
Dr.-Ing. B. Korn,	Deutsches Zentrum für Luft- und Raumfahrt
Prof. dr. M. Lind,	Danmarks Tekniske Universitet

Dr. ir. C. Borst heeft als begeleider in belangrijke mate aan de totstandkoming van dit proefschrift bijgedragen.

Printed by: Ipskamp Printing, Enschede

Copyright © 2017 J. Comans

ISBN 978-94-028-0751-6

An electronic version of this document is available at
<http://repository.tudelft.nl>

*To Kamiel,
my grandfather*

Summary

Increasing automation in aviation has played a key role in the rapid development of the aviation industry in the last decades. Without any doubt, it has vastly increased safety and efficiency. Automation will continue to play an important role in the future. Consensus exists that this does not necessarily mean only increasing the level of automation, it also means developing automation in such a way that the human operator keeps a central role in the system.

One of the main lessons that the aviation community and beyond has learned in the past, is that in the process of moving a task from the human operator to an automated system, mistakes and accidents can indeed be prevented, but usually new opportunities for mistakes and accidents are introduced. These archetypical 'ironies of automation' are omnipresent in aviation.

Complacency and boredom are just a few problems that emerge in highly automated systems. One of the challenges in future automation design will therefore be to balance new increasing levels of automation and human involvement. A significant part of facing this challenge will be in the domain of human-machine interfaces and human-centered automation.

Ecological Interface Design (EID) aims at supporting human operators in complex socio-technical system domains that, at times, require elaborate operating procedures. Experimenting with a number of experimental aviation interfaces in human-in-the-loop studies has shown that they often provide better support for pilots in difficult situations. Ecological interfaces aim to visualize a complete overview of the possibilities and constraints of the system. However, this can sometimes invites operators to operate close to these constraints.

Operating close to boundaries can pose risks when pilots are close to physical constraints that could lead to severe incidents or accidents when being crossed. This possibility for risk migration is not necessarily the result of the EID methodology as it is sometimes claimed. Visualizing rules, regulations, and procedures could prevent limit seeking behavior, which lead to the following problem statement: *Can we, by clearly visualizing rules, regulations, and procedures, create ecological*

displays that lead to safer overall human-machine system performance?

This problem statement indicates that the display design methodology needs a way to explicitly incorporate rules, regulations, and procedures. Furthermore, an analysis is needed of the impact of such an updated methodology. This leads to the three main research questions in this thesis:

1. *How do rules, regulations, and procedures fit into the EID framework?*
2. *Will pilots be able to distinguish between physical constraints, and constraints introduced by rules, regulations, and procedures when they are visualized in the interface?*
3. *Will pilots make better decisions based on the additional information?*

The first research question addresses the foundation of an EID display, the Work Domain Analysis, leading to the Abstraction Hierarchy (AH). The AH represents the work domain at different levels of abstractions. Functions at adjacent levels are linked through means-end links. It will represent the *functional* structure, which is characterized by the goals and the constraints limiting the dynamic behavior of the system. The constraints on the system under study can be divided in two categories.

The first category encompasses the *causal* constraints. Constraints that are governed by the laws of nature and the physical processes involved. These constraints enforce the physical boundaries on the system. Causal constraints can either be violable or inviolable. Violable constraints represent boundaries that can be crossed and that will lead to a severe degradation of performance, or even accidents. An example of a violable constraint could be an obstruction on a road that a car could drive into. Inviolable constraints represent limits that represent asymptotic behavior. They will not necessarily lead to system degradation, but they pose limits on what can be achieved. An example could be the maximum velocity of a car. It is not possible to drive any faster, but it does not impact the system's integrity.

The second category of constraints is driven by the actors' intentions, values, rules, regulations, etc. They are referred to as *intentional* constraints. Intentional constraints do not represent physical boundaries. They are in place to shape behavior. This can be driven by numerous goals, like for example, safety, efficiency, profit, etc. In the context of this thesis, we focus on the intentional constraints stemming from rules, regulations, and procedures related to safety. An example of such an intentional constraint could be the speed limit imposed on cars. A car might be perfectly capable to exceed the speed limit, but adhering to the speed limit will increase the safety of all traffic participants.

The implications of both types of constraints are, however, very different. Causal constraints dictate how one *can* act, intentional constraints represent how one *should* or *would want to* act. This difference can be crucial in unexpected and emergency situations. As an example, when flying an aircraft, avoiding a tall

Summary

structure is important under any circumstances, but the goal of avoiding a noise sensitive area can in principle be ignored when entering that area would make the resolution of an emergency safer or easier. In the existing EID displays, there is no explicit distinction between causal and intentional constraints. Some displays only model causal constraints, others mix both causal and intentional constraints without making the distinction clear to the operator.

Based on the above distinction, two existing experimental flight deck displays have been analyzed in this thesis, to identify the types of constraints that have been incorporated, and also to investigate how these displays could be adapted to incorporate an explicit split between causal and intentional constraints.

An Ecological Synthetic Vision Display (ESVD) was used as the baseline display of the first case study. This display consists of a traditional Primary Flight Display on top of a three-dimensional synthetic view of the terrain with an indication of the available climb performance. This climb performance indication can be easily related to the synthetic terrain, to quickly assess whether the aircraft is indeed able to climb over the terrain.

The most important elements for the terrain avoidance task, for which the display is designed, are: the virtual terrain, the maneuver distance indicator, and the minimum required bank angle. These three elements are all based on physical constraints. The case study focused on incorporating both horizontal and vertical separation criteria in the Enhanced Vertical Situation Display (EVSD). Based on these criteria, a visualization of the separation buffers was added to the display in such a way that pilots can see both the required safety margins as well as the underlying causal constraints.

The second case study used an EVSD as its baseline display. The EVSD shows a side view of the ownship, together with upcoming terrain and traffic. On top of this presentation, a performance envelope is drawn, representing the limits of ownship performance. In this performance envelope, triangular conflict zones representing velocities that lead to a traffic conflict are drawn together with a terrain angle line, which indicates the climb angle required to clear the upcoming terrain with a specific minimum terrain clearance.

The analysis shows that the current EVSD display design contains a mix of causal and intentional constraints, without making the distinction clear. The traffic constraint is represented by a filled polygon visualizing the safety zone around an intruding aircraft. This represents a pure intentional constraint and in the way it is visualized obscures the underlying causal constraint. The terrain constraint is represented by the virtual terrain combined with a terrain line indicating the minimum required height above the terrain. This mixes both the causal and intentional constraint, but in two different representations that are difficult to relate to each other. Furthermore, the filled polygons representing the traffic are much more present than the terrain line.

Based on this analysis, an attempt is presented to make a clear split in a causal and an intentional part with a uniform visualization, yielding an updated display.

After the theoretical analysis, the thesis continues with two chapters describing an experimental pilot-in-the-loop evaluation of the two re-designed ecological displays.

In the first experiment, we compared the additions in the novel Intentional Synthetic Vision Display (ISVD) with the baseline ESVD. Sixteen commercial pilots were, after a training period, presented with four measurement scenarios where they were put into a terrain conflict. They performed their runs with a baseline ESVD and with the improved ISVD.

Analysis of the results shows an increase in minimum terrain clearance during their evasive maneuvers. Visualizing the intentional constraint resulted in better compliance with the intentional constraint. Furthermore, pilots were able to intuitively use the additional information. In conflicts that were easy to solve, they used the display to fine-tune their performance to minimize deviations and fuel consumption, while satisfying the minimum safe altitude constraint. In more difficult scenarios, with less options, pilots were able to directly use the representation of the intentional constraint to execute their chosen resolution strategy. For example, after deciding a turn is required, the pilot can immediately see which direction will be the most beneficial, or when the turn can be stopped.

In the second experiment, we compared the EVSD baseline display with the novel augmented Intentional Vertical Situation Display (IVSD) visualizing both causal and intentional constraints. Sixteen subjects, of which eight experienced pilots and eight inexperienced students, were divided in two groups. After training, one group used the baseline display to resolve four mixed terrain and traffic conflicts. The other group used the IVSD to resolve the same conflicts.

The experiment did not show an anticipated increased terrain clearance, nor did it show an increase in protected zone violations. Furthermore, the experiment did not show a clear difference in strategy between the experienced and inexperienced subjects. However, the trajectories flown by the subjects do show a reduced variation in performance when comparing the baseline display with the IVSD display, which indicates that subjects using the IVSD were indeed more aware of the constraints and were able to execute their strategy more precisely.

To conclude, this thesis showed how rules, regulations, and procedures can be fitted into the EID by properly identifying them and representing them in the AH. This thesis proposes a distinct split in the AH between the causal domain and the intentional domain. In this way, there is a clear distinction between the rigid causal constraints and the more flexible intentional constraints. This in turn allows for the design of an interface that presents the same distinction which can help pilots in prioritizing constraints, especially in unexpected circumstances.

The ISVD experiments shows that pilots were able to benefit from a clear visualization of intentional constraints. The IVSD experiment, on the other hand, makes it difficult to claim a clear benefit of an explicit visualization of intentional constraints. However, this is more likely a result of insufficient training and scenario design than of the visualization itself. Therefore, the ISVD results indicate

Summary

that explicitly visualizing intentional constraints allow pilots to distinguish between physical constraints and constraints originating from rules, regulations, and procedures, but more validation is still required.

The ISVD experiment indicated that when pilots are able to distinguish between causal and intentional constraints their decision making improves, while the IVSD—due to training and scenario issues—failed to show a clear improvement. A single experiment is not sufficient to conclusively confirm a hypothesis, but it does show that future validation experiments could be worthwhile.

A number of potential reasons were identified as to why the IVSD experiment failed to show a clear difference between the baseline condition and the addition of explicit intentional constraints. First and foremost, subjects lacked familiarity with the basic principles of the interface due to insufficient training. Next to this, the scenarios failed to put subjects into the kinds of situations that would truly benefit from the improved visualization. Finally, a greater number of subjects would be required.

Contents

Summary	i
1 Introduction	1
1.1 Background	1
1.2 Problem Statement	3
1.3 Extending the Work Domain Analysis	4
1.4 Research Approach	7
1.4.1 Display Analysis	8
1.4.2 Experimental Evaluation	8
1.5 Research Scope and Assumptions	10
1.6 Thesis Outline	11
2 Theoretical Foundations	13
2.1 Safety in Aviation	13
2.2 Cognitive Systems Engineering	17
2.2.1 The Abstraction Hierarchy	17
2.2.2 The Skills, Rules, Knowledge Framework	20
2.3 Intentional Constraints	21
2.3.1 Work Domain Constraints	21
2.3.2 The Causal-Intentional Abstraction Hierarchy	25
2.4 Implications for Design	27
3 Case Studies	31
3.1 Case One: Synthetic Vision Display	31
3.1.1 Purpose and Goals	32
3.1.2 Current Ecological Interface	32
3.1.3 Extending the Work Domain Analysis	40
3.1.4 Augmented Interface	42
3.2 Case Two: Vertical Situation Display	48

3.2.1	Purpose and Goals	48
3.2.2	Current Ecological Interface	49
3.2.3	Work Domain Analysis	54
3.2.4	Augmented Interface	56
4	Experimental Evaluation of an Intentional Synthetic Vision Display	61
4.1	Background	61
4.2	Intentional Constraints	63
4.3	Experiment Method	65
4.3.1	Subjects	66
4.3.2	Apparatus	66
4.3.3	Display	67
4.3.4	Scenario	67
4.3.5	Independent Variables	69
4.3.6	Dependent Measures	70
4.3.7	Procedure	70
4.4	Results	71
4.4.1	Minimum Clearance	71
4.4.2	Clearance Violations	72
4.4.3	Final Altitude	73
4.4.4	Strategy Adaptation	73
4.4.5	Questionnaire Data	75
4.5	Discussion	76
4.6	Conclusion	78
5	Experimental Evaluation of an Intentional Vertical Situation Display	79
5.1	Background and Motivation	80
5.1.1	The Enhanced Vertical Situation Display	81
5.1.2	Visualizing Intentional Constraints	84
5.2	Experiment Displays	85
5.2.1	Primary Flight Display	85
5.2.2	The Baseline VSD	86
5.2.3	The Intentional VSD	89
5.3	Experiment Method	89
5.3.1	Subjects	89
5.3.2	Apparatus	89
5.3.3	Scenario	91
5.3.4	Independent Variables	94
5.3.5	Procedure	95
5.3.6	Dependent Measures	95
5.3.7	Hypotheses	96
5.4	Results	97
5.4.1	Strategy Observations	97
5.4.2	Terrain Clearance	109

CONTENTS

- 5.4.3 Minimum Distance to Intruder 111
- 5.4.4 Protected Zone Violations 111
- 5.4.5 Debriefing 113
- 5.5 Discussion 113
- 5.6 Conclusion 119

- 6 Discussion & Conclusions 121**
- 6.1 Discussion 121
 - 6.1.1 On the Research Questions 121
 - 6.1.2 On Related Topics 127
- 6.2 Conclusions 133
- 6.3 Recommendations 133

- A ISVD Briefing 135**
- B ISVD Questionnaire 141**
- C IVSD Briefing 143**

- Bibliography 149**
- Glossary 155**
- Samenvatting 157**
- Acknowledgements 163**
- Curriculum Vitae 165**

Introduction

1.1 Background

In aviation, an ever increasing degree of automation has played a key role in the rapid development of the industry in the last decades (Billings, 1996b; Lovesey, 1977). Even though this has resulted in a very reliable and safe air transportation system, predictions about required capacity and operations show that the current system will need to be improved further to cope with future requirements (EUROCONTROL, 2010). Automation will play an important role in dealing with these challenges (JPDO, 2011). This does not only mean increasing the level of automation, but to keep the system feasible, also developing automation in such a way that the human keeps a central role in the system (EUROCONTROL, 2006).

One of the lessons learned by the aviation community, is that in the process of moving a task from the human operator to an automated system, old mistakes and accidents can be prevented, but usually new opportunities for mistakes and accidents are commonly introduced through the interaction with the automation (Woods, Johannesen, Cook, & Sarter, 1994). Taking tasks out of the pilots' hands makes them less involved actually flying the aircraft, potentially depriving them from required information in situations that the automation can not solve and where control is handed back to the pilot (Sarter & Woods, 1995).

For a large part, these problems can be attributed to the fact that humans are not well suited for a monitoring task and therefore some accidents become inevitable (Bainbridge, 1983; Perrow, 1984). One of the challenges in future automation design will therefore be to balance new and increasing levels of automation with sufficient human involvement. A significant part of facing this challenge will be in the domain of human-machine interfaces and human-centered automation.

Over the past decades, an approach to interface design has emerged that has a pronounced appreciation for the *human contribution* to technical systems: Ecological Interface Design (EID), introduced by Vicente and Rasmussen (Vicente &

Rasmussen, 1990, 1992; Burns & Hajdukiewicz, 2004; Bennett & Flach, 2011). This perspective cautions that humans still are an essential component in technical systems, as they can bring adaptivity and creativity to a system at a level that computers are not yet able to provide. Such abilities are occasionally seen when a human performs *heroic* acts and recoveries, as in the Hudson River landing where captain Sullenberger successfully landed an Airbus A320 on the Hudson River after losing thrust in both engines due to multiple bird strikes (Reason, 2008; NTSB, 2009).

The EID perspective represents a fundamentally new view on how humans work with technology. Rather than striving exclusively to replace human weaknesses with technical systems, the emphasis shifts to explore ways that technology can facilitate human adaptivity and flexibility to help operators cope with unforeseen events. This recognizes that in complex domains, there will always be a potential for problems that cannot be anticipated in the design of automatic control systems. Thus, the creative human expert becomes an important resource for dealing with unanticipated variability.

The key implication for design is that this requires that decision support systems work together with humans in order for the system to respond robustly to the complex work environment (Billings, 1996a). EID's answer to this challenge is to promote coordination between humans and automatic systems through interface representations that are grounded in constraints that reflect the deeper structure of the work domain demands. As such, ecological information systems present the space of possibilities by explicitly mapping the system's means-ends relationships and the operating limits on the interface. The system user can then use this presentation to find any solution that lies within the boundaries of safe performance. However, the explicit presentation of operating limits also raises the concern that it may invite operators to seek out these limits and promote migration of activities towards the boundaries of safe system performance, a phenomenon that was termed *risk migration* by Rasmussen (1997).

This 'limit seeking' behavior has indeed been noted in previous experiments in aviation with several ecological flight deck displays. For example, Borst, Mulder, and van Paassen (2010) and Rijnveld, Borst, Mulder, and van Paassen (2010) evaluated ecological synthetic vision terrain displays. In those experiments a number of pilots flew considerably closer to the terrain in conditions without an ecological display. From a physical perspective there is no problem, as there is no difference in clearing terrain by 50 *ft* or by 2,000 *ft*. From a safety perspective, however, the lower clearance is much less favorable. When a pilot only uses a minimal margin there is no margin left to deal with unexpected events like, e.g., non-optimal performance or unexpected obstacles. In aviation, rules, regulations, and procedures are our primary means to—among other things—ensure that pilots maintain a safe margin above the physical limitations of the aircraft. They are a set of *intentional* constraints.

Using terrain clearance as an example, regulations can specify that a pilot needs to maintain a clearance of 500 *ft* at all times. In nominal situations, the

pilot will need to comply and will have a safety margin when unexpected events happen. If on the other hand, the pilot would find himself in a situation where the only way out is to clear the terrain with less than 500 *ft*, he is allowed to disregard the regulations and use a lower clearance. In such a situation, avoiding the *physical* terrain becomes the prime focus.

1.2 Problem Statement

A clear visualization of constraint boundaries could invite pilots to operate closer to those boundaries than when using conventional displays. This can pose risks when pilots are close to physical constraints that could lead to severe incidents or accidents when being crossed. This possibility for risk migration is, however, not necessarily the result of the EID methodology as is sometimes claimed (Borst, Flach, & Ellerbroek, 2014). Usually rules, regulations, and procedures are in place to prevent pilots from operating too close to a limit. This suggests that adding, or clearly visualizing these rules, regulations, and procedures could prevent limit seeking behavior, which leads to the following problem statement:

Can we, by clearly visualizing rules, regulations, and procedures, create ecological displays that lead to safer overall human-machine system performance?

This problem statement indicates that the display design methodology needs a way to explicitly incorporate rules, regulations, and procedures. Furthermore, an analysis is needed of the impact of such an updated methodology. This leads to three questions that will form the main research questions in this thesis:

- 1. How do rules, regulations, and procedures fit into the EID framework?*
- 2. Will pilots be able to distinguish between physical constraints, and constraints introduced by rules, regulations, and procedures when they are visualized in the interface?*
- 3. Will pilots make better decisions based on the additional information?*

Before addressing the research approach, the next section will provide a brief introduction on the fundamental principles of the EID framework that are closely related to the first research question shown above.

1.3 Extending the Work Domain Analysis

Interface design needs to answer two fundamental questions. *What* is the information content that needs to be presented to the operator and *how* should this information be presented on the interface? Vicente and Rasmussen presented the EID framework to answer these questions for complex human-machine systems. EID uses the Abstraction Hierarchy (AH) to represent the work domain which will represent the information content of the display. They argue that the major cause for life-threatening accidents is related to unexpected events that could not be foreseen at the (interface) design for which there is limited support on traditional interfaces (Vicente & Rasmussen, 1992). Through the AH functional relations will be explicitly represented. This *should* make it possible to determine when constraints are broken.

The AH describes the constraints of the work domain at various levels of abstraction, and therefore determines what information needs to be presented. The higher levels of abstraction describe the governing principles for the work domain and can help an operator determine which lower level constraints are broken and what the implications for the system are. This information is considered vital for dealing with unexpected events.

For the presentation of the information, the EID framework relies on the Skills, Rules, Knowledge (SRK) taxonomy (Rasmussen, 1983). This taxonomy describes the three ways in which information is interpreted—in terms of signals, signs, or symbols—and how the way in which information is interpreted determines which level of cognitive control is activated—Skill-Based Behavior (SBB), Rule-Based Behavior (RBB), or Knowledge-Based Behavior (KBB). All levels of the SRK taxonomy should be supported by the interface while at the same time cognitive control is not forced to be at a level higher than demanded by the task. In this way, the SRK helps the designer in determining how the information needs to be presented.

The rest of this section will focus on the Work Domain Analysis (WDA), the technique used to obtain an AH for the work environment under study. The reason for this is that if we want to address the problem statement—including rules, regulations, and procedures in ecological information aids—we chose to start by extending the WDA to include the relevant rules, regulations, and procedures.

Within a work domain we can distinguish between causal constraints, governed by the laws of nature and the physical processes involved, and intentional constraints, that are governed by actors' intentions, values, procedures, rules, etc. Rasmussen, Pejtersen, and Goodstein (1994) describe a taxonomy for classifying work domains based on the relative degree of intentional and causal constraint (Naikar, Hopcroft, & Moylan, 2005).

The types of work domain can be seen as a continuum, shown in Figure 1.1, with on the right hand side purely causal systems, that are only governed by the laws of nature and on the left hand side purely intentional systems, that are only

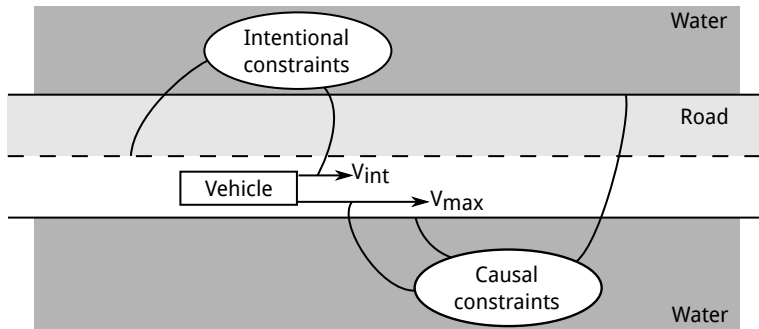


Figure 1.2: An example of causal and intentional constraints in the vehicle domain.

imum velocity is higher than the maximum speed restriction, the operator can decide to go faster than allowed. The difference with violating a causal constraint is that violating an intentional constraint *might* be a direct safety hazard, but it does not have to be. The same holds for the division in the middle of the road, vehicles should stay on the right hand side of the road—in the Netherlands and most other countries in the world—but are allowed to temporarily drive on the left hand side to overtake another vehicle. Such a maneuver on the left hand side of the road is a safety risk, especially when the operator is not paying enough attention, but will not necessarily lead to an accident.

In the majority of cases, intentional constraints will impose stricter constraints on top of the underlying causal constraints. This is illustrated in Figure 1.3. In these cases, the available work space will be further limited by the intentional constraints. It is possible for intentional constraints to be less strict than causal constraints. The maximum speed for a road, for example, might be higher than the maximum speed that can be achieved by the vehicle. In this case the causal constraint will be the relevant constraint.

This distinction between causal and intentional constraints will play an important role when looking back at the problem statement. Rules, regulations, and procedures are all intentional constraints imposed by actors in the work domain to shape behavior. By extending the work domain analysis to include these intentional constraints they can, in principle, be identified and also mapped onto the interface. With this information, operators should be able to clearly distinguish, however, between behavior required by rules, regulations, and procedures and behavior required to satisfy physical constraints.

The example above indicates some key differences in how an operator can treat the constraints. Causal constraints should always be respected, intentional constraints can be crossed if necessary, i.e., when the situation requires this. Only showing physical constraints may lead to unsafe situations when the operators'

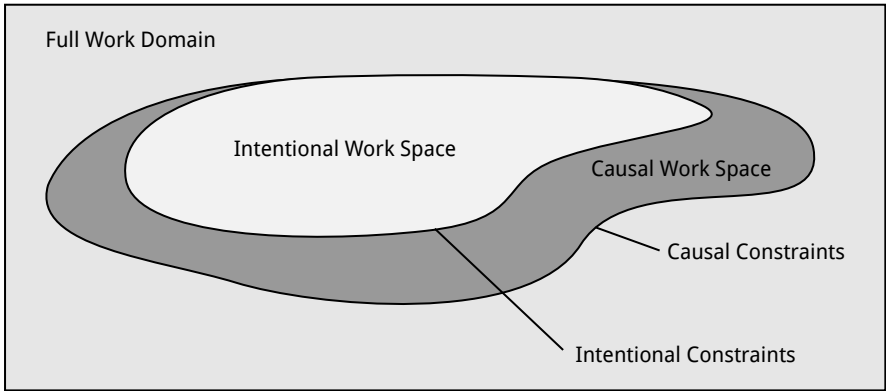


Figure 1.3: Relation between causal and intentional constraints.

actions migrate to these boundaries. On the other hand, when causal and intentional constraints are lumped together and are presented without distinction, the control space for the operator will be shown as being narrower than it actually is. Especially under unexpected events this might lead to situations where the display shows no way out of a situation as rules, regulations, etc., prohibit it while in reality there still is a solution within the physical constraints. Hence, the key point to improve both performance and safety could be to visualize both sets of constraints, while keeping the distinction between causal and intentional constraints clear at all times.

1.4 Research Approach

This thesis aims to address the research questions formulated in Section 1.2. As explained in Section 1.3, intentional constraints map to rules, regulations, and procedures. Visualizing both causal and intentional constraints can provide optimal support for pilots assessing the opportunities and risks associated with specific choices.

This leads to a three-phased approach. First the theoretical foundations of the explicit split between causal and intentional constraints will need to be described. This will address the first research question: *How do rules, regulations, and procedures fit into the EID framework?*

Based on these principles, two existing ecological interfaces that have been developed at the TU Delft will be analyzed to identify possible intentional constraints that were left out or that have, implicitly, already been included and been presented as intentional constraints. Based on the analysis, and an extended

WDA, improved displays will be designed that take into account an explicit split between causal and intentional constraints.

Finally, the two improved interfaces will be used in two validation experiments that will attempt to assess the benefits and drawbacks of differentiating between both types of constraints. The experiments will attempt to answer both the second and third research question: (a) *Will pilots be able to distinguish between physical constraints, and constraints introduced by rules, regulations, and procedures when they are visualized in the interface?* (b) *Will pilots make better decisions based on the additional information?*

1.4.1 Display Analysis

In a number of current aviation displays based on EID, intentional constraints are not taken into account (Borst et al., 2010) or there is no clear distinction between causal and intentional constraints (Rijneveld et al., 2010; Ellerbroek, Mulder, & van Paassen, 2011). As discussed in Section 1.3, visualizing this distinction may be important to improve the pilots' awareness of the applicable causal and intentional constraints.

The Vertical Situation Display (VSD) is a display that assists pilots in managing their vertical profile. The ecological version of this display introduced by Rijneveld et al. (2010) visualizes the constraints imposed by terrain and traffic and is shown in Figure 1.4. In this VSD an intentional constraint is implicitly introduced by including the Minimum Safe Altitude (MSA). The traffic constraint on the other hand is presented as a single causal constraint. This leaves room to properly separate the causal and intentional constraints.

The Synthetic Vision Display (SVD), which is a perspective Primary Flight Display (PFD) has also been adapted into an ecological version (Borst et al., 2010) as shown in Figure 1.5. The ecological additions provide pilots with a better understanding of the relationship between the performance of their aircraft and the performance required to clear terrain obstacles. The display was shown to be successful in preventing terrain collisions, but the visual representation makes it difficult to judge distances and altitude. To improve safety, pilots should be aware of the margins they should maintain with respect to the terrain to ensure safe crossing. Treating this MSA as an intentional constraint should help pilots in dealing with this margin.

1.4.2 Experimental Evaluation

Two experiments will be performed to evaluate the effects of explicitly visualizing intentional constraints in practical use. The first experiment is performed with the SVD as a baseline display in a terrain avoidance task. Pilots are put in situations where they are below surrounding terrain and the only viable option is to climb over it. The SVD is a valuable tool in these situations because it clearly shows the relation between the available climb performance of the aircraft and the climb

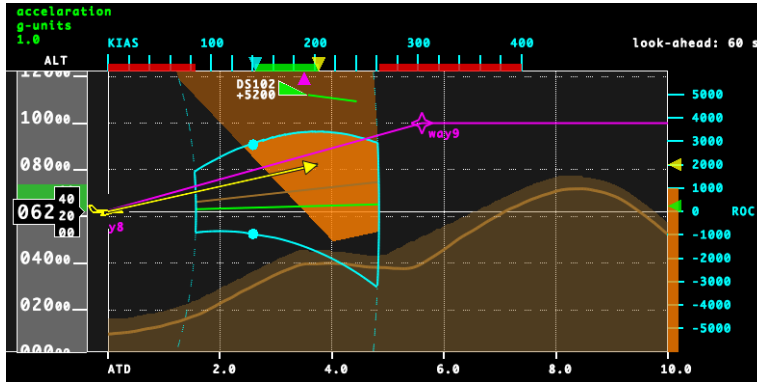


Figure 1.4: The ecological Vertical Situation Display (VSD)(Rijneveld et al., 2010).



Figure 1.5: The ecological Synthetic Vision Display (SVD)(Borst et al., 2010).

performance required to clear the terrain. While this is sufficient to avoid terrain collisions, from a safety point of view it is preferable to build in some margin with respect to the terrain. This margin is added to the display as an intentional constraint. It should be adhered to if possible, but can be ignored if deemed necessary by the pilot. The goal of this experiment is to investigate how different pilots interpret and use the intentional addition, and if safety can improve by visualizing intentional constraints.

The second experiment will look at how pilots manage multiple constraints at the same time, and how the visualization of the constraints influences their decision making. A VSD will be used to present scenarios with a simultaneous traffic and terrain constraint. The baseline version of this display will be based on the VSD that has been used in previous experiment, with a non-uniform presentation of the constraints. In the augmented version, both traffic and terrain constraints will be presented in a uniform way that attempts to show boundaries of equal risk. The goal of this experiment will be to investigate if pilots will be better informed about the difference between the causal and the intentional part of the constraints, and about the relationship between the risks posed by both constraints.

1.5 Research Scope and Assumptions

A number of assumptions have been made to limit the scope of the research:

Aviation domain Constraint-based displays can have applications in a number of different fields. The work in this thesis will focus mainly on the aviation domain and more specifically on the aircraft cockpit automation and interface.

Low levels of automation Because of the fundamental nature of this research, it will focus on simple dynamics and relatively simple constraints. The focus is on manual pilot control tasks without any automation assistance. In this perspective, the work is most applicable in a general aviation context, but can be a stepping stone towards implementing these systems in a commercial aviation setting.

Existing ecological interfaces No new interfaces will be developed. The work is focused on analyzing and improving existing displays developed at the Control & Simulation section.

No technological constraints The designs of the EID displays assume that there is no limit to available processing power and that all required data are available all the time without errors. Practical limitations like navigation inaccuracies, signal noise, etc., are omitted.

Intentional constraints limited to safety Intentional constraints encompass a whole range of rules, regulations, and procedures. This thesis focuses only on risk perception and safety. Therefore the scope of the WDAs and visualizations will be limited to constraints related to safety.

Exceptional situations The work will focus on exceptional situations that are not likely to occur under normal procedures. This is mainly because it is hypothesized that the strength of EID displays lies in supporting operators during unforeseen situations that are usually very rare.

1.6 Thesis Outline

This section will provide a brief overview of the structure of this thesis. A graphical overview is given in Figure 1.6.

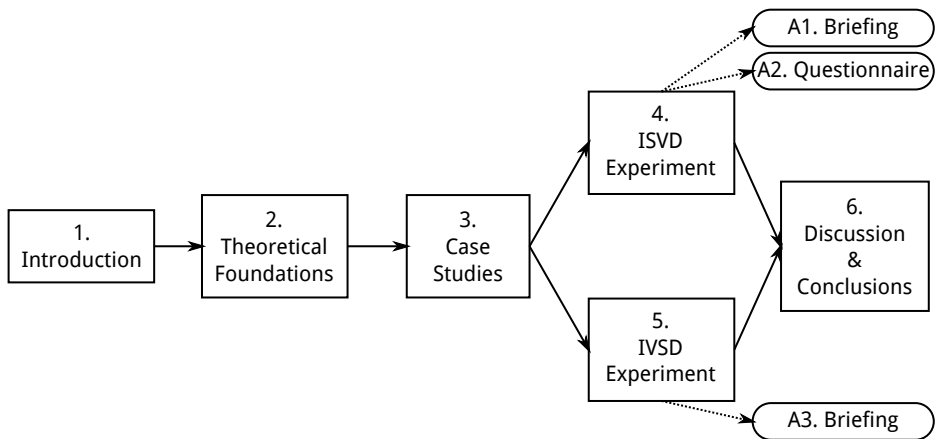


Figure 1.6: The structure of this thesis. In this figure, ISVD stands for Intentional Synthetic Vision Display and IVSD stands for Intentional Vertical Situation Display.

Chapter 2: Theoretical Foundations

This chapter will cover the theoretical foundations of the WDA taking into account intentional constraints and the resulting implications for the interface mapping. This chapter will also describe how the visualization of both causal and intentional constraints is hypothesized to affect safety and risk perception.

Chapter 3: Case Studies

Two case studies will be presented, following the concepts put forward in the theoretical foundations in Chapter 2. Each case study will give an overview of the AH and visualization of original display. After the original constraints are classified in causal and intentional constraints, the scope of the work domain will be re-evaluated to include a number of rules, regulations, and procedures. Based on the new AH, the interface will be re-evaluated and adaptations and improvements will be discussed.

Chapter 4: Experimental Evaluation of an Intentional Synthetic Vision Display (ISVD)

This chapter will describe an experiment to investigate how pilots deal with the addition of an intentional constraint to an existing EID. Based on the analysis in Chapter 3 one intentional constraint is added to an ecological SVD. The goal of this chapter is to demonstrate that making intentional constraints explicit can help pilots to comply with rules, regulations, and procedures if possible and make informed decisions on how to ignore them when they are unable to comply.

Chapter 5: Experimental Evaluation of an Intentional Vertical Situation Display (IVSD)

Similar to Chapter 4, this chapter presents an experiment with an ecological display augmented with explicit intentional constraints, but this time for a VSD. In the VSD information is provided about terrain and traffic. In a previous experiment, pilots showed a tendency to fly relatively close to the terrain. A reason for this could be that the way in which the constraints are presented puts more emphasis on the traffic constraint pushing pilots towards the terrain. An experiment is described in which both constraints are shown based on the principles from Chapter 3, to evaluate whether pilot strategy and safety improves by making the distinction between causal and intentional constraints more explicit.

Chapter 6: Discussion and Conclusions

This chapter will wrap up the theoretical and experimental work, return to the research question and goals, and provide recommendations for future research. The main conclusions of this thesis will be stated.

Theoretical Foundations

This chapter will provide the theoretical foundation for the analysis and experiments in the remainder of this thesis. It will start with an overview of some topics on aviation safety that are relevant for the rest of the chapter. The next section will provide an overview of the Cognitive Systems Engineering (CSE) framework that lies at the basis of this thesis. The final section will explore what the implications of this new approach will be on the design of the actual interface.

2.1 Safety in Aviation

In just over one century, aviation has developed into one of the most efficient and safest modes of long distance transportation. The first five decades of this century were characterized mainly by structural improvements, creating more reliable and flyable aircraft. Around 1950, the mass transportation era began with the introduction of the de Havilland Comet. While the number of flights and number of passengers transported increased year by year, aviation safety also steadily increased. The number of fatalities decreased while the air traffic volume increased. Figure 2.1 shows the evolution of fatal accidents throughout the last fifty years.

The Comet also revolutionized the aviation industry in a tragic way. Only one year into service, three Comets broke up in mid-air due to metal fatigue, a phenomenon that was not well known at that time (Withey, 1997). The fact that a well-designed and well-tested aircraft could just break up in mid air sent a shock wave through the industry and formed the basis for rigorous accident investigations and the development of an unprecedented safety culture. Lessons were learned with every accident or serious incident, and rules, regulations, and procedures were and are still continuously fine-tuned to improve safety (European Transport Safety Council, 2001).

From the pilot's perspective, at the operational side of modern aviation systems rules, regulations, and procedures are the backbone of this safe system.

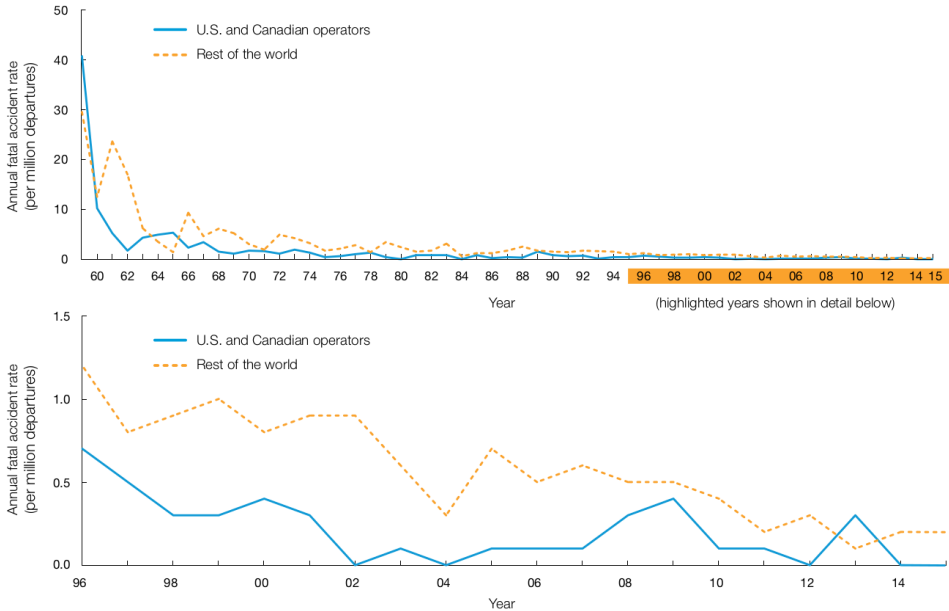


Figure 2.1: Fatal Accidents - Worldwide Commercial Jet Fleet - 1959 through 2015 (Boeing Commercial Airplanes, 2016).

Pilots have some leeway to use their knowledge, skills, and insight to operate the aircraft, also known as ‘good airmanship’ (Langewiesche, 1944). But they are guided by rules, regulations, and procedures to ensure efficiency and safety. For the majority of the situations, this works very well. Problems can arise, however, in off-normal situations. Rules, regulations, and procedures are designed for the *known* off-nominal conditions, not for the *unknown* off-nominal conditions. When unexpected or unanticipated situations arise, procedures might not be available or even be wrong. General experience, incidents, or accidents can lead to new insights and additional or improved rules, regulations, and procedures, but there is no way to explicitly design for the *unknown*.

The human operator plays a critical role in these unanticipated off-normal situations. System knowledge combined with creativity and flexibility is of great importance to deal with unanticipated problems. In unanticipated situations, procedural information becomes less relevant and pilots require more information about relevant limitations for that situation. Two examples illustrate these points:

Example One: Prioritizing Limitations All rules, regulations, and procedures have their reasons, usually rooted in experience in dealing with physical limita-

tions. They do not, however, necessarily carry the same weight. Some leave no room for errors or interpretation. As an example, consider the Decision Height, the height at which the pilots must have a visual reference to the runway when landing the aircraft. If not, they are required to initiate a Go-Around maneuver. One of the main factors in determining the decision height is the spool up time of an average aircraft engine. If the decision to go around is postponed until after reaching the decision height, there will be insufficient time left for the engines to deliver the desired (high) level of thrust, leading to a situation with a high risk of colliding with. This makes the decision height an important and strict constraint that should always be adhered to.

Other rules, regulations, and procedures might be less critical. One of the rules employed by Air Traffic Control (ATC), for instance, is that aircraft should be horizontally spaced by 3 *NM* to 5 *NM*, depending on the phase of flight. One of the driving factors for this requirement is the accuracy of radar surveillance equipment available to the Air Traffic Controller (ATCo) to ensure adequate separation (Nolan, 2010). This results in a no-go area with a radius around an aircraft which is huge compared to the physical size of the aircraft. The chances of actually colliding with another aircraft are relatively small when violating this separation constraint, especially when compared with the risk of ignoring decision height.

Even though pilots have to comply with all rules, regulations, and procedures, they bear ultimate responsibility for the safety of the flight and its passengers. This implies that they are allowed to deviate from rules, regulations, and procedures if this is deemed absolutely necessary. While their ability to do so is critical in truly unanticipated situations, this does not always have the desired effect. In 2010, a Royal Air Maroc Boeing 737 suffered a bird strike on takeoff from Amsterdam (The Dutch Safety Board, 2011). The pilots misinterpreted the condition of their aircraft and ignored a number of procedures. Instead of climbing straight out to reach a safe altitude, the pilot immediately initiated a turn and decided not to retract the landing gear. As a result, the aircraft flew over a number of built-up areas at an altitude much lower than some of the obstacles in the area, creating a very dangerous situation. In this case, not following the procedures proved to be the wrong decision.

An example of the opposite case—not complying with the procedures—that proved to be successful was US Airways Flight 1549 that had a dual engine failure and was forced to ditch in the Hudson river. According to the National Transportation Safety Board (NTSB) report, Airbus designed the *Engine Dual Failure Checklist* for the occurrence of a dual-engine failure above 20,000 *ft*, and did not consider developing a dual-engine failure checklist for use at low altitude (NTSB, 2009). The first part of the checklist is focused on attempting to restart the engines. The third item on the checklist—establishing an optimal relight speed of 300 *kts*—was already inappropriate for the situation, they would lose too much altitude. They skipped this item and continued on to find an item that required them to wait for 30 *s*. By this time they were only at an altitude of about 1,000 *ft* which made it impossible to continue the checklist. From this point on they aban-

done the checklist/procedure and used their best judgment to guide the aircraft towards the Hudson river for landing. Right before ditching in the water, the pilots lowered the flaps. This was only possible because the captain had instinctively started the Auxiliary Power Unit (APU). Starting the APU was in the procedure, but in a part they never reached. Without the APU they would not have been able to lower the flaps, which would have led to a much higher landing speed, making it more difficult to land the aircraft in one piece. In this example, adhering strictly to the procedure could have made the situation much worse.

Example Two: Filtering Alerts All modern commercial aircraft are equipped with a number of systems that monitor the safety of the aircraft. Examples are the Traffic Collision Avoidance System (TCAS) to avoid surrounding traffic, over-speed and stall warnings, and the Enhanced Ground Proximity Warning System (EGPWS) to avoid terrain collisions, etc. There is no communication between these different systems. Each system monitors its own set of rules and produces warnings if necessary. Under most circumstances, conflicting warnings are rare. For the rare cases where multiple conflicting warnings do occur, a ‘filter’ is implemented to prioritize the warnings and only pass the most important ones. When, for example, a terrain warning and a traffic warning produce conflicting resolution advisories, the terrain warning will take priority because of the higher risk of a terrain collision. The individual systems have limited awareness and share no situational awareness. It is up to the pilot to keep track of the complete picture.

The problem with this kind of ‘filtering’ is that it relies on an up-front assessment of the risks involved and the context in which the warnings occur. Modern systems have come a long way and are performing well, but they do not support unanticipated situations. If pilots are aware of the rationale behind the warning systems, we expect them to be able to filter the information themselves in a way that is more appropriate for their situation. With the combined terrain and traffic warning example, the default strategy of giving terrain the higher priority will be suitable for most situations. But in a situation where the terrain alert is due to a high obstacle in the immediate surroundings, a pilot can choose to give priority to the traffic warning and descend towards the terrain if he has a good visual reference and can assess whether the aircraft is clear of all obstacles.

These two examples highlight some of the main challenges in commercial aviation. The human contribution is crucial in dealing with critical and unexpected events, but the complexity of modern aviation makes it difficult to be continuously aware of all relevant information required to handle any possible situation. In modern aviation, both aircraft and pilots are well equipped for normal operations, and procedures are in place to keep the aircraft’s operation within acceptable limits. But the complex interactions and uncertainties can sometimes—very occasionally—open up opportunities for intricate and complex unexpected events where, for example, procedures do not work anymore or the automation’s actions are not appropriate in the given context, which results in a need for cognitive

work. Supporting pilots during these unexpected events requires a fundamentally different approach when designing interfaces and automation.

In process control, another safety-critical socio-technical systems domain, Ecological Interface Design (EID) emerged as a framework to help in the design of support interfaces for unexpected events; it is described in Section 2.2. This thesis investigates the applicability of EID to support pilots in unforeseen situations. In order to make the kind of decisions highlighted above, pilots need to be aware of both the physical limitations and the procedural and regulatory requirements. Section 2.3 will describe how these two sets of constraints can be modeled. Finally, Section 2.4 will highlight some challenges in visualizing both sets of constraints.

2.2 Cognitive Systems Engineering

In the 1980's, when socio-technical systems became more and more complex, Cognitive Systems Engineering (CSE) emerged as an analysis and design framework that tried to improve the more traditional frameworks that mainly operated on the physical and physiological level. CSE is built on the notion of a cognitive system: an adaptive system that functions using knowledge about itself and the environment (Hollnagel & Woods, 1983). The main need for a CSE approach stems from the nature of the tasks performed by operators in Human-Machine Systems (HMSs). Work has shifted from being based on perceptual-motor skills toward cognitive tasks like problem solving and decision making.

In 1992, Vicente and Rasmussen presented EID as a new framework for designing interfaces for complex human-machine systems (Vicente & Rasmussen, 1992). EID attempts to extend the benefits of Direct Manipulation Interfaces to more complex work domains, especially focusing on the challenges posed by unanticipated events. The goal of EID is to support the *entire* range of cognitive activities operators will be faced with and to not contribute to the difficulty of the task.

Vicente defines the core of the interface design problem around two fundamental questions, as shown in Figure 2.2. The first question is how to describe the complexity of the work domain. This question is primarily related to the fundamental characteristics of that domain. The tool selected by Vicente to answer this question is the Abstraction Hierarchy (AH) of Rasmussen (1985). The second question is how to communicate the information in a useful and meaningful way to the operator. This question is primarily related to the characteristics of the operator. The Skills, Rules, Knowledge (SRK) framework was selected to describe the ways in which operators process information and to aid in the design of the interface (Rasmussen, 1983).

2.2.1 The Abstraction Hierarchy

Vicente and Rasmussen (1992) have shown that the AH is a useful framework to represent a work domain in a way that is relevant to interface design. The AH is a

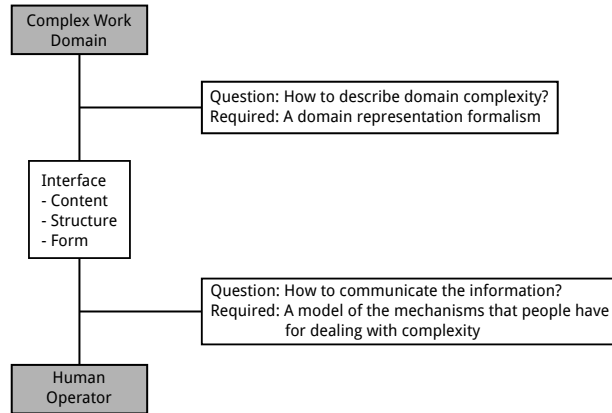


Figure 2.2: The structure of the interface design problem (Vicente & Rasmussen, 1992).

stratified hierarchy that models the same system under investigation at different levels of abstraction. The relation between the different abstraction levels are not arbitrary, they are linked through a *why-what-how* relationship. When focusing on a specific level, this level describes the *what*. The more abstract level above then indicates the *why* and the more concrete level below specifies *how*.

This holds for the levels in general, but can also be applied to individual functions in the AH. Each individual function at a specific level can be connected to one or more functions at the higher and lower levels. The functions at the higher level describe the *why* for this function, the functions at the lower level again describe the *how*. The relations between functions on different levels are usually referred to as *means-end* links. Specific functions on a lower level form the *means* to a specific function at the higher level, the *end*.

To clarify the *why-what-how* relationship, Figure 2.3 shows a partial abstraction hierarchy for a hypothetical scientific research program in human factors. It shows the part related to workload and performance measurement. The actual levels are not important for the example, it is chosen to show the hierarchical relations between adjacent levels. With this AH we can for example select the reaction time as the *what* under consideration. The reason *why* the reaction time shows up on this level is found in the level above, to measure performance. The lower abstraction level defines *how* the reaction time will be measured. Either with a stopwatch or with a computer. The lines between the functions represent the 'means-end' links. In this case the reaction time can be seen as a *means* to measure performance—the *end*—at a higher level of abstraction. At the same time, the reaction time can also be seen as an *end*. The computer and stopwatch are both

means to achieve this end.

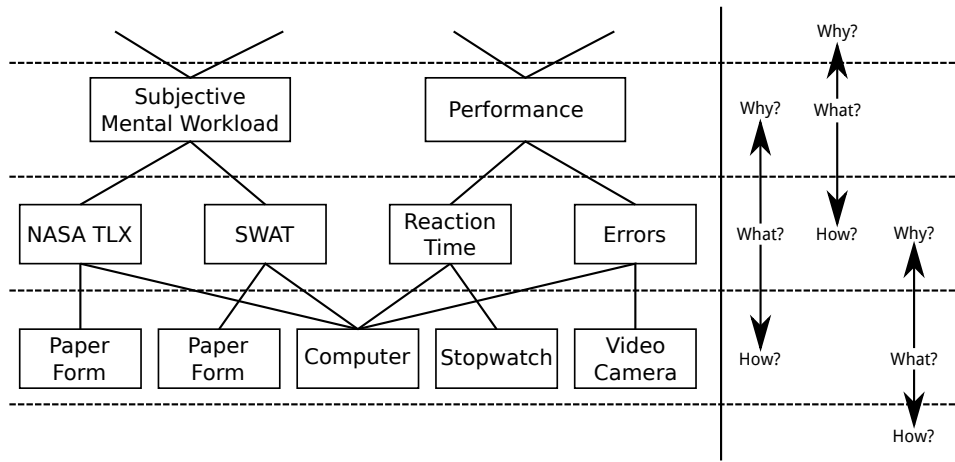


Figure 2.3: A part of a hypothetical abstraction hierarchy (Vicente, 1999) showing the why-what-how relationship and the means-ends links.

The number of levels and actual information captured by these levels varies and depends on the type of analysis being performed and the system under study. The form used in the remainder of this thesis is based on the levels used by Rasmussen (1985). Rasmussen uses five levels that have been identified from verbal protocols related to energy-conversion systems and digital computers. This representation maps well to modern aviation systems. Aircraft can be, for some aspects of their functioning, considered as energy conversion systems, and modern avionics controlling aircraft are mainly implemented as digital computers. The five levels are defined as follows (Rasmussen, 1985; Naikar & Sanderson, 1999):

Functional Purpose

This level of abstraction contains objects describing the high level purpose of the system, the purpose for which it was designed. The concepts at this level describe the highest level objectives of a work domain.

Abstract Function

At this level, the system is modeled in terms of the fundamental principles that are required to achieve the functional purpose described in the level above. The system can for instance be expressed in terms of mass and energy balances, and in general terms as fundamental physical processes.

Generalized Functions

The generalized function level describes the system's overall processes independent of the underlying physical implementation.

Physical Functions

This level describes the functionality afforded by physical devices in a work domain. It can also describe environmental conditions that have an impact on the system.

Physical Form

The lowest level of abstraction describes the physical appearance and physical location of all components of the system.

During problem solving and troubleshooting, operators will move up and down the hierarchy. The different levels of abstraction and the *means-ends* links between them can show how the lower level system components are related to the goals of the system, and vice versa. If a goal of the system is not met, the operator can trace a path to specific system components through the links to identify the cause. In the other direction, a component failure can also be traced back to the higher levels to discover the impact of the failure on the overall system goals.

An important advantage of the AH is that the higher levels provide a less detailed view of the system which makes systems look less complex. This is especially important in complex socio-technical systems with multiple interdependencies. Moving up a level removes complexity while a complete overview of the system is maintained.

The information captured in the AH will depend on the choice of system boundaries. These boundaries are somewhat arbitrary, however, and depend primarily on the scope and purpose of the analysis.

Looking back at Figure 2.2, the AH provides one of the two inputs for the design process. It tries to capture all information relevant to the problem that needs to be visualized. Any error, mistake, or omission in the AH will result in missing information on the resulting interface. Generally, multiple design iterations are needed to flesh out the design and capture as much information as possible.

2.2.2 The Skills, Rules, Knowledge Framework

EID proposes to use the SRK framework to model the mechanisms operators have for dealing with complexity. This framework describes three levels of cognitive control employed by operators. At the lowest level, there is Skill-Based Behavior (SBB), behavior that requires no or limited conscious control. At a slightly higher level of cognitive control, operators use Rule-Based Behavior (RBB). This is characterized by the use of rules and procedures to accomplish a goal in familiar circumstances. In novel and unexpected situations, operators revert to Knowledge-Based Behavior (KBB). At this level, there are no straightforward rules that can

be applied. Operators need to apply analytical problem solving skills and mental simulations to control a system.

These three levels can be divided in two categories: RBB and SBB relate to perception and actions, KBB relates to problem solving. The first category usually happens fast and can be done in parallel with little mental effort. Examples of this are controlling the pitch and roll attitude angles of an aircraft by looking out the window and selecting a flap setting at a certain speed by observing the cockpit instruments. In the ideal case, the majority of the tasks for the operators should be on these two perceptual levels. The second category, the analytical problem solving, happens slower and can only be done in a serial fashion. An example of this can be found in the United Airlines Flight 232 incident where pilots landed the aircraft by only using throttle inputs after a severe hydraulic failure (NTSB, 1989). This type of control is required for dealing with unanticipated events.

The way in which the levels above are presented should not give the impression that they are mutually exclusive. Any kind of task in a complex environment will rely on all three levels simultaneously. There is, however, a preference for using the lower levels of control—SBB and RBB—over KBB whenever possible. Operators will try to use techniques that they are familiar with as much as possible, even if these techniques might be inappropriate at times. This leads to three guiding principles in the design of ecological interfaces (Vicente & Rasmussen, 1992):

- The interface should directly support SBB. Time-space signals on the display should directly support an operator's subconscious control.
- A consistent one-to-one mapping between work domain constraints and cues or signs should be provided to support RBB.
- Use the abstraction hierarchy to analyze the work domain, and create a work domain representation to serve as an *externalized mental model* to support KBB.

2.3 Intentional Constraints

As described in Section 2.2.1, the AH is a fundamental part of the EID process. The process of obtaining the AH is not always easy. Vicente (1999) provides a systematic approach to Cognitive Work Analysis (CWA). The first step in this framework is Work Domain Analysis (WDA), a technique to obtain an AH for a specific system. The work domain is the system being controlled, independent of any particular workers, automation, event, task, or interface.

2.3.1 Work Domain Constraints

During the WDA, the functional structure of the work domain is analyzed. This functional structure is characterized by the constraints that limit the dynamic be-

havior of the system. The constraints determine the space of action possibilities that is available to perform work. This is represented in an abstract way in Figure 2.4. The concrete action space is a multi-dimensional space with dynamic and context-dependent boundaries in every dimension.




Figure 2.4: A constraint-based view on the work domain (Vicente, 1999). The causal constraints represent the physical limitations of the system.

Workers are able to choose any trajectory within the action space to satisfy their goals, they are all equally valid. No one trajectory is necessarily better than the other. However, all trajectories are constrained by the boundaries. How the trajectories interact with these boundaries depends on the type of constraint. Some constraints imply limits that can't be crossed. An example is the absolute ceiling of an aircraft, the maximum altitude at which an aircraft can sustain level flight. This type of constraint would *bend* the trajectory to become tangent to the boundary. The other type of constraint will result in a physical failure of the system. An example of this type of constraint is a mountain. A pilot can steer an aircraft into the mountain, which would cause physical damage to the aircraft. This is the equivalent of a trajectory crossing the constraint boundary.

All trajectories in Figure 2.4 are physically viable. Physical viability is a necessary, but not a sufficient prerequisite for successful operations. Not all trajectories are equally desirable. Some trajectories can be unsafe and will come too close to the boundary which can result in incidents or accidents when mistakes are made. Other trajectories can be inefficient from an economical perspective or even illegal. These issues can also be interpreted as new sets of constraints, but this time not originating in the laws of physics and the properties of the designed systems and its environment, but in the intentions of stakeholders in the work domain. In this thesis, these stakeholder intentions are represented by the rules, regulations, and procedures.

Rasmussen et al. (1994) characterized work domains based on the relative degree of causal and intentional constraint. Causal constraints constitute inevitable

boundaries on action originating in the laws of physics. They determine all possibilities for action afforded by the work domain. Intentional constraints, in turn, originate in actors' intentions and try to shape behavior in a work system by limiting the physical action affordance space to a reduced intentional space. Causal constraints dictate how one *can* act, while intentional constraints condition how one *should* or *would want to* act. Figure 2.5 shows an augmented version of Figure 2.4 with intentional constraints included. The complete space of action possibilities is still the same. Workers are still able to utilize the whole space, but the intentional constraints limit the amount of space they should use.




Figure 2.5: A constraint based view on the work domain including intentional constraints to direct workers behavior.

The physical implications of causal constraints can be split in two broad categories: inviolable and violable. An inviolable constraint represents a physical state that is impossible to attain and a path towards the constraint usually shows asymptotic behavior. An example of an inviolable constraint is the maximum velocity of a car. The exact value of this maximum speed might vary with wind speed, road gradient, etc., but under these specific conditions, the car will never be able to go faster than the theoretical maximum speed. Violable constraints are constraints on a state that can be attained or even crossed. This will inevitably lead to a breakdown of the system. As an example, the maximum structural load factor—the load at which a wing will start to deform or even break—of an aircraft is a violable constraint. In most aircraft, at high speed it is possible to apply enough control input to increase the load factor well above the maximum structural load factor, resulting in serious damage to the aircraft, usually ending in a crash.

The three constraint types are clarified by means of the traffic example introduced in Chapter 1 and, with the addition of the distinction between inviolable and violable constraints, again illustrated in Figure 2.6. This figure shows a single vehicle driving on a straight road, surrounded by water, with a road division, and

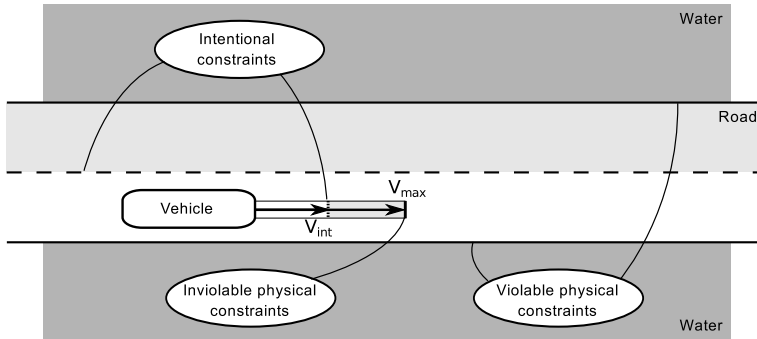


Figure 2.6: A traffic example illustrating physical and intentional constraints.

a speed restriction. The road boundaries form violable physical constraints. An accident occurs when the vehicle crosses these boundaries. The maximum speed of the vehicle constitutes an inviolable physical constraint. Other examples of inviolable physical constraints are the minimum break distance and turn limits of a vehicle. Notice how the physical constraints dictate how one can act. Finally, the road division and speed restriction form intentional constraints. Other examples of intentional limitations are traffic lights and priority rules. These intentional limitations guide how one should act, and may be respected or violated by cognitive agents.

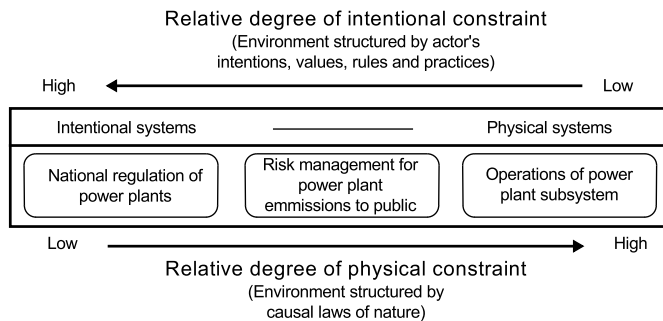


Figure 2.7: The intentional-physical continuum of a work environment (Hajdukiewicz et al., 1999).

The relative degree of causal constraints with respect to intentional constraints that needs to be taken into account in a WDA depends strongly on the purpose of the analysis. When looking at direct control of physical processes, it is sufficient

to take into account only causal constraints. On a slightly higher level, looking at the efficient use of a physical process in an economic setting would require additional inclusion of intentional constraints that would try to enforce the most efficient way of using the process. As discussed in Chapter 1, Figure 2.7 shows how different work domains related to power plants could be mapped onto the continuum depending on the type and scope of the analysis.

2.3.2 The Causal-Intentional Abstraction Hierarchy

This section will show two simple abstraction hierarchies for the demonstration system shown in Figure 2.6. The first one will be a purely causal abstraction hierarchy modeling only the physical aspects of the example. The second example will correspond to a work domain with a broader scope that includes intentional constraints.

The goal of our demo system, and most transportation systems in general, is to safely transport people and cargo from A to B. This constitutes the Functional Purpose of the system in Figure 2.8. The main way to achieve the transportation goal at the Abstract Function level is locomotion. The motion of the vehicles makes them effective as a means for transportation. Safety in this system is ensured by maintaining separation. The way in which the vehicle moves also contributes to the safety of the system, thus locomotion is also linked to safety.

The Generalized Function level describes *how* the vehicles motion is governed by the dynamics, and *how* the required separation depends on a combination of these dynamics and the obstruction posed by the road and traffic.

The more concrete description at the Physical Function level shows that in this system, the dynamics mainly impose a constraint on the velocity of the vehicle. The obstructions in the work domain of this system are the road boundaries and the other traffic.

The AH shown in Figure 2.8 shows a—simplified—physical view of a vehicle transportation system. The low level constraints—velocity, road boundaries, and traffic—will always need to be satisfied by the operator of the vehicle, but it is important to note that these are still abstract representations. The road boundaries will only depend on the position of the vehicle, but the influence of the surrounding traffic changes based on the trajectories chosen by the operators of the other vehicles. The constraint on velocity can be split in two different constraints. One is the maximum achievable velocity of the vehicle. This constraint can for example prevent an operator from over taking another vehicle with opposing traffic. The second constraint on velocity is derived from the separation requirement. Traffic and road boundaries consist of physical obstructions, but the velocity of the vehicle will ultimately determine whether the vehicle will stay clear of them or not. If a vehicle would be driving on the right-hand side of the road, turning right will not be possible anymore while turning to the left is still possible.

The system sketched above, although greatly simplified, is representative for simple low-volume traffic systems comparable to the early days of the car. With

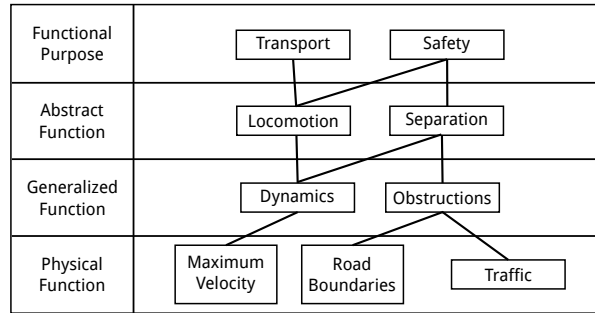


Figure 2.8: A simple causal Abstraction Hierarchy based on the system sketched in Figure 2.6.

low traffic density and low speeds, staying on the road and avoiding other traffic is rather straightforward. With increasing traffic densities and increasing velocities, maintaining a steady flow and avoiding collisions becomes more difficult. To solve this, rules, regulations, and procedures can be implemented to achieve a high traffic flow while maintaining a required level of safety. One of these rules could be to always drive on the right hand side of the road. In this way, bi-directional high-density traffic is possible with a low chance of collisions.

Adding rules, regulations, and procedures is equivalent to adding extra constraints to the work domain and these will limit the available action space as described above. The requirement to drive on the right hand side will be an intentional constraint, as described in Section 2.3.1.

Adding a set of intentional constraints to impose rules, regulations, and procedures broadens the scope of the work domain. This result in an extended AH shown in Figure 2.9. In this thesis, a strict separation is used between the causal domain and the intentional domain. This is not strictly necessary, both domains could be mixed, but since the objective of this thesis is to explicitly visualize the distinction between causal and intentional constraints, an explicit split in the AH is used to keep track of the domain each object belongs to. By doing this, the left hand side of Figure 2.9 is identical to Figure 2.8.

Adding rules, regulations, and procedures expands the Functional Purpose of the system. The ultimate goal of the rules, regulations, and procedures for the demo system is to achieve safety and efficiency through structuring of the transportation system. This structuring can be achieved by imposing safety criteria and flow management, which is found at the Abstract Function level. The safety criteria attempt to prevent incidents and accidents, or at least limit their impact. Flow management aims to achieve steady throughput.

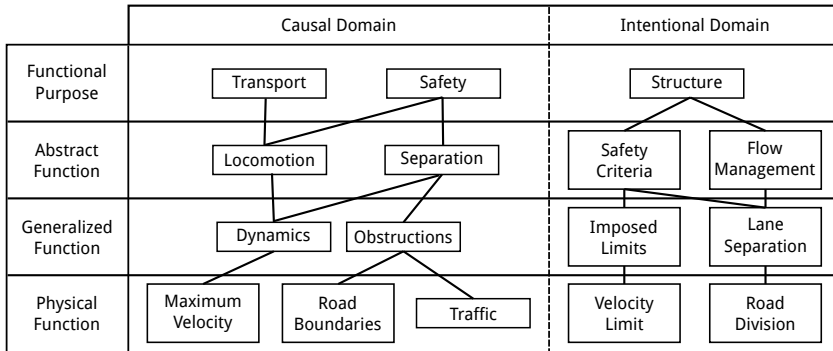


Figure 2.9: A simple causal and intentional Abstraction Hierarchy based on the system sketched in Figure 2.6.

At the Generalized Function level, the safety criteria are implemented as imposed limits on the vehicle capabilities and in terms of lane separation. Obviously, lane separation—keeping opposing streams on different sides of the road—is also a means to establish flow management. At the Physical Function level, the generalized functions result in velocity limits and in the division of the road in multiple lanes.

2.4 Implications for Design

There are numerous examples of constraint-based displays for the aviation domain, many of which were developed at the TU Delft (Van Dam, Mulder, & van Paassen, 2008; Klomp, van Paassen, & Mulder, 2011; Amelink, van Paassen, & Mulder, 2003; van Paassen, Gernaey, Veld, & Mulder, 2007). The majority of these examples either do not include intentional constraints or they lump the causal and intentional constraints together. Explicitly incorporating intentional constraints in the analysis has implications, however, for the design phase of EID.

As explained in Section 2.3.1, causal constraints form the hard limits of the system. Either because they can not be physically violated, or because they can not be violated without severe consequences. This property should be clearly reflected in the interface. Depending on the type of constraint, visualizing the difference between a violable and an inviolable causal constraint can also be beneficial. On the one hand, inviolable constraints are mainly informational artifacts, they can be used to see if a solution is feasible, and what the ultimate limit of the system

in this dimension is. Violable causal constraints, on the other hand, require active attention from the operator who must ensure to not get into a situation that results in the violation of this constraint.

Reiterating the traffic example, the maximum velocity can be used by the driver to estimate the absolute minimum time required to cover a specific distance. There is, however, no need to allocate any cognitive resources to avoid reaching this maximum velocity. A simple tick mark can be sufficient to indicate this type of constraint to the operator. This is different when a physical obstacle is present. The driver must either make sure he is not driving towards the obstacle, or if he is, needs to constantly monitor his velocity and remaining distance to decide when to initiate an avoidance maneuver. This requires a much more complex representation on the interface to show the relation between the car's dynamics and the obstacle.

Typically, intentional constraints are violable since they do not represent a physical aspect of the system. It is, however, possible that they do not satisfy the causal constraints. A speed restriction on a road can always be violated, as long as the vehicle's maximum speed is higher than the restriction. If it is not, the intentional constraint can not be violated.

One approach to include intentional constraints on an interface could be to treat them as if they were causal constraints. This will drive operators towards perfect procedural compliance, but would undermine one of the fundamental strengths of the ecological approach. With unexpected events, operators would be unaware of the actual physical constraints. It is possible to end up in a situation where violating intentional constraints is the *only* path out of a dangerous situation. When intentional constraints are treated as causal constraints there are two possibilities in these situations. Either the operator is unaware that there is still some control space left and will not find a way out of the situation, or the operator will be aware that the boundary is not causal, but will have no indication how much margin is left to the causal constraint. This makes it important to clearly visualize the difference between causal and intentional constraints.

By making the distinction between causal and intentional constraints apparent on the display, it is hypothesized that operators have the tools to reason about their work environment in a better way and as such can better assess risks. In normal operations, they will comply with the intentional constraints and look for trajectories in the available intentional control space. If they end up in a situation where there is no intentional control space left, they can evaluate which intentional constraint(s) can be violated.

In the mapping on the interface, based on this hypothesis, the violable nature of the intentional constraint should be clear. There should be a strong drive towards satisfying the intentional constraint, but the available margin with respect to the underlying causal constraint should be clear. In the traffic example, the speed restriction could be indicated to the driver through e.g., shading the part of the speedometer above the restricted speed in a specific color or with a certain pattern. This would allow the driver to immediately see the speed restriction, but

also be aware of the extra speed margin available when overtaking traffic.

In the next chapter, examples of two ecological aircraft interfaces will be described that have been developed in previous research. The chapter will use these two examples to explain, in more detail, how a mixture of causal and intentional constraints has affected the design of those interfaces. It will then show how creating an explicit distinction between causal and intentional constraints can further improve the design of these interfaces.

Case Studies

This chapter will present two case studies of existing ecological interfaces in aviation, and will consider them in the light of the theoretical foundations discussed in the previous chapter. After a short introduction, each section will discuss the purpose and goals of the interface relevant to this research. The current interface is presented and the relevant elements are discussed. Following this, the corresponding Abstraction Hierarchy (AH) is shown and analyzed with an explicit split between a causal and an intentional domain in mind. Based on this analysis, an augmented version of the interface will be presented, to show how a better distinction can be made between causal and intentional constraints.

3.1 Case One: Synthetic Vision Display

In the past decades, various Synthetic Vision Display (SVD) systems have been introduced in aviation. One of the major drivers for the research in SVDs was that there is still a significant amount—14 out of 65 fatal accidents between 2006 and 2014—of Controlled Flight Into Terrain (CFIT) accidents linked to a lack of situational awareness (Boeing Commercial Airplanes, 2016). The core of an SVD is a virtual three-dimensional representation of the surrounding terrain with a traditional Primary Flight Display (PFD) on top. This provides vital information in low visibility conditions, but even under good visibility important features of the surrounding terrain can be highlighted. Typical examples of this include coloring terrain that is higher than the aircraft's current altitude, indicating a predicted point of impact, and displaying terrain evasion maneuver commands.

Borst et al. (2010) noted that while an SVD brings a number of improvements, it still lacks a number of crucial features for successful and safe terrain avoidance. For example, the combination of the synthetic terrain and a flight path vector indicator will clearly show if the current trajectory will eventually result in a collision with the terrain or not. There is no direct way, however, to relate the aircraft per-

formance capabilities to the requirements imposed by the terrain. Pilots can see the terrain, but are not supported in their decision making whether the terrain can be passed, under current conditions, or not. They need to interpret the raw data and form a mental model of the situation themselves. In practice, aviation authorities still require that synthetic vision systems are backed up by last resort Terrain Awareness Warning Systems (TAWS). These systems can help to avoid a terrain collision, but do not necessarily give the pilot more insight in the situation.

To resolve these issues, Borst et al. (2010) applied the Ecological Interface Design methodology to augment an SVD with cues to better assist pilots in choosing strategies to avoid the terrain. Although experiments show that these ecological additions were indeed useful, the design did not make an explicit distinction between causal and intentional constraints. This section presents an analysis and proposes additions to better visualize this distinction.

3.1.1 Purpose and Goals

Borst et al. (2010) proposed to extend synthetic vision displays by making the terrain and aircraft performance constraints perceptually evident on the interface, to better relate the *internal* and *external* constraints on the aircraft motion. The SVD shows the external constraints, imposed by the terrain. The internal constraint, the aircraft performance, is not shown. It is hypothesized that it is beneficial to show both sets of constraints on the display. This was formalized in two main goals:

1. Pilots must be aware of what their aircraft allows them to do relative to the terrain;
2. Pilots must perceive terrain features relevant to the safety of their flight.

3.1.2 Current Ecological Interface

The SVD is a three-dimensional representation of the world mapped on a two dimensional PFD. Due to the nature of this mapping, pilots can not perceive absolute distances, heights, and locations of objects very well (Bolton & Bass, 2008). These features can only be perceived in a relative way through relative angles, occlusion and relative size of objects (Borst et al., 2010). In the case of aircraft displays, the angular features of the PFD are an advantage since a number of important aircraft performance characteristics can also be naturally expressed in an angular form.

Constraints

Terrain Angle Significant terrain peaks can be characterized by a combination of the distance to this peak, D_{peak} , and the relative height with respect to the own aircraft, ΔH_{peak} . As Figure 3.1 shows, this results in a corresponding terrain angle, γ_T , that can be expressed as :

$$\gamma_T = \arctan \left(\frac{\Delta H_{peak}}{D_{peak}} \right) \tag{3.1}$$

The above formula only holds for relatively sharp peaks. The terrain angle is defined as the tangent to the highest point, which means that with flatter peaks, the distance and height may need to be adjusted slightly; this is illustrated in Figure 3.1. In these cases, ΔH_{peak}^* and D_{peak}^* should be used in Equation 3.1.




Figure 3.1: The terrain angle.

As noted above, the disadvantage of common SVD displays is that accurate distance and height are difficult to estimate from the display. From Equation 3.1 it can be seen that when two peaks have the same ratio between distance and height, the terrain angle will be the same. When not applying special drawing techniques, like a grid overlay, to explicitly visualize distance and height, drawing a high mountain at a large distance may look exactly the same as a low mountain at a short distance, as shown in Figure 3.2.

The terrain angle can be interpreted as an angular constraint on the flight path of the aircraft. It indicates the minimum flight path angle required to climb over terrain peaks located in the current direction of travel. Interpreting the virtual terrain as an angular constraint requires no special additions or computations. By its nature, the outline of the terrain relative to the horizon is the direct representation of the angular constraint. Adding the traditional pitch ladder to the display will allow pilots to accurately quantify the angle. Adding a flight path vector indicator to the display will further allow them to directly relate the current flight path to the angular constraint imposed by the terrain.




Figure 3.2: Virtual Terrain Scaling.

Climb Performance The flight performance of an aircraft can be expressed in many different forms (Ruijgrok, 2009). In the context of terrain avoidance, the most important performance variable is the climb angle. More specifically, the optimal climb angle, the angle that yields the most altitude gained within a particular distance. In steady symmetric flight, the climb angle, γ , is expressed as:

$$\gamma = \arcsin \left(\frac{T}{W} - \frac{1}{(C_L/C_D)} \right) \quad (3.2)$$

In this equation, the weight, W , is quasi steady. It changes due to fuel burn, which means that during long flights the climb performance will improve, but in the short term the change is too small to be noticeable while controlling the aircraft. Both thrust, T , and the lift-to-drag ratio, C_L/C_D , can be directly controlled by the pilot. In aircraft with turbojet engines, thrust can be considered constant with airspeed and directly depends on the throttle setting. C_L/C_D can not be directly controlled, it depends on the angle of attack. During steady flight, the angle of attack is directly related to the aircraft velocity.

The relationship between aircraft velocity and C_L/C_D is shown in Figure 3.3. As can be seen, the way C_L/C_D changes depends on the aircraft velocity itself. In the lower speed region, C_L/C_D increases with increasing velocity. In the higher speed region, C_L/C_D decreases with velocity. Pilots typically control the aircraft velocity by changing the aircraft's pitch through the elevator. As an example, assume the pilot pitches up the aircraft without changing the throttle, and assume the aircraft is flying in the high speed region. As a result of the pitch up, the velocity will decrease. After the transients are gone, the aircraft will have stabilized at a new, lower, velocity. In response, C_L/C_D will have increased. Looking at Equation 3.2, the result will be an increase in flight path angle.

When pilots change the throttle setting, they can decide how the aircraft will respond by simultaneously controlling the pitch of the aircraft. If they maintain their current velocity, the flight path angle will change as a result. But when they control the pitch angle such that the flight path angle remains the same, the aircraft's velocity will change. Obviously, they can also opt for a combination of




Figure 3.3: Typical relation between aircraft velocity and C_L/C_D . The graph starts at stall speed and ends at the maximum operating speed.

changing the flight path and the velocity at the same time. It is important to note, however, that in the end, the flight path angle will always be correlated with the velocity according to Equation 3.2.

Equation 3.2 shows that the maximum climb angle occurs when C_L/C_D has its largest value, regardless of the magnitude of the thrust. Looking at Figure 3.3, the maximum C_L/C_D corresponds to a specific velocity. This velocity is referred to as the *best angle of climb velocity* and is denoted by V_x . Flying at V_x will result in the best possible climb angle for the current power setting.

Based on this, the maximum climb angle, γ_{max} , can be expressed as:

$$\gamma_{max} = \arcsin \left(\frac{T_{max}}{W} - \frac{1}{(C_L/C_D)_{max}} \right), \quad (3.3)$$

where T_{max} is the maximum available thrust, and $(C_L/C_D)_{max}$ is the C_L/C_D at V_x . This shows that maximum thrust in combination with V_x results in the maximum attainable climb angle. When presented on a synthetic vision display, the maximum climb angle can be directly related to the terrain. As described above, the outline of the terrain represents the minimum angle required to clear the terrain in that direction. When adding an indication of the maximum climb angle, the angular distance between the terrain and the maximum climb indication represents the margin that is available or required to clear the terrain. If the maximum climb angle is above the terrain, it will be possible to climb in a straight line. In the other case, a turn will need to be made to find a direction where the terrain peaks lie below the maximum climb indication.

Next to presenting the maximum climb angle, it can also be useful to show the maximum climb angle for the *current* power setting. With this information, pilots are able to perform climbs over terrain with reduced power while maintaining confidence that they will actually clear the terrain.

Distance to Maneuver Combining the climb power and the required terrain angle shows more than just a binary indication of whether the aircraft can climb over the terrain, or not. The situation in which the climb angle is larger than the terrain angle can be interpreted in two ways. Either by assuming that there is excess power to climb, meaning that at full power the aircraft will clear the terrain with an extra margin, or the pilot can use a reduced power setting to climb. The other way is to treat the margin as a horizontal distance that can still be traveled until the climb really needs to be initiated.

The former interpretation is directly supported by the angular representation, but the latter is more difficult to perceive from the display. It is straightforward to calculate, it is defined as the distance over which the aircraft is still able to climb over the terrain:

$$D_m = D_{peak} - \frac{\Delta H_{peak} - \Delta H}{\tan(\gamma_{max})} \quad (3.4)$$

This distance can be visualized to assist pilots in deciding when to start their terrain avoidance maneuver.

Turning Performance In the horizontal plane, the terrain can also limit the turn radius of the aircraft. R_{max} is the largest turn radius possible that does not lead to a terrain collision, as illustrated in Figure 3.4.

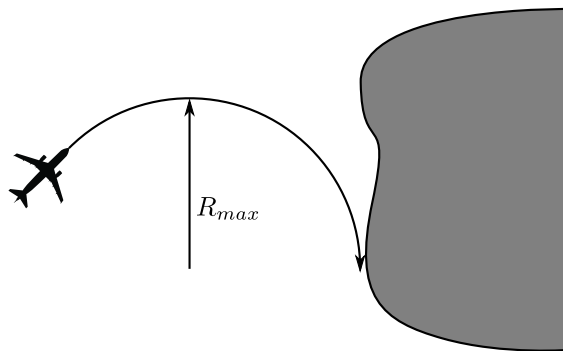


Figure 3.4: Maximum Turn Radius.

As before, the turn radius is a spatial constraint that is difficult to represent intuitively. Fortunately, the turn radius is directly related to the bank angle and the aircraft's velocity. When the maximum turn radius and aircraft velocity are known, the corresponding minimum bank angle can be calculated as follows:

$$\Phi_{min} = \arctan\left(\frac{V^2}{gR_{max}}\right) \tag{3.5}$$

Both R_{max} and Φ_{min} represent the same constraint, but expressing the constraint in the form of a bank angle is easier to present on a PFD and translates better to the control strategies employed when flying an aircraft.

Interface

The different constraints can be mapped to a typical PFD and the resulting ecological SVD is shown in Figure 3.5.

The virtual terrain can be drawn in various ways, with different kinds of texturing or shading indicating important properties of the terrain. In the context of this research, the focus lies on the outline of the terrain ①. It is this outline that indicates the required terrain angle to climb over the terrain. Next to immediately seeing the required climb angle in the direction of travel, a pilot can also immediately see if there are lower climb angles available by turning left or right.

The current flight path itself is indicated by the flight path vector ②. This indicator shows the actual direction of travel of the aircraft with respect to the outside world. The pilot has direct control over the flight path by using the flight controls. This will be his primary means to position the flight path vector in a way that will lead to the most efficient and safest solution.

The maximum climb angle bars ③ show the maximum climb angle that can be sustained when the engine is delivering maximum power. It is the main vertical limitation on the flight path vector. It is not impossible to point the flight path vector above these bars, but this will decrease the airspeed and may eventually result in a stall. Another useful feature of the maximum climb bars is that they can be directly related to the terrain outline. Since they show the maximum attainable flight path angle, their position with respect to the terrain outline immediately shows whether the aircraft is at all capable of clearing the terrain ahead. When the bars are above the terrain, as shown in the figure, the pilot can apply full power and clear the terrain.

Similar to the maximum climb performance bars, the current climb performance bars ④ indicate the maximum climb angle that can be sustained, but this time with the *current* power setting. The same relationships with the flight path vector and terrain outline hold as with the maximum climb angle. It also provides more information to the pilot. With the situation as indicated in the figure, a pilot can immediately see that with the current power setting it is impossible to clear the terrain. But since the bars move when the power setting changes, the pilot can




Figure 3.5: A schematic ecological SVD. With, ① the 3D terrain outline, ② the flight path vector, ③ the maximum climb angle bars, ④ the current climb angle bars, ⑤ the maneuver distance box, ⑥ the minimum maneuver distance bracket, and ⑦ the minimum required bank angle.

increase the throttles until the current climb bars are above the terrain, indicating that it is possible to clear the terrain at that power setting. In the example, this power setting will be close to the maximum power setting. In situations where the terrain is less high a pilot can use the current climb bars to climb over the terrain at reduced power settings which in turn can save the engine a bit.

As indicated before, the margin between the maximum climb angle and the terrain peak is directly proportional to the distance that can be traveled before the pilot has to respond. In the situation sketched in the figure, continuing on this flight path, the terrain will rise slowly, and at some point it will be higher than the maximum climb bars. This can be difficult to observe because of the display scaling. To improve the perception of this distance to maneuver, an alternative visualization is included. It consists of an expanding distance to maneuver box ⑤ and a minimum maneuver distance bracket ⑥. They are centered around the flight path vector. When the distance to maneuver drops below a certain threshold, the distance to maneuver box starts to expand proportionally to the remaining distance, calculated with Equation 3.4. When the box reaches the minimum maneuver distance brackets, the distance to maneuver has reduced to zero and the pilot will need to find another way out of the terrain conflict.

Finally, the turning performance constraints are shown on the bank indicator at the top of the display. As described above, the required turn radius can be expressed as a minimum required bank angle. This is indicated by a bank angle band ⑦ on the bank angle indicator. In this example, the pilot can immediately see that a right turn is only possible with a bank angle larger than approximately 25 degrees. With the bank angle in the gray band, the aircraft will turn into the mountain side eventually.

Experiment

Based on these additions, Borst et al. (2010) designed an ecological SVD and validated it in an experiment. The interface was compared with a baseline display that provided simple escape maneuver command cues. The experiment used a number of normal scenarios where pilots could fully focus on the terrain avoidance task and two unanticipated scenarios where pilots had to deal with unanticipated system failures while navigating their way through a mountainous area.

In the normal scenarios, the ecological display helped pilots understand the situation, at the cost of a higher response time and less safe maneuvers, i.e., the aircraft flew closer to the terrain. However, in the unanticipated scenarios, the ecological interface indicated a clear benefit over the command interface. The ecological cues immediately showed that aircraft performance was less than expected and directly triggered problem-solving activities. With the command based display, pilots did not always notice the lack of performance and were inclined to follow the commanded cues even though they were not appropriate, or even invalid.

The slower response time is not necessarily a problem. The ecological interface is not intended to be a last minute alerting display. On the contrary, pilots will be able to see a normal situation gradually evolve into a more difficult or even critical one. Operating close to the terrain on the other hand can indeed be less favorable behavior. Borst et al. (2010) noted that in this experiment pilots had a natural tendency to operate at the boundaries of system performance. Making the terrain and performance constraints visible allows pilots to do that. As discussed in the previous chapters, this is not a shortcoming in the Ecological Interface Design (EID) approach per se, but merely a result of a Work Domain Analysis (WDA) that did not explicitly take into account some of the safety margins that are usually applied through rules, regulations, and procedures.

3.1.3 Extending the Work Domain Analysis

Section 2.3.2 outlined the impact of adding intentional constraints on the Abstraction Hierarchy (AH). This section will start by presenting the AH used by Borst et al. (2010). Based on this AH, an extended AH will be discussed that can be used for a display with an explicit split between causal and intentional constraints.

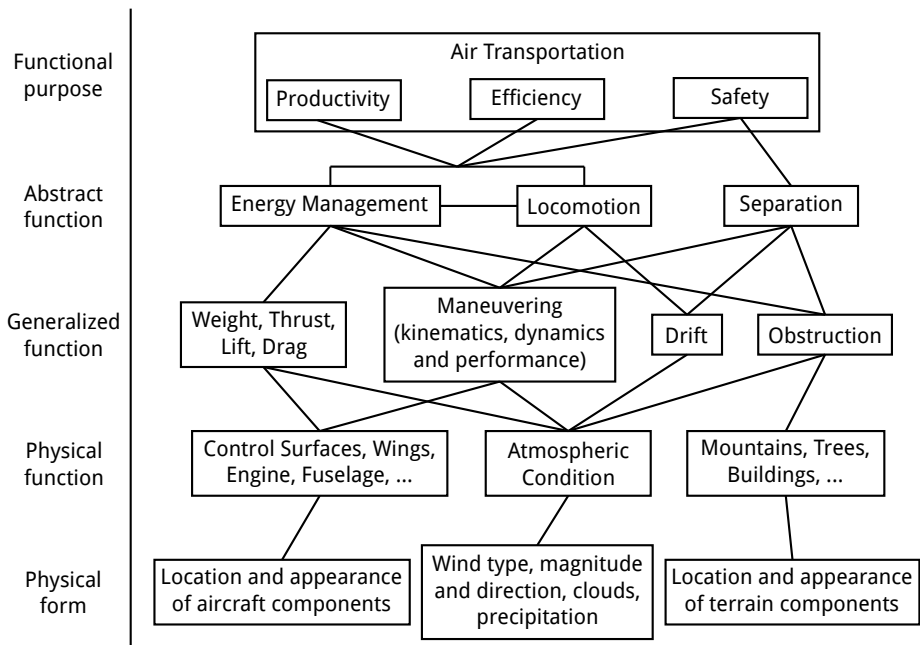


Figure 3.6: The basic abstraction hierarchy for the synthetic vision display (Borst 2010).

The original AH used by Borst et al. (2010) is shown in Figure 3.6. The main purpose of an aircraft and the crew in the context of the SVD is to provide air transportation in a safe, productive, and efficient way. This is represented at the Functional Purpose level. The underlying causal relationships—energy management, locomotion, and separation—are found at the Abstract Function level. The link between energy management and locomotion indicates that energy management can be seen as an abstract representation of speed- and altitude-based locomotion.

The Generalized Function level describes how the causal relationships of the Abstract Function level can be achieved, independent of the physical implementation of the system. On this level, weight, thrust, lift, drag, and thrust functions describe the abstract energy management function. The obstruction functions of the terrain and atmosphere relating to separation and energy management are also found on this level. Drift due to atmospheric conditions affects both locomotion and separation. Finally, the maneuvering function affords energy management, locomotion, and separation.

The functionality of the different system components are described on the Physical Function level. Control surfaces, wings, engines, a fuselage, etc. are means that afford maneuvering, and yield weight, thrust, lift, and drag forces. The atmospheric conditions have an influence on all functions at the generalized function level. They constrain the thrust, lift, and drag. They have an effect on maneuvering, they cause drift, and they form obstructions in certain cases. Finally, terrain features also define the obstruction function.

Finally, at the lowest level of abstraction, the Physical Form level, a detailed description of the physical appearance and location of all system components is found.

Borst et al. (2010) showed that a traditional SVD mainly maps the safety, separation, drift, obstruction, and terrain functions to the display. In this way, the display fails to communicate the *functional meaning* of the constraints. It shows the aircraft status together with the obstructions, but fails to show the relationship between the internal performance constraints of the aircraft and the constraints imposed by the terrain. In their ecological approach, they added these relationships in a way that is directly perceived by the pilots.

The ecological SVD is able to properly support the productivity, efficiency, and safety goal of the air transportation system from a causal perspective. In practical operations, however, pilots also have to take rules, regulations, and procedures into account. In the context of this thesis, the focus will be on the regulation concerning the Minimum Safe Altitude (MSA). The goal of these regulations is to improve safety by establish a safety margin above the terrain. Typically, above congested areas the MSA is 1,000 *ft*, and above populated areas the MSA is 500 *ft* above the highest object within a radius of 2,000 *ft* (ICAO, 2016).

The MSA is an intentional constraint. It does not have a direct relationship to the performance or constraints of the aircraft, but it is a blanket rule that ap-

plies to any aircraft. For this reason, in the extended AH shown in Figure 3.7, a new function is defined at the functional purpose level, rather than linking a new *intentional* abstract function to the existing safety function. This emphasizes the difference between both functional purposes. Safety in this case reflects the *pure physical* safety of the system, it results in constraints that *can never* or *should never* be violated. The regulations also introduce a layer of safety, but as long as they are in line with the causal constraints. When doing so, a pilot is technically breaking the law, but in specific situations, this could be the only option.

At the Abstract Function level the regulations give rise to safety criteria, which can be seen as priority measures needed to meet the system purpose. At the Generalized Function level, for this system, the safety criteria result in a buffer function. A generalized way of specifying separation. The more concrete Physical Function for the buffer is a horizontal and vertical separation function. The concrete implementation of the separation can be achieved in different ways. In this example, a virtual terrain is used and shows up at the Physical Form level. The implementation of this virtual terrain will be discussed in the next section.

3.1.4 Augmented Interface

Based on the augmented WDA the display can be improved. This section revisits the affected constraints and shows how the display changes when incorporating an explicit split between causal and intentional constraints.

Constraints

Terrain Angle The causal representation of the terrain remains the same as in the original display. The intentional constraint, the safety buffer around the terrain commonly referred to as the MSA, needs to be added. Conceptually, adding a buffer around the terrain can be seen as calculating a new virtual terrain based on the required safety margins. Figure 3.8 shows how such a virtual terrain can be created.

From a purely causal terrain as in Figure 3.8 (a), an additional virtual terrain representing the intentional MSA separation constraint can be created in two ways. A first option would be to only apply a vertical separation requirement. In this case, the intentional terrain is obtained by translating every point of the causal terrain up with the specified separation amount, as in Figure 3.8 (b). The second option is to use both a vertical and horizontal separation distance, in this case, each point is shifted away from the causal terrain by an amount determined by the separation constraints as shown in Figure 3.8 (c).

The effect of adding an intentional virtual terrain can be seen in Figure 3.9. Looking back at Figure 3.1, not much has changed. An intentional layer has shown up, and an additional set of describing parameters has shown up to describe the features of the intentional constraint. The most important change is

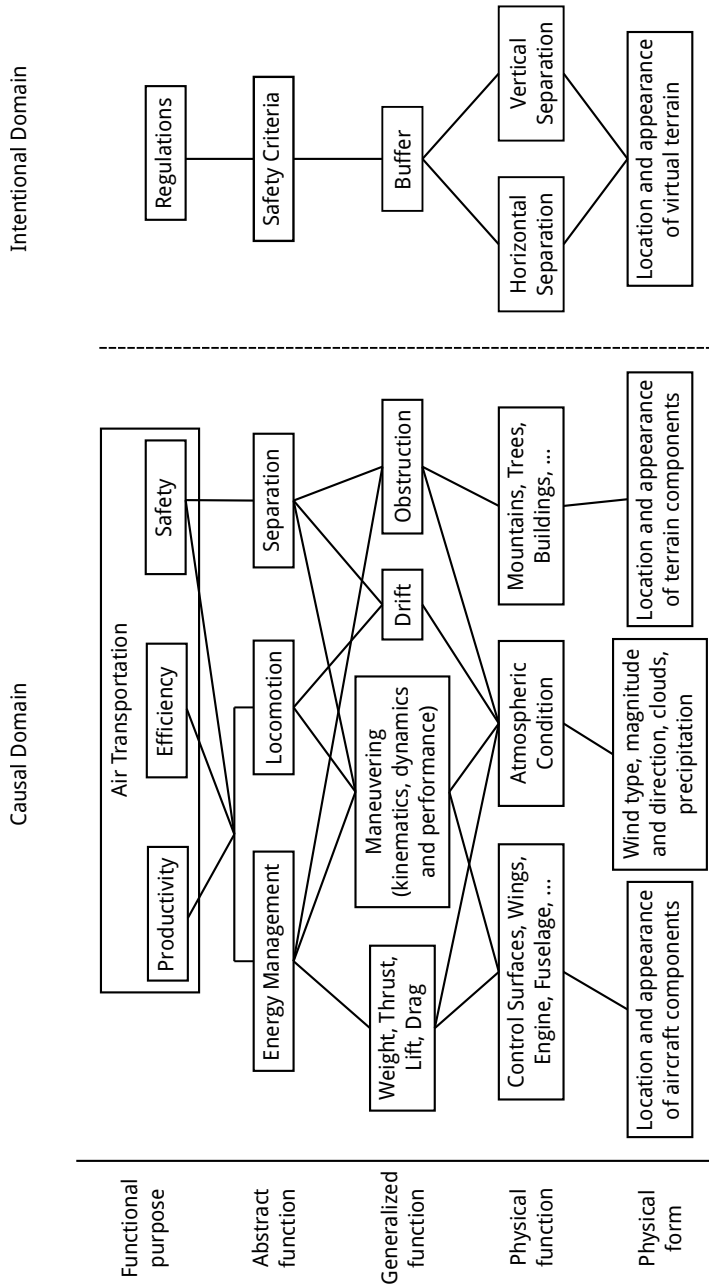


Figure 3.7: The intentional abstraction hierarchy for the synthetic vision display.




Figure 3.8: The virtual terrain, with (a) a purely causal representation, (b) adding a pure vertical separation, and (c) combining a vertical and a horizontal separation.

that now there are two terrain angles. One corresponding to the causal terrain, and one corresponding to the intentional terrain.

Distance to Maneuver In the beginning of this section it was shown how the remaining distance to maneuver can be calculated and what the relevance is to the terrain avoidance task. The same observations as before are valid. But since there are two terrain angles when incorporating intentional constraints, as described in the previous paragraph, there will also be two different values for the remaining distance to maneuver. One corresponding to the causal terrain angle, the other corresponding to the intentional terrain angle. Since the intentional terrain angle will always be larger than the causal terrain angle, the remaining maneuver distance for the causal terrain will always be larger than the distance for the intentional terrain. In other words, pilots will always run out of intentional maneuver distance before running out of causal maneuvering distance. When violating the intentional constraint, pilots move into the safety buffer.

Turning Performance Adding an intentional layer over the terrain also affects the constraints on the turning performance of the aircraft. The intentional layer will put an additional constraint on the maximum turn radius and thus the minimum required bank angle, as shown in Figure 3.9. Since the intentional terrain will always be closer to the aircraft, the turn radius required to avoid the intentional constraint will always be smaller than the radius required to stay clear of the causal terrain constraint.

As before, the turn radius constraint can be converted into a constraint on the bank angle. Once again this results in a causal bank angle constraint and an intentional bank angle constraint. The required bank angle for the intentional constraint will always be larger than the causal bank angle since the associated turn radius will always be smaller.




Figure 3.9: The causal and intentional terrain angles.

Interface

The adjustments to the constraints described above do not change the fundamental structure of the display. Figure 3.11 shows an example of what the explicit split between causal and intentional constraints could look like. Comparing with the previous example in Figure 3.5, the same features are present, but in a slightly refined form.

The causal terrain ① is still rendered in the same way as in the original display. The intentional separation constraint is drawn as a layer on top of the causal terrain ②. The same relationship with the flight path vector ③, current climb bars ⑤, and maximum climb bars ⑥ still holds. The intentional layer can be treated as if it were actual terrain. In situations as sketched in Figure 3.11 where climb power is not sufficient to clear the intentional layer, the pilot can see that he can violate the intentional constraint while still satisfying the causal constraint.

When the added intentional layer has a constant thickness, there are some additional benefits. Even though there is no proportional relationship between the eventual clearance and the position of the flight path vector in the intentional terrain layer, a pilot can make a rough estimate of the consequences of putting the flight path vector in a specific position. Keeping it close to the top of the layer




Figure 3.10: The causal and intentional turn radius.

will make sure that the eventual clearance is always above half the intentional separation. Furthermore, with a constant intentional separation, the thickness of the layer is related to the distance to the terrain. Adding the intentional layer can also assist in recovering some of the depth perception that is lost in synthetic vision systems.

In a similar fashion, the distance to maneuver indication can be augmented to incorporate the intentional terrain separation. The causal distance to maneuver box (6) and minimum distance to maneuver bracket (8) are the same as before, but there is an additional intentional distance to maneuver box (7) which is always going to be larger than its causal counterpart. Initially, both boxes will be much smaller than the brackets and will expand when the aircraft gets closer to the terrain. Once the intentional maneuver box touches the brackets, the pilot needs to start a maximum angle climb immediately to avoid violating the intentional terrain separation constraint. If this does not happen, the intentional distance to maneuver box will expand beyond the brackets and the causal distance to maneuver box will eventually reach the brackets. At this point, climbing is not an option anymore and a different solution needs to be found altogether.

Experiment

Based on the concepts sketched in the display above an experiment was set up to evaluate how pilots deal with the explicit visualization of an intentional terrain constraint. Since the experiment was meant as an initial investigation into the explicit use of intentional constraints, only a virtual terrain based on a vertical separation distance in combination with the climb performance indications were used. The distance to maneuver box and required bank angle indications were left out to prevent subjects from using different sources of information when flying




Figure 3.11: A schematic ecological SVD. With, ① the causal 3D terrain outline, ② the intentional terrain outline, ③ the flight path vector, ④ the maximum climb angle bars, ⑤ the current climb angle bars, ⑥ the causal maneuver distance box, ⑦ the intentional maneuver distance box, ⑧ the minimum maneuver distance bracket, ⑨ the minimum required causal bank angle, and ⑩ the minimum required intentional bank angle.

the scenarios. In other words, only the aircraft vertical/longitudinal motion will be investigated, turns and heading changes are not included. This experiment will be described in full detail in Chapter 4.

3.2 Case Two: Vertical Situation Display

In future aviation systems, there is a trend towards more responsibility for flight crews in terms of preferred trajectories and self separation. Although numerous separations aides exist at this time, they will probably not be able to provide the support necessary to cope with the new situations. Current alerting systems are mainly short-term conflict avoidance system. The Traffic Collision Avoidance System (TCAS) and Enhanced Ground Proximity Warning System (EGPWS) are systems that provide last minute advisories and resolutions. Pilots need to respond immediately, following the directions of the system. While these systems have been proven successful in preventing accidents, they are not designed with the purpose of *preventing* pilots getting into potentially dangerous situations. Furthermore, the rules used by these systems are rather opaque and not based on quantities that are easily perceivable by pilots.

This has two important implications. First, alarms can come as a surprise, raising stress levels in already dangerous situations (Burian, Barshi, & Dismukes, 2005). Second, for pilots it is sometimes difficult to distinguish between false alarms and proper alarms (Bliss, 2003). To address these issues, Rijnveld et al. (2010) proposed an integrated traffic and terrain awareness display with a large time horizon. The constraint-based display, derived from EID principles, directly show the relations between aircraft performance and constraints from traffic and terrain. With this system, pilots are able to see dangerous situations develop and are constantly aware of the control space that is available to avoid or get out of a dangerous situation.

As before with the SVD, this approach proved successful in enhancing situational awareness, but there is still room for improvement by implementing a better distinction between causal and intentional constraints.

3.2.1 Purpose and Goals

Rijnveld et al. (2010) showed how three different sets of constraints—aircraft performance, traffic, and terrain—can be combined in one display. The main challenge was to integrate these constraints in both the velocity and position domain. The main goals for this display can be formalized in four main goals:

1. Pilots must be aware of the performance constraints of their aircraft;
2. Pilots must be able to perceive terrain properties relevant to the safety of the flight;

3. Pilots must be able to perceive traffic properties relevant to the safety of the flight; and
4. The relationships between all sets of constraints should be perceived.

3.2.2 Current Ecological Interface

The Vertical Situation Display (VSD) shows a two-dimensional side view of the terrain and airspace directly in front of the ownship. The display maps both the position and velocity domain onto each other. Current-day VSDs present external constraints like traffic and terrain in the position domain, and present the current performance as a velocity vector in the velocity domain. Rijnveld et al. (2010) has shown that the representation can be adjusted to make the relationships between the different domains more meaningful.

Constraints

Terrain Angle The terrain constraints have been described in Section 3.1.2. In contrast with the discussion about a Synthetic Vision Display, however, the terrain angle is not directly perceivable in an VSD. The terrain angle will need to be visualized in the display. The most straightforward way is to draw a line segment that is tangent to the terrain which has the same origin as the ownship flight path vector. In this way, the flight path angle indicated by the angle of the flight path vector can be directly related to the terrain angle.

Traffic Conflict Zone Traffic in the neighborhood imposes a constraint on the maneuvering space of the ownship. This constraint can be formalized by defining the concept of the Protected Zone (PZ). The PZ is an area centered around an airplane that should not be used by other airplanes. Figure 3.12 shows the usual implementation of the PZ in two dimensions. The resulting shape of the PZ is a disc with radius R_{PZ} and height $2h_{PZ}$. Typical values for R_{PZ} are between 3 NM and 5 NM, $2h_{PZ}$ is usually 1,000 ft (ICAO, 2016). Hence, the shape of a typical PZ resembles that of a flat, coin-like shape also referred to as ‘puck’ or ‘hockey puck’. It is important to note that when one aircraft enters the PZ of another aircraft, the other aircraft is also inside the PZ of the first aircraft.

Two aircraft are defined to be *in conflict* with each other when they are on trajectories that *eventually* lead to a protected zone violation (Ellerbroek, Visser, Van Dam, Mulder, & van Paassen, 2011). Considering an ownship and one intruder, Figure 3.13 provides a schematic representation of a conflict in the vertical plane. The ownship is moving with velocity \mathbf{V}_{own} , the intruder moves in the opposite direction \mathbf{V}_{int} . With these two velocities, it is easy to calculate the relative velocity of the ownship with respect to the intruder:

$$\mathbf{V}_{rel} = \mathbf{V}_{own} - \mathbf{V}_{int} \tag{3.6}$$




Figure 3.12: Definition of the horizontal (left) and vertical (right) protected zones in terms of protected zone radius [R_{PZ}] and protected zone height [h_{PZ}].

The graphic representation of this operation is shown in the top part of Figure 3.13. If the direction of this relative velocity crosses the PZ centered around the intruder, the ownship will eventually enter the PZ, and both aircraft are defined to be in conflict. The two gray lines tangent to the PZ show the bounds on the relative velocity that lead to a conflict. In this example, the relative velocity points into these bounds indicating that with the current velocities, the two aircraft are indeed in conflict (van Paassen, 1999; Van Dam et al., 2008).




Figure 3.13: Definition of the vertical conflict zone.

Using the *relative* velocity to examine traffic situations is a straightforward way to determine if two aircraft are in conflict, but it has two drawbacks. First of all,

it would add extra information to a display that would be difficult to interpret. It is not always trivial to reason about the relationship between ownship velocity and the relative velocity. Pilots can only change their ownship velocity and may have difficulties in predicting what effects this change may have on the relative velocity. A second drawback is that in the case of multiple intruders, there is no directly visible correlation between the relative velocities with respect to *all* intruders. Moving one relative velocity vector out of the area resulting in a conflict could potentially move the other relative velocity into its respective conflict area.

The bottom part of Figure 3.13 shows a simple approach to remedy both drawbacks. Instead of calculating the relative velocity, it is possible to shift the triangle created by the tangents of the PZ with an amount equal to V_{int} . Instead of subtracting V_{int} from the ownship velocity, we *add* it to the points making up the triangle. The shifted triangle is called the Conflict Zone (CZ).

In this representation, the origin of the relative velocity is in the tip of the CZ and the tip of the relative velocity vector coincides with the tip of the ownship velocity vector. In this way, it is possible to directly relate the absolute ownship velocity to the CZ. If the tip of the velocity vector is in the CZ, both aircraft are in conflict. When it is outside the CZ there is no conflict. Since the pilot directly controls the ownship velocity vector, it becomes almost straightforward to avoid or resolve conflicts.

When multiple intruders are present, multiple CZs exist. One for each intruder, each with the origin corresponding to their respective V_{int} . Regions where CZs overlap are velocities that lead to conflicts with multiple intruders at once, albeit not necessarily at the exact same time. In this way, avoiding a conflict equals finding a position of the velocity vector outside all conflict zones.

Performance Envelope In a VSD, the performance constraints discussed in Section 3.1.2 still apply. The side view employed in a VSD allows for a more complete overview of aircraft performance than the representation used in an SVD. The velocity vector on a VSD shows both the magnitude and the flight path angle. This makes the VSD suitable to represent the full range of constraints on the velocity and the climb angle.

The relationship between velocity and flight path angle was already shown in Equation 3.2 and Figure 3.3. This relationship can be used to calculate the constraints on the velocity vector. The most obvious limits are the minimum and maximum velocity. They are represented by the stall speed, V_s , the speed at which flow separation occurs and the aircraft loses lift. And by the maximum operating limit speed, V_{MO} , defined by the manufacturer to avoid exceeding structural limits.

A pilot can choose to fly any velocity between V_s and V_{MO} , but the steady state velocity is always constrained by Equation 3.2. For a given throttle setting the equation can provide the steady state flight path angle for the full range of velocities between V_s and V_{MO} . When using T_{max} for all velocities, it is possible

3.2. CASE TWO: VERTICAL SITUATION DISPLAY

to create a curve representing the maximum climb angle for each velocity. In the same fashion, using $T = 0$ will give a curve representing the glide performance when all engines fail.

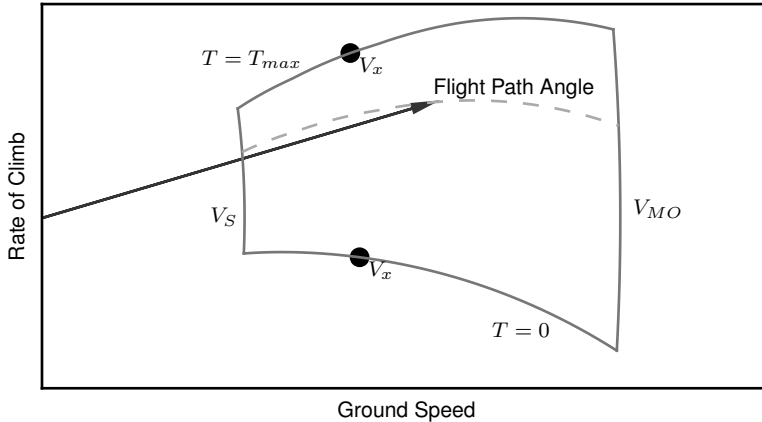


Figure 3.14: The performance envelope and the velocity vector using a typical aspect ratio of 2:1. The dotted curve indicates the climb angles at 70% power.

Combining V_s , V_{MO} , and the maximum climb and glide performance curves results in the performance envelope shown in Figure 3.14. The top and bottom curve represent the maximum climb and glide performance, respectively. Both ends of the curves are connected by semi circles representing V_s and V_{MO} . Together, these four curves represent the full achievable steady state performance.

It must be noted that each point inside the envelope is the result of a specific combination of throttle and elevator. As an example, Figure 3.14 shows the flight path angles when applying approximately 70% power. At this power setting, any velocity selected by the pilot will result in the flight path angle on the dotted line corresponding to that velocity. If the pilot would deliberately keep the flight path angle above the dotted line, the velocity will keep on decreasing until falling below V_s .

As a final observation, the maximum attainable climb angle and the best glide angle, shown in Equation 3.3, correspond to the tangent point of the (extended) velocity vector and the best angle of climb and glide curve. Both points can be indicated on the curves, and will correspond to V_x , see Figure 3.14.

Interface

A schematic overview of a typical ecological VSD is shown in Figure 3.15. The current state of the ownship is shown by the velocity vector ①. It shows the aircraft velocity in the vertical plane. The length of the arrow represents the True Airspeed (TAS) and its angle represents the geometric climb angle γ . The horizontal and vertical component of the flight path angle represent the ground speed and the vertical speed, respectively. For this type of display, the focus lies on the relationship between the geometry of the constraints and the position of the flight path vector. For this reason, the actual quantities of the variables indicated by the flight path vector are not directly perceptible, but they can be added in any number of forms to the display, if necessary.




Figure 3.15: A schematic ecological VSD.

The performance limits of the ownship are indicated with the performance envelope ②, already illustrated in Figure 3.14.

The resulting envelope provides an indication of the control space available to the pilot. Temporary excursions of the velocity vector above the maximum climb angle and below the glide angle are possible, but in steady flight, the velocity vector will always be inside the envelope. As such, the envelope represents the ownship performance constraints.

The virtual terrain ③ shows the vertical profile of the terrain relative to the ownship in the direction of travel of the aircraft. Especially when terrain features are further away, it is difficult to properly relate the flight path angle indicated by the flight path vector to the required terrain angle needed to clear the obstacle. To resolve this, a terrain angle line ④ is drawn inside the performance envelope. Keeping the flight path vector above the terrain angle line will result in a trajectory over the terrain.

3.2. CASE TWO: VERTICAL SITUATION DISPLAY

Like the terrain, traffic is also represented at its location relative to the ownship (5). Its speed and flight path angle are indicated by an arrow (6). The protected zone around the traffic is indicated by a box around the aircraft symbol (7). As explained in the previous section, there is no direct correlation between the traffic location and the ownship velocity vector. To provide this information, a conflict zone (8) corresponding to the traffic is drawn inside the performance envelope. Although temporary excursions above the maximum climb angle and below the glide angle are possible, the conflict zone is drawn between the stall speed arc and the minimum operating speed arc.

Experiment

Rijneveld et al. (2010) used an ecological VSD in an experiment to study the effects of combining terrain and traffic constraints in a single display. The experiment compared a baseline display, only showing the terrain outline and the traffic location, with an ecological version of the display described in the previous subsection.

They reported better situational awareness for pilots using the ecological version of the display. Pilots were more aware of the capabilities of the aircraft when avoiding traffic and terrain, but a significant effect on performance was not measured. They also noticed an increase in terrain intrusions with the ecological VSD. During the experiment, a 1,000 *ft* safety layer was superimposed on the terrain. With the baseline display, no subject entered this zone, but with the ecological VSD a number of subjects flew closer to the terrain, resulting in safety zone intrusions. There was no in depth analysis of the cause of this difference.

3.2.3 Work Domain Analysis

Similar to the analysis presented in Section 3.1.3, this section will start by discussing the original AH used by Rijneveld et al. (2010) which was discussed in Rijneveld (2010). Based on this AH an extended AH will be presented that uses an explicit split between causal and intentional constraints.

Figure 3.16 shows the original AH on which the work by Rijneveld et al. (2010) was based. Comparing it with the AH of the VSD in Figure 3.6 shows many of similarities. This should not be surprising, since they both describe collision avoidance work domains.

At the Functional Purpose level the same three functions—Productivity, Efficiency, and Safety—are found, with the exception that Productivity is called Performance in this AH. On the Abstract Function level, Separation is called Collision Avoidance and Force Balance is added. Locomotion and Energy Management are also repeated on this level. The means-ends links between both levels are also similar between the VSD and SVD.

Again, the Generalized Function level is similar to the SVD. Only the Drift function was left out in this case since the effect of wind was not included in the

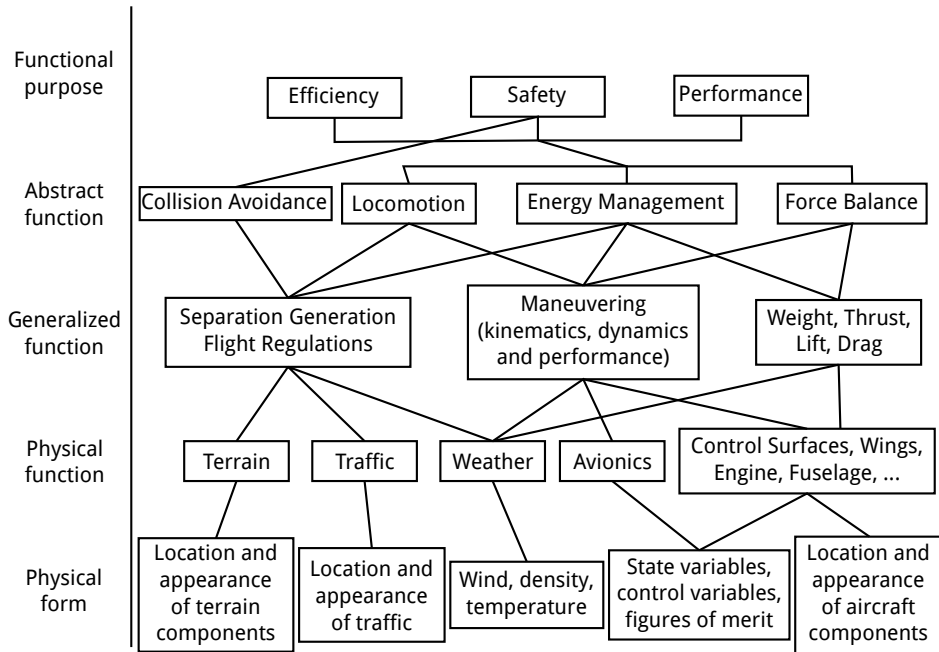


Figure 3.16: The basic abstraction hierarchy for the vertical situation display.

work of Rijnveld et al. (2010) and the Obstruction function is called Separation Generation and Flight Regulations. The means-ends links are also equal to the SVD ones, with the sole exception of the omission of the link between Collision Avoidance and Maneuvering by Rijnveld et al. (2010).

The Physical Function level is extended slightly. Control surfaces, wings, etc. are repeated. Atmospheric Condition and Mountains, Trees, Buildings, etc. are grouped into Weather and Terrain respectively. Since the VSD aims to provide information on terrain *and* traffic, Traffic is added to this level and is linked to the Separation Generation level. Finally, Rijnveld et al. (2010) decided to also add Avionics to this level.

The Physical Form will not be discussed in detail, but it provides the concrete physical representation of the functions on the Physical Function level.

In Section 3.1 it was shown how the original AH for the SVD was based on a purely causal WDA. This does not hold for the AH of the VSD. Although not specifically mentioned by Rijnveld et al. (2010), the Separation Generation and Flight Regulations function on the Generalized Function mixes a causal and an intentional constraint. Since the reason for the combination of separation and

3.2. CASE TWO: VERTICAL SITUATION DISPLAY

regulations is not discussed, it is not possible to know the exact reasoning behind this choice.

Looking back at the history of ecological terrain and traffic displays can give an insight, however. The early incarnations of ecological traffic displays have always used a protected zone around traffic to assure separation (Van Dam et al., 2008). The constraints used have always been intentional in nature. The terrain related displays on the other hand, have mainly focused on the physical constraints imposed by terrain (Borst, Suijkerbuijk, Mulder, & van Paassen, 2006). Bringing both constraints together also resulted in combining them on the Generalized Function level.

The significance of this mixed constrained was not lost in the display designed by Rijnveld et al. (2010). Instead of showing a purely causal representation of the terrain, a safety margin was added, not unlike the MSA discussed before. The traffic constraint, however, was not treated the same way. It only represented the intentional part of the constraint.

Based on the above, a revised and extended AH can be created. The first step is to transform the original AH shown in Figure 3.16 into a purely causal AH. This is done by removing the Flight Regulations at the Generalized Function level. The left side of Figure 3.17 shows the purely causal AH, the bold box shows where the Flight Regulations were removed.

The intentional additions for the VSD are very similar to the additions in Section 3.1.3 for the SVD. The same philosophy applies here, the causal part of the AH reflects the physical baseline that always applies. The intentional part reflects the rules, regulations, and procedures that should be adhered to but that allow some leeway under extraordinary circumstances.

This results in the AH shown on the right of Figure 3.17. The regulations become an additional Functional Purpose of the system. They result in Safety Criteria, that in this system are represented by Buffer functions. The Physical Function resulting from the Buffer function is an additional Terrain Separation and an additional Traffic Separation. At the Physical Form level these can be represented by a virtual terrain reflecting the MSA and the protected zone reflecting the separation minimums.

3.2.4 Augmented Interface

Based on the augmented WDA the existing ecological display can be improved. This section revisits the affected constraints and shows how the display changes when incorporating the explicit split between causal and intentional constraints.

Constraints

Terrain Angle The changes to the terrain angle are exactly the same as those discussed in Section 3.1.4. The same additions discussed and illustrated in Figure 3.8

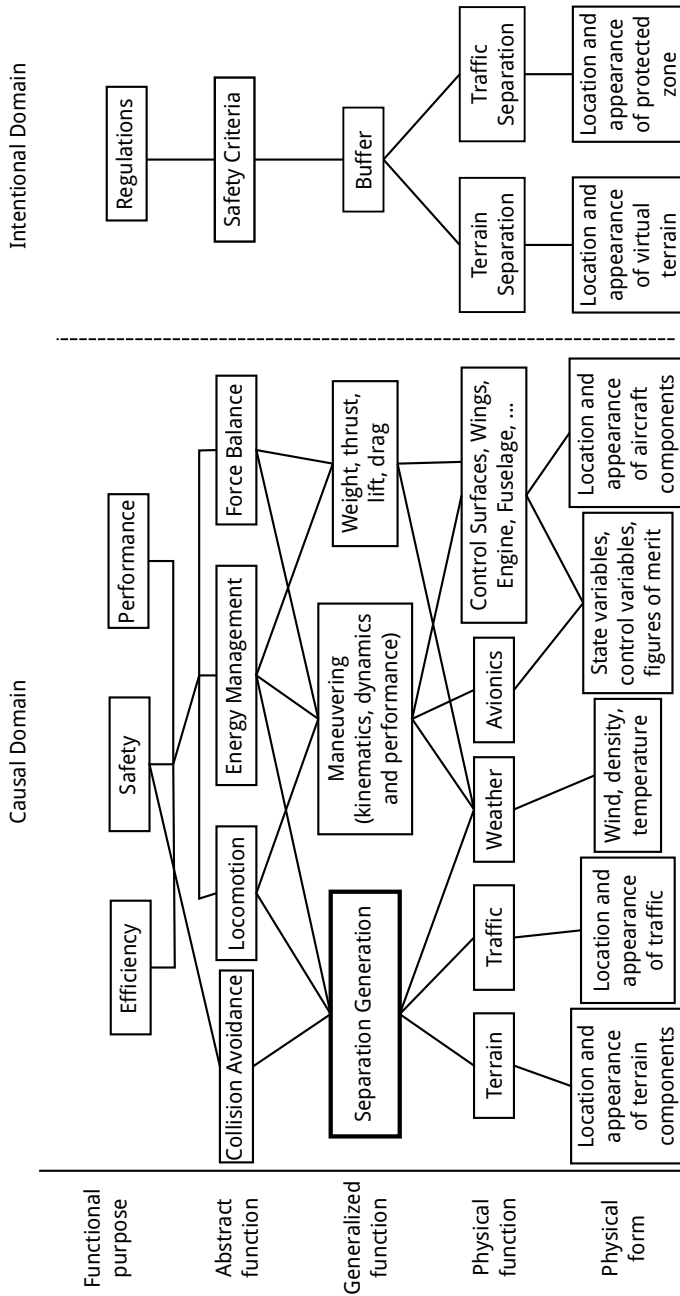


Figure 3.17: The intentional abstraction hierarchy for the vertical situation display

and Figure 3.9 apply to the augmented VSD as well.

Traffic Conflict Zone In Figure 3.12 the size of the aircraft is greatly exaggerated compared with the protected zone. In reality, an aircraft occupies much less than 0.001% of the total volume of the protected zone when the R_{PZ} is 5 NM and the h_{PZ} is 1,000 ft.

This means that while entering the protected zone is without a doubt a safety violation, the probability of actually encountering the other aircraft is extremely small. During normal operations there is no reason to enter the protected zone, but under extraordinary circumstances, the available space in the conflict zone could be used resolve the situation.

In light of the main topic of this thesis, the traditional way of determining the protected zone for traffic results in the visualization of an intentional constraint. Inside the intentional protected zone, there is a causal zone. One that when entered, results in serious incidents or even accidents. The question than becomes: how can the causal constraint be defined?

The initial starting point would be to define the space occupied by the aircraft as the causal protected zone. This is however not sufficient. Avoiding a physical collision is only one part. All aircraft generate an area of strong turbulence behind their wings, called *wake turbulence*. Figure 3.18 shows the typical position of the wake. It spreads out laterally behind the wings and descends. The magnitude of the wake is directly related to the weight of the aircraft (Gerz, Holzäpfel, & Darracq, 2002). The wake turbulence region spreads out behind and below an aircraft. Flying into this zone can result in serious damage to the aircraft entering the area.

To illustrate the importance of wake turbulence, in air traffic control, it has lead to the definition of three wake turbulence categories under ICAO: light, medium, and heavy (NATS, 2014). Based on these categories, separation minimums are defined for lighter aircraft trailing a heavier aircraft. The larger the difference in weight, the larger the required separation, either in time or distance.

Next to avoiding the wake turbulence area behind traffic, it is also important to pass well in front of other traffic in such a way that it will be not forced to pass through the ownship wake turbulence area. To do this, an area can be defined in front of the traffic that will be part of the causal constraint.

Performance Envelope Typically, the performance envelope shows the physical limits of the aircraft, and therefore represents causal constraints. In unexpected situations, a pilot should be able to use the full extent of the performance envelope if necessary. But during normal operations, parts of the performance envelope will result in less efficient operations than other areas. Especially in commercial operations, there is a lot of emphasis on saving money and cutting down on emissions by operating in a narrower operational envelope.




Figure 3.18: Wake turbulence areas behind an aircraft. The area expands equally in lateral direction with increasing distance. In the vertical plane the area also expands, but due to gravity the wake also descends behind the plane with increasing distance. The intensity also decreases with distance, the wake turbulence is stronger close to the aircraft.

As an example, all modern flight management systems use a cost index. The cost index is a value usually computed by the airline and represents the ratio of time related costs to fuel cost. A high cost index will result in faster climbs, cruises, and descents. A lower cost index will result in more fuel efficient speeds, but will increase flying time.

A cost index and its resulting speed profile represents an intentional constraint. It would be fairly trivial to add these restrictions to the performance envelope by for example using different colors to mark different regions. Usually, the autopilot tracks the values computed by the flight management system during normal operations, but the visualization in the performance envelope can help pilots monitor the performance of both systems.

Interface

Based on the above analysis, this section presents a concept of an intentional VSD. Figure 3.19 presents the enhanced version of Figure 3.15.

The causal terrain ① is still the same as in the original display, but in this version, an additional intentional layer ② is added. To bring these constraints into an angular form, both the causal terrain angle ③ and the intentional terrain angle ④ are then drawn in the performance envelope. This time, the terrain conflict zones have been drawn as filled polygons. Drawing them as lines would understate their importance with respect to the filled triangle used for the traffic constraints.

The traffic constraint is also split in the causal and intentional part. The causal constraint is represented by a small rectangle ⑤. No in-depth wake turbulence analysis is used here, the full causal constraint is approximated by a cylinder similar to the full protected zone used in the original display, but smaller. This should




Figure 3.19: A schematic ecological VSD with an explicit split between causal and intentional constraints.

be a good enough approximation of the causal constraint for demonstration purposes. The full intentional constraint is represented by the larger rectangle (6), this corresponds to the existing protected zone 'puck'. The difference in size between the causal and intentional protected zone is not to scale in this figure, in reality the difference is much larger.

The conflict zones corresponding to the causal and intentional protected zone are drawn inside the performance envelope. The smaller causal zone (8) is drawn inside the larger intentional conflict zone (7). Next to providing direct information about the causal constraint, the geometry of both traffic conflict zones can provide more information about the conflict geometry, even when the traffic is not shown on the display.

Experiment

Based on the concepts discussed in this section, an experiment has been set up to investigate how pilots would use the additional information added by creating an explicit split between causal and intentional constraints. Again, only vertical/longitudinal motion is taken into account. This experiment will be described in full detail in Chapter 5.

Experimental Evaluation of an Intentional Synthetic Vision Display

One issue that regularly occurs in the context of ecological information systems is that these systems can invite operators to migrate to the limits of system performance. This could lead to the assumption that ecological systems are thus inherently unsafe. We argue, however, that the source of this issue is tied to a modeling problem of the work domain. That is, the majority of ecological systems predominantly model the physical or causal structure of the work domain, thereby neglecting the intentional structure. Many complex socio-technical systems contain a mix of causal and intentional constraints—rules, regulations, and procedures—that contribute to safe operations of those systems. The work described in this chapter examines how visualizing intentional information in an ecological synthetic vision display affects pilot behavior, decision-making, and safety in a terrain avoidance task. An experiment with 16 professional pilots showed that adding an intentional constraint increased the clearance during terrain avoidance and gave them more insight into the terrain avoidance task, which enabled them to make better trade-offs between safety and performance.

4.1 Background

Cognitive Systems Engineering (CSE) (Rasmussen et al., 1994) and Ecological Interface Design (EID) (Vicente, 1999) paradigms are commonly regarded as guiding frameworks to develop ‘transparent’ automation, allowing human agents to monitor the machine and fluently re-direct machine activities warranted by the

This chapter is based on the following publication: Comans, J, Borst, C, Mulder, M & Paassen, MM van (2014). Risk Perception in Ecological Information Systems. In MA Vidulich, PA Tsang & JM Flach (Eds.), *Advances in Aviation Psychology* (pp. 121-138). Farnham: Ashgate.

demands of the situation. The rationale is to let the computer provide the human agent a set of constraints, rather than an explicit solution, that is directly visible on the human-machine interface, providing a problem space where any action within the constraint can solve a problem. Such an approach is seen as more robust and resilient, and can be contrasted with a 'brittle' automation design that provides optimal advice most of the time, but fails spectacularly in a few cases.

Although empirical studies have shown that such information aids enable the human to have a better system understanding and a better notion of the physical limitations, possibilities, and relationships within the work domain, humans often tend to propose actions that are sub-optimal, good enough, or even 'pushing the envelope' (Rasmussen, 1997). For example, Borst et al. (2010) showed that their ecological synthetic vision display invited pilots to routinely violate minimum terrain clearances. This could lead to the assumption that ecological information systems are unsafe and can 'promote' risky behavior.

Although we share the same concern about seeking the limits of performance, we argue that the risky behavior is tied to the scope of the work domain analysis of the system that is modeled rather than the EID framework itself. That is, the majority of ecological systems predominantly model the physical or causal structure of the work domain, thereby neglecting the intentional structure like rules, regulations, and procedures (Hajdukiewicz et al., 1999). For example, aviation safety is not only accomplished by the technical systems on board an aircraft, but also by standardized communication and coordination protocols, procedures, and airspace organization. So when the scope mainly includes the causal constraints, the 'physical structure' in the environment will be made compelling and this can cause people to pursue these causal boundaries, leaving little room to prevent accidents. On the other hand, however, when the scope mainly includes intentional constraints, the system will generally be safer, but the operational range of causal systems can be significantly limited to effectively solve problems in novel situations. The EID approach, however, can also be used to make both the causal and the intentional constraints visible and it can also manipulate the relative salience of those constraints.

In this chapter it is investigated how visualizing intentional constraints in addition to causal constraints affects pilot behavior and decision-making in a terrain avoidance task when utilizing an enhanced synthetic vision display. As such, it aims to answer the following question: when pilots are explicitly confronted with intentional constraints in addition to causal constraints, will they make 'better' decisions and will they better understand the risks involved in those decisions? The work in described in this chapter is essentially a repetition—in some aspects—of the experiment reported by Borst et al. (2010) with the addition of an explicit visualization of the required minimum safe altitude above terrain.

4.2 Intentional Constraints

Traditionally, display design is driven by task and work analyses to discover the information required on the display. The advantage of this approach is that it provides a display that is well suited for the tasks that need to be performed and usually requires limited mental effort. The downside of this approach is that the produced displays are not necessarily well suited for novel situations that were not included in the task analysis. Off-normal situations are usually extremely rare occasions, but in the context of aviation, they can have severe consequences when they are not handled correctly. Further more, designers might have an inaccurate or incomplete model of the world which will influence the task analysis and resulting design.

Ecological interface design was introduced almost 25 years ago as a design framework that aims to mitigate some of the drawbacks of task-based decision support systems (Vicente & Rasmussen, 1992). Its goal is to create a system that facilitates human creativity and flexibility to resolve novel situations unanticipated in the design of automated systems. The starting point for an ecological display is the work domain analysis. The goal of this analysis is to identify relationships and constraints within the system under study that together determine the space of possibilities in how the system can be controlled. In the display, in addition to providing the goal states of the system, the complete control space is shown to the operator. This allows operators to adapt to the specific tasks they have to perform and the context in which they appear. Although this may require a slightly higher mental workload, it can enable operators to find creative solutions to unexpected events.

The first step in a work domain analysis is to determine the scope of the analysis. When dealing with process control or the dynamic aspects of controlling vehicle motion, the analysis will be mainly focused on the physical aspects of the system under study. When looking at a complete airline operation, however, the focus will be less on the physical aspects of flying, but more on the financial and regulatory aspects, the intentional constraints.

Up until now, the majority of the work on ecological displays in aviation have focused mainly on causal constraints (Rijneveld et al., 2010; Ellerbroek, Brantegem, van Paassen, de Gelder, & Mulder, 2013). One recurring observation when experimenting with these displays is that operators will operate at the boundaries set by the constraints. In terrain avoidance experiments, for example, pilots will be able to clear the terrain obstacle with an ecological display, but they sometimes clear it with a very low margin (Borst et al., 2010). From a physical point of view, there is no problem, the obstacle is cleared. From a safety point of view, a low clearance is not favorable because it leaves little room for error to respond robustly to unexpected events, such as an engine failure during a climb. This boundary-seeking behavior may lead to the notion that systems employing EID may be inherently unsafe, because they invite operators to 'push the envelope'. Although we are concerned about this behavior, we argue that it is not necessar-

ily a property of ecological information systems.

The safety issue described above is part of a larger problem. Actions that satisfy causal constraints are not necessarily optimal or desired actions. In aviation systems, a pilot's decisions are also strongly influenced by rules, regulations, and procedures. In order to satisfy these constraints, they need to be included in the work domain analysis and presented on the display. The lack of safety found in experiments with ecological information systems can be traced back to not complying with the rules, regulations, and procedures. In the context of terrain avoidance, for example, pilots apparently seemed to disobey the minimum safe altitude restrictions to clear obstacles. To include these rules, regulations, and procedures, the scope of the work domain needs to be extended to also capture the intentional structure of the work domain. Just like their causal counterpart, intentional constraints limit action possibilities and can shape operators' behavior in the work domain.




Figure 4.1: A fully causal action space (left) and an action space with both causal and intentional constraints (right).

Figure 4.1 shows an abstract example of the resulting action space—all possible ways in which an operator can reach a goal state—with and without intentional constraints. The purely causal action space shows a situation where an operator can choose any trajectory as long as it satisfies the causal constraints. The action space with the addition of intentional constraints shows how the intentional constraints limits the space available to the operator to choose a trajectory. When operating within the intentional constraint boundaries, both sets of constraints are satisfied. Unlike with causal constraints, an operator can choose to ignore intentional constraints. A causal constraint is always a hard limit to the system performance and forms the accident boundary, whereas intentional constraints are soft limits that can be violated and therefore form the incident boundary.

In a well-designed system, the intentional constraints will be well tuned to the tasks the operator has to perform. It will allow for all actions required to deal

with a system that performs as expected. In complex domains, like the aviation domain, unexpected events will occasionally happen. These unexpected events can potentially force trajectories outside the intentional constraint boundaries. At this point, it is important for the operator to be able to clearly distinguish between causal and intentional constraints. As an example, consider a map display that shows noise sensitive areas that should not be overflown together with physical obstacles like radio towers that can not be over flown. Under normal situations, a pilot will choose a trajectory that avoids all areas that need to be avoided. Under emergency conditions—for example, when an engine fails—the pilot can choose to ignore the noise sensitive areas and fly over them if it will increase the safety of the flight. In the meantime, the pilot will still have to avoid the radio towers. Overflying the noise sensitive areas does not even have to be an all or nothing situation, the pilot can still decide to try to minimize his impact and find a trajectory that produces the least amount of nuisance.

It will be important for the operator to be able to clearly distinguish between both sets of constraints. As discussed before, only showing causal constraints can lead to boundary-seeking behavior. On the other hand, only showing intentional constraints would decrease the ‘apparent solution space’ which can result in situations where there is no visible action space left, but the remaining causal space is hidden. A non-display related example of this situation can be the Flight Envelope Protection system implemented in all modern Airbus aircraft. This system limits the actual control inputs given by the pilot before they are sent to the control surfaces to make sure the aircraft does not surpass any structural limitations. The advantage of this system is that pilots can not accidentally overstress the aircraft, and this works well in most day-to-day situations. However, in 1985 a crew of a China Airlines Boeing 747 ended up in a situation where the only way out of a steep dive was to actually overstress the aircraft. The aircraft got damaged because of this maneuver, but was still flyable (NTSB, 1985). If a Flight Envelope Protection system had been in place to enforce the theoretical design limit, the aircraft would probably have crashed into the sea. This example shows that while under most circumstances pilots should be conforming to intentional constraints, situations may develop where they need to violate those constraints and focus on the causal limits of their vehicle. In this light, we argue that operators should be presented with both sets of constraints clearly visible. By doing this operators will be able to see both required behavior tied to the causal constraints and expected behavior tied to the intentional constraints. With this they should be able to make good trade-offs between satisfying intentional and/or causal constraints. The following section describes an experiment to test this hypothesis.

4.3 Experiment Method

Borst et al. (2010) introduced an ecological terrain awareness display. By combining a synthetic terrain display with a visualization of the aircrafts’ climb per-

formance, pilots were able to quickly and correctly determine their options when having to perform an emergency climb over a terrain obstruction. By visualizing the climb performance data, the number of terrain collisions reduced to zero. The terrain clearance, on the other hand, decreased significantly, indicating that pilots were flying on the edge of the safe flight envelope. As discussed above, we argue that this is a result of limiting the scope of the work domain analysis to causal constraints only. To test this hypothesis, we redesigned the display to include an intentional constraint: the Minimum Safe Altitude. By showing this constraint, in addition to the causal constraints, it was expected that the minimum terrain clearance would increase and that pilots would opt for more robust and thus safer control strategies to resolve terrain conflicts. By placing pilots in situations that are difficult to resolve without violating intentional constraints, we also expect that they will deliberately violate the constraints while still maximizing safety.

4.3.1 Subjects

A mix of seven recently graduated and nine commercial pilots participated. Their average age was 40 years (SD 16.13) with an average experience of 3,370 flight hours (SD 3,923.07). Four of them were TU Delft/NLR test pilots. One of them was a former military F16 test pilot. Business jet pilots were chosen to be the group of commercial pilots because they are usually used to more dynamic and changing operations than air transport pilots. The recently graduated pilots were chosen for their lack of procedural habits.

4.3.2 Apparatus

The experiment was conducted in a fixed base flight simulator. The display was shown on an 18 inch monitor located in front of the pilot. An outside visual consisting of fog and cloud fragments was projected on the front and side walls to provide some sense of motion. The aircraft model was controlled by a right hand hydraulic side stick and a throttle quadrant on the left. The throttle contained the trim switch, autopilot disconnect switch and Horizontal Situation Indicator center button. A mode control panel on top of the instrument panel was used to control the Horizontal Situation Indicator course. A non-linear six degree of freedom Cessna 172 model was used for the experiment. Pitch, roll and throttle commands were directly controlled by the pilot. To compensate for the lack of rudder pedals, a side slip controller was implemented to minimize side slip and engine torque effects. Two different performance settings were used during the experiment. During normal performance runs, the model operated in a normal International Standard Atmosphere giving the normal performance at the altitudes flown. In the reduced performance mode, the aircraft performance corresponded to what would be expected at low density altitude conditions. In this mode, climb performance decreases significantly with altitude. In high performance conditions, the maximum climb angle was always between 7.5 degrees and 6 degrees,

while in the low performance conditions the maximum climb angle deteriorated from 4.5 degrees to 1 degree while climbing to the final altitude.

4.3.3 Display

The display used in the experiment is shown in Figure 4.2. It is similar to a Garmin G1000 NAV III augmented with a synthetic vision system. Three additional cues were added for the baseline display. The flight path vector indicating the geometric flight path which provides immediate feedback to the pilot about his current trajectory. If the flight path vector is pointing at the synthetic terrain, the aircraft will eventually impact the ground at that position. If the flight path vector points above the terrain, the aircraft will clear the terrain. The maximum sustained climb angle at full power is shown by a wide green bar (no. 3 in Figure 4.2). This indication immediately shows whether the aircraft is able to clear the terrain at maximum climb performance. When this line is below the synthetic terrain, the pilot will not be able to climb over it and will have to use a different maneuver. Similarly, the current maximum climb angle (no. 2 in Figure 4.2) is the maximum climb angle that can be sustained with the current power setting/throttle position. With this indication, the pilot gets immediate feedback about his current climb performance and is able to climb at lower power settings while still being sure to clear the terrain. The result is an ecological baseline display showing the causal constraints of the terrain avoidance task.

The baseline display is augmented with an intentional layer indicating the minimum safe clearance above the terrain (no. 4 in Figure 4.2). This layer is created by shifting the synthetic terrain up and drawing it in amber behind the physical terrain. In this way, the layer has the same relationship to the flight path vector as the original terrain. If the flight path vector is above the layer, the terrain clearance will be at least the required minimum clearance. An additional advantage of adding the intentional layer to the display is that it improves distance perception of terrain features, something that is very difficult in traditional synthetic vision displays. Because the layer has a fixed height, its resulting thickness on the display is an indication for the distance to the terrain. Even though the relationship between distance and thickness is non-linear, it can be used for a crude estimation of the actual distance of terrain features.

4.3.4 Scenario

The scenario used for the experiment consisted of an artificial terrain with a number of narrowing fjords. The base of the fjords was at sea level, the tops were around 3,000 *ft*. The pilots were told that they flew into the wrong fjord on their way to the airport and are low on fuel by the time they realize their mistake. Each experiment run started in one of the predetermined initial locations at an altitude below the surrounding fjord tops. From this starting position it was impossible to get to the airport without climbing over the terrain. A navigation beacon was




Figure 4.2: The synthetic vision display, showing: 1) the flight path angle, 2) the current maximum sustained climb angle, 3) the maximum sustained climb angle at full power, and 4) the intentional terrain constraint.

placed in-between the initial aircraft's position and the airport to provide a navigation reference that could be reached in three to five minutes eliminating the need for long cruise segments to reach the airport.

Pilots were instructed to navigate to the waypoint and keep clear of the terrain in a way they considered safe and comfortable. Pilots were not given minimum altitude instructions, but they received a map of the area that showed 4,000 ft as the minimum safe altitude for the area they were navigating.

4.3.5 Independent Variables

The experiment used three within-subject variables each having two levels: display configuration, scenario difficulty, and aircraft performance. The display configuration was either a baseline display without the intentional layer, or a display with the intentional layer visible.

Two levels of difficulty were used, as illustrated in Figure 4.3. Easy conditions started with enough margin for a straight climb towards the beacon. The hard conditions required immediate full power and a maximum performance climb to avoid transgressing the intentional layer.




Figure 4.3: Scenario difficulty. Easy conditions (left) start with enough margin to easily perform a straight climb above the intentional layer. Difficult conditions (right) start with the maximum climb angle close to the top of the intentional layer.

Two levels of aircraft performance were used. With normal performance, the climb performance remained almost constant during the first 4,000 ft of the climb. With the reduced performance, climb performance severely deteriorated above 2,500 ft.

The order in which the display configuration was tested has been treated as a between-subject variable. One group of pilots started with eight runs using the baseline display and then moved to the augmented display. The other group started with the augmented display and moved to the baseline display afterward.

Table 4.1: Condition Overview

Name	Difficulty	Performance	Display Configuration
Easy High Performance	Easy	High	Baseline
	Easy	High	Intentional
Hard High Performance	Hard	High	Baseline
	Hard	High	Intentional
Easy Low Performance	Easy	Low	Baseline
	Easy	Low	Intentional
Hard Low Performance	Hard	Low	Baseline
	Hard	Low	Intentional

4.3.6 Dependent Measures

Three objective measures are used to quantify pilots behavior. The minimum terrain clearance and number of clearance violations are used as a measure of safety. The final altitude is used as a measure of procedural compliance.

Next to these objective measures, notes were taken during each run describing the pilots choices. After each run, pilots provided feedback about their strategy and choices.

4.3.7 Procedure

The experiment started with a training phase to familiarize the pilots with the display, flight controls, and aircraft model. During training, the pilot could fly around freely in a training scenario and was given an explanation off the display features. Once the pilot was familiar with the added features, the measurement runs began. No task specific training was done, only display familiarization.

The pilots were divided in two groups, one started without the intentional additions, the other started with the intentional additions enabled. Initially there were 18 pilots divided equally among both groups, but during the experiment, two pilots failed to complete the experiment. This resulted in nine pilots starting without intentional additions and seven pilots starting with intentional additions.

Each pilot flew eight conditions per block (two difficulty levels, two performance levels, and two repetitions). The conditions were randomly distributed based on a Latin square matrix to avoid effects based on the condition order. Different Latin squares were used for both blocks. At the end of the experiment, 256 samples were collected

Before each measurement run started, the pilots were instructed to set the throttle to the trim position. Once the run started, the autopilot maintained al-

titude and airspeed for five seconds. During this time, the pilot was asked to observe the situation. After the autopilot disconnected, the pilot had to confirm the disconnect by pushing a button and navigate the aircraft towards the navigation beacon. Once they were sufficiently close to the beacon, the run ended and the pilots provided feedback about their strategy during the run.

At the end of the experiment the pilots were asked to complete a questionnaire to evaluate the overall experiment.

4.4 Results

4.4.1 Minimum Clearance




Figure 4.4: Minimum clearance per condition and repetition. The whiskers indicate the lowest and highest datum within the 1.5 IQR of the lower and upper quartile.

The minimum clearance for all conditions and repetitions is shown in Figure 4.4. The box plots show that the spread of the minimum clearance decreases and the lower bound shifts towards the 1,000 ft line when the intentional con-

straint is shown. Both repetitions of the Easy High Performance condition show quite some variation in the data, but with the intentional constraint visible all clearances are well above 900 *ft*, except for one outlier, a pilot who chose to fly at less than 500 *ft* but indicated after the run that he felt compelled to increase his clearance in the next runs. In the Hard High Performance condition, the spread in clearances is large when the intentional constraint is not shown. The lower bound is at 500 *ft* or lower. This changes significantly when intentional constraints are shown, the clearances cluster around the top of the intentional layer with the exception of a few outliers. In the Easy Low Performance condition, a similar effect is shown, clearances are more clustered and generally above 1,000 *ft*. In the final Hard Low Performance condition, the change in spread is less than in the other conditions, but there is a clear shift towards the top of the intentional layer.

To further analyze the objective clearance measure, a repeated measures Analysis of Variance (ANOVA) has been performed. To simplify the analysis, the results from the repetitions were averaged per condition resulting in eight data points per pilot. This assumption should not distort the results too much since the majority of the pilots showed reasonable consistency between the repetitions in terms of strategy and minimum clearance. The increase in clearance is confirmed with the ANOVA showing a significant effect for the display type ($F(1,14)=5.44$, $p < 0.05$). No significant interactions were found for the difficulty and performance variables with the display type.

During the post experiment questionnaire one of the pilots remarked that his strategy without intentional layer had changed by starting with the intentional layer enabled. Analyzing the display order did not show a significant effect ($F(1,14)=0.687$, $p=0.064$) but was close to the $p=0.05$ significance level. Further research with more pilots might reveal an actual influence of the display order.

The only other significant effects were the main effects of the scenario complexity ($F(1,14)=23.446$, $p < 0.01$) and performance ($F(1,14)=15.332$, $p < 0.01$) variables. This confirms that the task became more difficult both with decreasing performance and with increasing difficulty.

4.4.2 Clearance Violations

Figure 4.5 shows the clearance violations at two clearance altitudes, one at the intentional clearance of 1,000 *ft* and one slightly lower at 900 *ft* to filter out minor violations. In the Easy conditions there is a clear distinction between the runs with and without the intentional constraint visible. At the 900 *ft* level, the violations drop to 1 in 32 runs with the intentional constraints visible at both performance levels.

In the Hard conditions there is also a decrease in violations when the intentional constraints are shown, but the difference is smaller than in the Easy conditions.

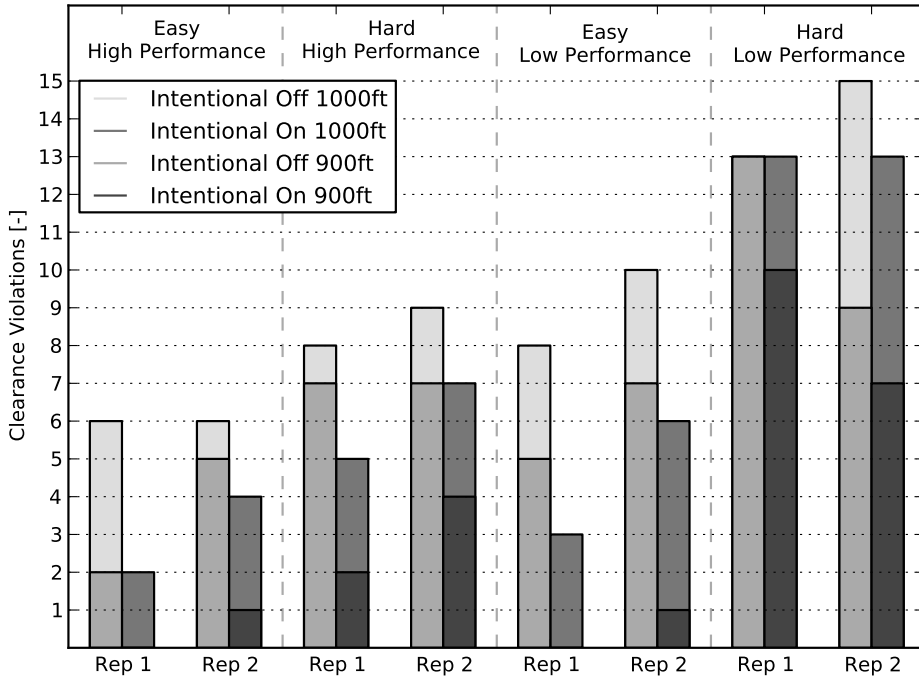


Figure 4.5: Clearance Violations per condition and repetition.

4.4.3 Final Altitude

For the final objective measure, Figure 4.6 shows the final altitude at the end of the run for all conditions. With the average terrain height at 3,000 ft, the 4,000 ft level indicates the minimum safe altitude above the terrain. The majority of all runs end at or above 4,000 ft, but mainly in the Hard conditions without the intentional constraint visible, a number of pilots end up below 4,000 ft. With the intentional constraint visible, all pilots except a few outliers have a final altitude above 4,000 ft and show a much lower spread compared to the runs without intentional constraints.

4.4.4 Strategy Adaptation

During the experiment, pilots were not instructed to use a specific strategy. This resulted in a wide range of different strategies employed during the experiment. In order to analyze how pilots' strategies changed when the intentional constraints are shown, all time traces were classified according to the strategy that was used.




Figure 4.6: Final Altitude per condition and repetition. The whiskers indicate the lowest and highest datum within the 1.5 IQR of the lower and upper quartile.

Two main criteria were used: the horizontal trajectory and the combination of power and pitch.

Five general trajectory strategies were used by most pilots: 1) flying straight until clear of the terrain, 2) flying directly towards the navigation beacon, 3) orbiting until clear of the terrain, 4) turning towards the lowest point of the terrain, and 5) flying parallel to the ridge. In terms of power and pitch, three different strategies were identified: 1) climbing at full power with the maximum climb angle, 2) climbing with enough power to achieve a clear margin above the terrain, and 3) climbing with the power required to keep the climb angle above the intentional layer.

When looking at the actual changes in strategies when the intentional layer was visible, no real differences are observed between the High Performance and Low Performance conditions. During the Easy scenarios, however, approximately 20 per cent of the runs resulted in a changed horizontal trajectory, and 29 per cent switched to a strategy where they used a combination of power and pitch to keep

the flight path vector just above the intentional layer. Not all trajectory changes can be attributed to the addition of the intentional layer. Some runs provided no clear indications of the reason for using a different strategy. The changes in engine power and path on the other hand are all directly influenced by the addition of the intentional layer since pilots actively used the intentional layer to choose an engine power setting and climb angle.

In the Hard scenarios, the way in which strategies changed differs substantially. During 40 per cent of the runs, pilots changed their trajectory when presented with intentional constraints. The changes observed in the Hard scenarios can be divided in two broad categories. A number of pilots changed to a strategy where they would actively search for the lowest point in the intentional layer and used it to direct their aircraft to these low spots. Another group used the intentional layer to determine that their previous strategy would not work and switched to a different strategy. In terms of power settings, all pilots used maximum power in combination with the maximum climb angle for all runs and no changes were observed in this strategy for the Hard conditions.

4.4.5 Questionnaire Data

At the end of the experiment each pilot answered four questions about their experience with the intentional constraints. Fourteen pilots responded positively to the question if the intentional addition makes the terrain avoidance task clearer. The main reasons they gave were the increase in situation awareness, better perception of height above the terrain, and a reduced mental load when planning a terrain avoidance maneuver. One pilot stated that the procedural addition makes the task more restrictive. Another pilot, however, indicated that the intentional layer was distracting and ignored the additional information.

The question whether the intentional layer changes their strategy, was confirmed by 12 pilots. The majority indicated that the additional information presented enables them to quickly see the lowest regions of the terrain and also provides immediate information about the engine power required to reach the safe altitude. One pilot also noted that he intentionally put his flight path vector in the intentional layer because he felt comfortable with a lower clearance and the layer gave him a good indication what pitch angle was required to satisfy this clearance. Among the pilots answering no, two indicated their strategy remained the same but they used the intentional layer as a confirmation of their strategy.

Every pilot, except the one that ignored the intentional constraints, felt that the perceived level of safety increased when the intentional constraints were visible. This mainly happened because the safety margin becomes explicit in the display enabling them to directly assess the risk involved. One pilot noted that adding the procedural constraints takes the guessing out of flying. The final question whether the pilots considered the procedural additions useful was answered positively by all but one pilot and is in line with the answers to the previous questions.

4.5 Discussion

The main objective of the experiment was to investigate whether visualizing intentional constraints in addition to causal constraints helps pilots in making better decisions. Better decisions in the context of this experiment meant respecting the minimum safe clearance as much as possible. Analysis of the objective clearance parameter confirmed this hypothesis. There is a significant increase in minimum clearance when comparing the intentional addition to the baseline display.

To get more insight in this change in clearance, Figure 4.4 presented the clearance values of all pilots per condition and Figure 4.5 showed the number of pilots that flew below the minimum safe clearance. In Easy High Performance conditions all but two pilots flew above the minimum safe clearance when the intentional layer was enabled. One pilot deliberately ignored the layer and accepted a clearance of less than 500 *ft*. The second pilot violated the minimum clearance and kept his flight path vector close to the intentional layer, resulting in a very brief excursion just below minimum clearance. In the post run feedback and through observations it became clear that with sufficient performance and margin pilots will treat the amber intentional constraint as if it is actual terrain and will avoid it.

In the Hard High Performance conditions, the same strategy surfaces. Pilots are more inclined to try to meet the minimum clearance with the intentional layer enabled. There are more violations than in the Easy High Performance condition, but they are all except one minor violations. By steering into the top of the intentional layer pilots could make an informed choice about sacrificing a small amount of clearance for a quicker route towards the airport.

In the low performance conditions, the same trends can be observed as in the high performance conditions. The main difficulty in the low performance conditions is that the climb performance significantly decreases during the climb. A number of pilots failed to note that there was not enough margin between the maximum climb performance and the intentional layer, but even in these cases the intentional layer shows that they are closer to the top of the minimum clearance and can continue in a relatively safe manner.

An analysis of the data showed that with the baseline display a number of pilots flew with less than 500 *ft* clearance. For two pilots this was a deliberate choice, the other pilots were mainly unaware of their actual clearance. Not counting the deliberate violations would leave approximately 10 instances where pilots flew below 500 *ft* above the terrain without being aware of this. This number is only a third of what Borst et al. (2010) reported in a previous experiment with an ecological terrain awareness display. The main reason for this difference is probably the fact that pilots had more freedom to perform an escape maneuver in this experiment while in the previous experiment pilots were asked to fly straight as much as possible.

The influence of the intentional constraint on the final altitude, shown in Figure 4.5, is also an interesting result. While the majority of the pilots definitely aims

to fly at 4,000 *ft*, there are a number of pilots that are slightly below this altitude. Enabling the intentional layer shifts the final altitude for all but five pilots, above 4,000 *ft*. As long as there is even a small part of the intentional layer above the horizon, indicating that there is still terrain that would be cleared by less than 1,000 *ft*, pilots will have a tendency to climb until the intentional layer is below the horizon. Without the intentional layer, some pilots accept an altitude slightly below 4,000 *ft*. One reason for this could be that once the actual ridge is cleared, the urgency of the situation decreases because the majority of the visual feedback is gone. With the intentional layer, this information will still be obvious as long as it is still relevant, i.e. as long as the clearance is less than 1,000 *ft*.

The way in which pilots adapted their strategy when presented with intentional constraints depended on the difficulty of the scenario. In the Easy conditions, the most obvious adaptation is when pilots directly use the intentional layer to determine the required engine power setting and corresponding climb angle to clear the terrain with the expected margin while preserving fuel as much as possible. In other instances, pilots used the added information to confirm their strategy. While doing the same as in the corresponding baseline conditions, a number of pilots indicated after the run that they felt more confident and were more aware of the safety margin. Finally a number of pilots did not change their strategy, but used the intentional layer to fine-tune it. For example, a number of pilots that performed a straight climb to 4,000 *ft* before making the turn towards the navigation beacon followed the same strategy but used the intentional layer as a cue to decide when to initiate the turn towards the navigation beacon.

In the Hard conditions, there is less leeway for the pilots, the scenarios forced pilots to quickly use all available performance. In these situations, the intentional layer in combination with the indication of the maximum climb performance gave the pilots an instant assessment of their options. With this information, a number of pilots decided to change strategy and for example make a few climbing turns. Other pilots stuck with their strategy and used the intentional layer to maximize their clearance. Either by keeping their flight path vector out of the intentional layer or at least in the top part. By doing this, they sacrificed some clearance, but were able to clearly assess the severity of this violation. Finally, a number of pilots used the intentional layer to guide their flight path towards the lowest parts of the intentional layer keeping their clearance close to or above 1,000 *ft*.

Through the observations, post run feedback of the pilots, and the questionnaire it became clear that the majority of the pilots used the intentional addition to either improve or change their strategy in solving the terrain conflict. The way in which they fit the intentional addition in their strategy can differ but they all indicated that it enhanced their analysis and awareness of the task at hand.

One drawback of the current intentional representation surfaced during the experiment. Once a pilot flies below the minimum clearance, the whole top part of the display is filled with the amber color. Once this happens, it is no longer possible to directly perceive the difference between a minor violation close to the top of the intentional layer or a dangerous violation close to the terrain. In the fu-

ture this could be resolved by using different shades of amber to indicate different clearance levels in the intentional layer.

Looking back at the hypothesis, this experiment has provided some insights into the usefulness of visualizing intentional constraints. The freedom given to the pilots to implement their own preferred solution resulted in the adoption of a wide variety of strategies, but the majority used the additional information to their advantage. Some used it to verify their strategy, some to fine-tune, and some used it to completely change their strategy. All of this is based on the visualization of the intentional constraint. Explicitly showing information that is otherwise presented through different modalities—charts, knowledge, other instruments, ...—enables pilots to focus on immediate use of this information instead of mentally piecing together all the pieces. In line with this, pilots were able to use strategies that would otherwise be impossible without elaborate calculations. As an example, some pilots decided to accept a small clearance violation in order to fly a shorter route. They would put their flight path vector in the intentional layer, but close to the top. In this way they could be certain that they would still have a sufficient margin above the terrain. Without the intentional layer, these types of strategies would require the pilots to accurately determine their position with the help of a map and a navigation beacon, calculate the climb performance and check if the performance is sufficient. These types of calculations quickly become time consuming and are not suited when quick decision making is required.

Finally the experiment showed that the current representation subtly drives pilots towards the intended margin. The pilots were not explicitly instructed to never fly into the intentional layer, they were only instructed on how the display worked and the kind of information presented to them. Even without explicit instructions, pilots showed a natural tendency to stay above the intentional layer if this was feasible.

4.6 Conclusion

The main objective of the experiment was to investigate whether the effect of visualizing intentional constraints in addition to causal constraints helps pilots in making better and more robust decisions. Better decisions in the context of this experiment meant respecting the minimum safe altitude as much as possible.

Analysis of the objective clearance parameter confirmed this hypothesis. There was a significant increase in minimum clearance when comparing the intentional addition to the baseline display that only portrayed causal constraints. Visualizing the minimum safe altitude resulted in better compliance with the intentional constraint and resulted in the pilots being able to make better trade-offs between performance and safety, confirming our hypothesis.

Based on the experiment results, we can conclude that incorporating intentional constraints can shape operators' behavior and can shift from operating close to the physical constraints towards a point closer to the intentional boundaries.

Experimental Evaluation of an Intentional Vertical Situation Display

Over the last decades a number of technologies has been developed to assist crews in detecting and avoiding external threats such as traffic, terrain, and weather. Although successful, they generally only provide support for short term Conflict Detection & Resolution (CD&R). In a future where long term trajectory planning will probably be favored over direct procedural control (JPDO, 2011), medium to long range collision avoidance tools will become more important.

When focusing on terrain and traffic, the relevant systems providing support are the Traffic Collision Avoidance System (TCAS) and the Enhanced Ground Proximity Warning System (EGPWS). Both systems provide a great deal of assistance in avoiding terrain and traffic, but they still have a few drawbacks. Because of the rather short term lookahead of both systems, it is difficult to properly anticipate conflicts and plan ahead to avoid them. This short lookahead also means that pilots need to respond very fast to alerts, raising stress levels considerably. As an example, with TCAS, about 48 seconds before an imminent traffic event the system issues a traffic advisory (TA). At that point the pilot should visually confirm the location of the traffic and take appropriate actions. At 35 seconds, a resolution advisory (RA) is issued, and the pilot should respond immediately by following the directions of the TCAS system. This means that a pilot has approximately 13 seconds to assess an unanticipated event (FAA, 2011).

In addition, both TCAS and EGPWS are isolated systems. They work on a fixed set of rules that apply to the scope of either system, but they are not aware of each other. In other words, they have no knowledge of the other applicable constraints. When multiple alerts occur at the same time, pilots will have to prioritize them based on their knowledge of both systems and have to quickly decide which alert has the highest priority. Under these circumstances, pilots' problem solving skills are more important than procedural knowledge (Hill, 1994).

The most important shortcoming of these types of systems, however, is that they hide the rationale for generating alerts. Pilots are expected to follow the advisory generated, without being able to validate and check the proposed solution. In most circumstances, this is not an issue. But in unanticipated and unexpected situations, understanding the rationale behind the decision and having the option to consider alternative solutions can be crucial. Especially when these systems would generate counter-intuitive solutions, or when circumstances not anticipated by the system designers make the suggested solution dangerous or even impossible.

The main drawbacks of current systems highlighted above can be summarized as follows:

1. Short lookahead gives pilots little time to anticipate difficult and potentially dangerous situations,
2. Different sets of constraints for different systems make interaction between different alerting systems complex, and
3. Hidden rationale keeps pilots out of the loop.

In this chapter, an approach to integrate traffic and terrain constraints adhering to the strategies outlined in the previous chapters will be presented. The starting point for the conceptual display is an existing Enhanced Vertical Situation Display (EVSD). Section 5.1 will present the rationale behind the EVSD and show the improvements that have been made. Next, Section 5.2 will explain all relevant display features in detail. Section 5.3 will describe the results of a preliminary validation experiment that has been performed.

5.1 Background and Motivation

Rijneveld et al. (2010) developed an ecological Vertical Situation Display (VSD) that presents a combined set of terrain and traffic constraints that addressed the three issues described above. They started with a baseline VSD showing traffic and terrain in the same display. By its nature, this display addresses the short lookahead time, the first issue in the previous section, by using an adjustable range on the horizontal and vertical axis. The ecological addition to the VSD consists of conflict zones that show all the combinations of ownship speed and flight path that will lead to a collision. This enables pilots to efficiently avoid obstacles, but when combined with an alerting system, can also visualize the rationale behind the traffic or terrain alert they generate, addressing the third issue. By combining terrain and traffic in one display, both sets of constraints are visible

and the coupling between the different alerting systems is shown, mitigating the second issue described above.

A between-subjects experiment was conducted that compared the ecological design with a baseline VSD. Twelve professional pilots were divided in two groups: one using the baseline VSD and the other using the EVSD. Both groups flew twelve scenarios consisting of traffic, terrain and mixed conflicts. An increase in self-reported situation awareness and a significant decrease in workload were reported. Both displays gave rise to conflicts with the constraints, as shown in Table 5.1. With the baseline condition, subjects violated the traffic constraint in 5 out of 66 possible events. With the EVSD, 3 out of 66 traffic violations occurred, and 4 out of 42 terrain safety constraints were violated by following the terrain too closely. No terrain crashes were reported, however.

Table 5.1: Safety in terms of traffic Protected Zone intrusions, terrain safety zone intrusions and terrain crashes. Between brackets are the number of possible conflicts (Rijneveld et al., 2010)

	VSD	EVSD
Traffic PZ intrusions (66)	5	3
Terrain safety zone intrusions (42)	0	4
Terrain crashes (42)	0	0

The reason for the additional terrain conflicts with the EVSD was not investigated in depth, but was attributed to pilots trying to follow the assigned trajectory as close as possible. Since both the baseline as well as the EVSD condition showed the trajectory there is probably more in play. In the next subsection we will use the schematic of Figure 5.1 to discuss in more detail what could have happened.

5.1.1 The Enhanced Vertical Situation Display

The baseline VSD used by Rijneveld et al. (2010), illustrated in Figure 5.1 (a), only presents the physical location of the terrain and traffic constraints, but not in the same way. The terrain is shown as two colored regions. One for the causal (physical) constraint, the dark shading representing the physical terrain. The other for the intentional constraint, a safety buffer with a certain thickness represented by the lightly shaded area. The traffic constraint, on the other hand, consists of an aircraft symbol and a rectangle showing a side view of the protected zone. Unlike the terrain constraint, here the distinction between causal and intentional constraints is less salient. The outside boundary represents the boundary of the intentional constraint, but even though we know the causal constraint will be located around the aircraft symbol, its actual shape and size is not visible.




Figure 5.1: The representations of the baseline and enhanced terrain and traffic constraints: (a) the VSD and (b) the EVSD.

This difference in the representation of the constraints—a filled polygon versus a rectangle—already shows the terrain as a *more prominent obstacle* than the traffic and in the experiment of Rijnveld et al. (2010) may have lead pilots, when using the baseline VSD, to favor solutions that keep a safe distance from the terrain. But there is another significant difference in the presentation of the constraints: Although they both represent the physical location of the constraints, they do not provide any information on the constraint *dynamics*.

With the terrain, these dynamics are trivial. The physical position of the terrain does not change, so every maneuver of the aircraft will be relative to this terrain. Once a pilot ensures that his flight path vector is aimed above the terrain, the terrain will always be cleared since the flight path indicator shows the geometric climb angle. This does not hold for traffic, however. The traffic itself moves *relative* to the ownship. Even if the direction of this motion is indicated, it is not straightforward for a pilot to relate the flight path vector of the ownship to the motion of the intruder since this motion depends on the relative velocity of the intruder with respect to the ownship. This relative velocity can not be easily determined from the display.

This makes it easier to predict the ownship's motion with respect to the terrain but may also bias pilots to choose for a maneuver away from the terrain. In this way, the baseline VSD again over-emphasizes the importance of the terrain constraint and gives the false impression that a maneuver away from the terrain is, overall, the safest way out of a situation.

The EVSD succeeds in addressing some of the described shortcomings. In the EVSD, Figure 5.1 (b), the traffic constraint is shown as a conflict zone (Van Dam et al., 2008). This conflict zone indicates the range of ownship velocities that will result in the ownship entering the protected zone of the intruder. There are two main advantages of showing this conflict zone. First, it incorporates the dynamics of the constraints on ownship motion caused by the intruder, and makes them explicit so no mental processing is required by the pilot. The relative motion of the intruder is captured by the conflict zone in such a way that it only relates to the ownship velocity and flight path vector. Because of this, all a pilot has to do is to change flight path and/or speed to resolve a conflict.

The second advantage is that it shows the constraints in an angular form (Borst et al., 2006). By representing *both* traffic and terrain constraints in an angular form, it is more straightforward for pilots to find solutions that satisfy both constraints if they are available since all areas can be added together. Finding a spot outside all zones will result in a conflict-free trajectory. The results of the EVSD experiment indeed reflect this by showing better situational awareness and lower workload. This does not necessarily explain the increase in terrain conflicts, however.

There still is a problem with the visualization of the constraints in this implementation of the EVSD. The terrain is presented with a clear distinction between causal and intentional constraints. Even though the angular constraint only shows the intentional constraint, the pilot can easily observe the available margin below this angle. While the preferred trajectory avoids violating the intentional con-

straint, it is still an option if the pilot deems this necessary.

The conflict zone, on the other hand, is drawn based on the protected zone—an intentional constraint—without an explicit visualization of the causal constraint. The protected zone creates a rather large buffer around an intruder. The actual area occupied by the intruding aircraft, combined with a possible allowance for wake turbulence, is still much smaller compared to the size of the complete protected zone. This means that, while not preferred, it may often still be possible to choose trajectories that cross the protected zone which still respect the causal constraint and not lead to a collision or any other incident, except a decreased separation. Not unlike flying through the intentional terrain constraint while staying clear of the physical terrain.

Figure 5.1 (b) shows the effect of not taking into account the distinction between causal and intentional traffic constraints. Showing the traffic constraint as one solid polygon misrepresents the weight of this constraint. Even though pilots might be aware of the nature of the conflict zone, there is no direct, reliable way of knowing how the causal constraint relates to the intentional constraint. This way of visualizing the traffic constraint could explain the unexpected increase in terrain conflicts with the EVSD.

5.1.2 Visualizing Intentional Constraints

To present the traffic causal and intentional constraints in the same way as the terrain constraints, a smaller causal protected zone can be defined that better represents the physical aircraft size and has an allowance for the severe wake turbulence area trailing the intruder. This causal protected zone can then be used to calculate the causal conflict zone, the range of velocities that will lead to a collision or a serious incident. This conflict zone has more similar properties to the causal terrain constraint, violation of either of them will result in an accident.

Figure 5.2 shows an example of what this distinction can look like. It shows the same situation as in Figures 5.1 (a) and (b), but it more explicitly distinguishes between causal and intentional constraints for both the terrain and intruder.

Visualizing all constraints like this would present a more *complete picture* of the situation, making clear where possible trade-offs can be made. The small white open space in the middle shows that there is some opportunity to clear both the traffic and the terrain with proper margins. Unlike the EVSD, both sides of the white area are now delimited by constraints of the same *severity*. It becomes more clear that both constraints have an intentional margin that can be used if the pilot deems necessary.

An experiment was designed to evaluate whether making this explicit split can indeed improve pilots' understanding of mixed terrain and traffic situations, and reduce the number of terrain incidents. The experiment was based upon the experiment of Rijnveld et al. (2010), using their EVSD as a baseline condition. In the next section the experiment displays are discussed in more detail. After this, the experiment setup is discussed.




Figure 5.2: The representations of the terrain and traffic constraints with an explicit split between causal and intentional constraints.

5.2 Experiment Displays

During the experiment runs, subjects used a Primary Flight Display (PFD) and a VSD (either the baseline VSD or the Intentional Vertical Situation Display (IVSD)). A traditional Navigation Display (ND) was not used.

5.2.1 Primary Flight Display

Since a PFD is a pilots' primary reference when operating an aircraft, a generic PFD was used next to the experiment display. The display is shown in Figure 5.3. The PFD does not show any information related to the terrain and traffic clearance task of the experiment. It only serves as a reference for the basic flight parameters.

The center of the PFD shows the traditional pitch ladder and flight path vector with the corresponding bank attitude indicator above. An altitude tape with vertical speed indicator is presented on the right. The barometer setting below the altitude tape is not used. On the left of the pitch ladder there is an airspeed tape showing calibrated airspeed with an optimal climb speed indication marked with an 'x'. The true airspeed is shown below the airspeed tape.

On the left of the airspeed tape the engine indications are shown. The top bars and digits provide both engines fan speeds, N1. And the bottom two digits indicate the turbine speed N2.

On top of the airspeed and altitude tape, the target airspeed and target altitude are shown in magenta. They depend on the target airspeed and altitude of the next waypoint. Next to the numerical representation, a magenta bracket also indicates the target value on the airspeed and altitude tapes.



Figure 5.3: The generic PFD, based on a Garmin G1000.

5.2.2 The Baseline VSD

Traditional VSD

The baseline VSD display is shown in Figure 5.4 and represents the EVSD as developed and tested by Rijnveld et al. (2010). A traditional VSD shows a schematic side view of the own aircraft, together with the flight plan and terrain and traffic obstacles in the vertical plane.

The origin of the VSD display is the ownship aircraft symbol shown in yellow ①. This symbol has a fixed position. The aircraft altitude is indicated with an altitude tape on the left hand side ② in combination with the horizontal grid lines at fixed altitude increments. When the aircraft altitude increases, both the altitude tape and horizontal grid lines will move down and vice versa. The altitude associated with the active waypoint is indicated by a magenta bracket ③.

The horizontal axis shows the along track distance ④ at which obstacles and waypoints are currently located. It is important to note that the horizontal and vertical scales are different. With an aspect ratio of approximately five, vertical distances are shown to be about five times larger than they actually are. The advantage is that obstacles become more pronounced, but the disadvantage is that some correlation is lost with respect to the PFD, especially for the angles.

A yellow arrow represents the aircraft flight path vector ⑤. The length of this arrow scales with true airspeed and its angle corresponds to the flight path angle, compensated for the display aspect ratio. The right hand side of the display shows the vertical projection of the velocity vector, representing the actual rate of climb.

The top part of the display shows the indicated airspeed ⑥ associated with the current flight path vector. The yellow indicator shows the current indicated

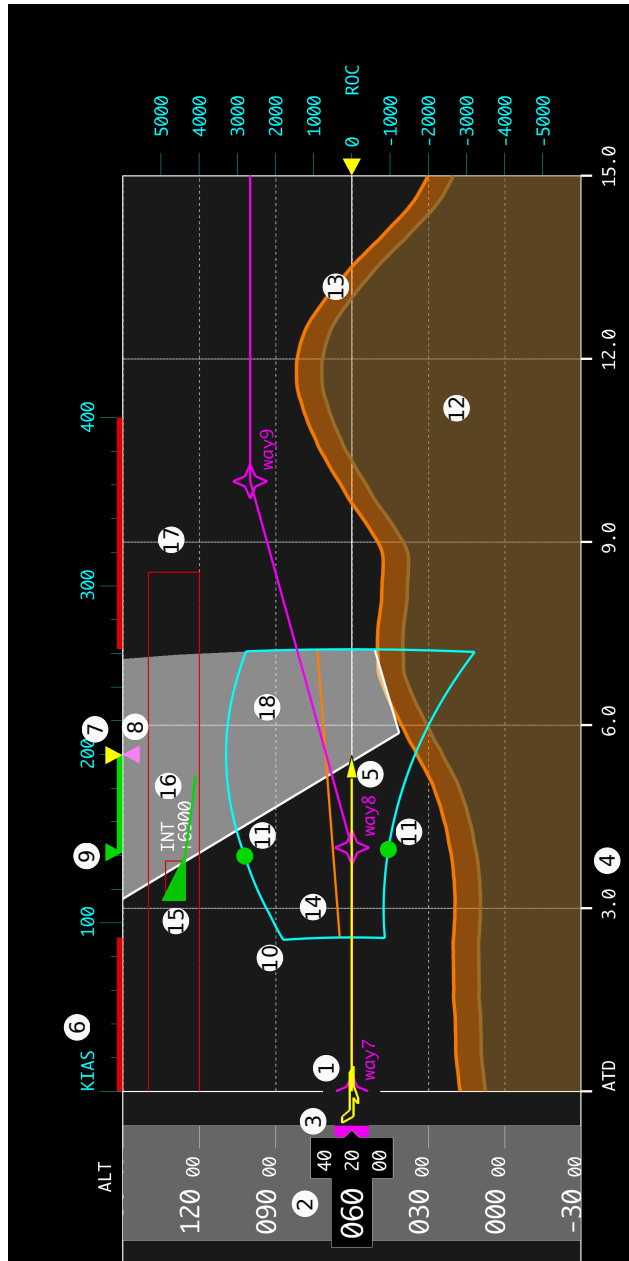


Figure 5.4: The baseline display.

airspeed (7). In horizontal flight, the yellow indicator will line up with the tip of the flight path vector. Since the display always shows true airspeed, and the conversion to indicated airspeed depends on altitude, the scale will (slowly) expand or contract with increasing or decreasing altitude, respectively. A magenta triangle indicates the required speed corresponding to the next waypoint in the flight plan (8) and a green triangle indicates the best angle of climb speed (9). A green line connects the current speed with the best angle of climb speed to visualize the remaining speed that can be traded in for altitude. Minimum and maximum speeds are indicated by red lines.

Ecological Overlay: EID Cues

The performance envelope of the aircraft is shown in cyan (10). This envelope shows the limits of (steady state, optimal) performance on the flight path vector and is treated as a causal constraint. The left and right side of this envelope represent the minimum and maximum velocities, respectively. The top part shows the maximum sustained climb performance. At full power, when the velocity vector is above this line, the speed will always decrease as long as it is above this line. The bottom part indicates the glide angles possible at idle power. At idle power, speeds above this line cannot be sustained. Two green dots indicate the optimal climb (full power, top) and optimal glide (idle power, bottom) speed/angle (11).

The terrain is drawn at the correct distance and height relative to the ownship, indicating its physical position (12). Since the flight path vector is drawn using true airspeed, the climb angle can be directly related to the terrain. If the vector points above the highest point, the aircraft will eventually clear the terrain. The terrain is drawn with two different colors. The brown color is used to indicate the physical terrain, the orange layer shows a 1,000 *ft* safety margin (13) which is an intentional constraint. As an additional cue, a terrain clearance angle is drawn in orange inside the performance envelope (14). This shows the angle necessary to clear the highest peak in the next 15 *NM* and includes the 1,000 *ft* safety margin.

Traffic is also drawn at the location relative to the ownship where it is located (15). It consists of a triangle with a stem that indicates the direction of travel of the intruding aircraft (16). A red box around the intruder aircraft indicates the protected zone (17), that should be avoided if possible. Since the intruder aircraft is moving, the relationship between the ownship flight path vector and the intruder is not as straightforward as it is for the terrain. To visualize this relationship, a conflict zone is drawn in the performance envelope (18). The conflict zone shows which combinations of ownship flight path vector angle and speed will result in a trajectory into the protected zone of the intruder assuming the intruder's trajectory does not change. Keeping the flight path vector out of this conflict zone will ensure a safe distance to the intruding aircraft.

5.2.3 The Intentional VSD

The augmented display, with the explicit distinction between causal and intentional constraints, as discussed in Section 5.1.2. is shown in Figure 5.5. The traffic constraint is now visualized by two areas. The intentional constraint (18) is shown as before. Inside this area a smaller causal constraint is drawn (19), which shows the area that will lead to a collision with an aircraft or lead to entering a severe wake turbulence area.

The terrain clearance angle is replaced by conflict zones to ensure visual conformance with the traffic constraint. Two terrain conflict zones are drawn: (20) a causal constraint, indicating which combinations of ownship flight path angle and velocity lead to a terrain collision, and (21) an intentional constraint representing the safety margin with a lookahead of 15 *NM*.

5.3 Experiment Method

Based on the experiment of Rijnveld et al. (2010), a new experiment was designed to investigate the effects of dealing with terrain and traffic constraints in combination with the two different risk boundaries.

5.3.1 Subjects

Sixteen subjects were divided in two groups. In the first group, there were eight experienced pilots with at least a Private Pilot License (PPL). Their average age was 38.2 years (SD 11.7) with an average flying experience of 3,412 hours (SD 2,965). All experienced pilots did also have a technical background, next to their aviation experience. In the second group there were eight subjects without any flying experience. Their average age was 24.4 years (SD 0.5) and they were all MSc graduate students in aerospace engineering with a basic notion of how an aircraft is controlled.

5.3.2 Apparatus

The experiment was conducted in a fixed base flight simulator. The experiment displays were shown on an 18 inch monitor located in front of the subject. A non-linear six degree of freedom Cessna 550 Citation II model flying in International Standard Atmosphere with no wind or turbulence was used for the experiment. The aircraft model was controlled by a right hand hydraulic side stick with a frozen roll input and a throttle quadrant on the left. The throttle directly controlled the throttle input of the aircraft model, the elevator input controlled the pitch attitude of the aircraft. A set of trim buttons on the stick controlled the elevator trim.




Figure 5.5: The IVSD display.

5.3.3 Scenario

The scenarios for this experiment used the same terrain database as Rijnveld et al. (2010). The four measurement scenarios were based on selected scenarios from that experiment, but were adjusted to better suit the intentions of the current one. The scenarios were chosen to represent situations where pilots have to navigate vertically through mountainous terrain while avoiding traffic.

Each scenario starts at a predetermined initial condition with the aircraft flying straight and towards the first waypoint. Until reaching the first waypoint, there is no traffic yet and the pilot has a chance to observe the terrain and flight plan. During this time, the pilot is also able to adjust the aircraft's speed and flight path. Each measurement scenario starts with a terrain obstacle that requires a climb, which can also be assessed by the pilot while on his way to waypoint one.

Upon passing waypoint one, an intruder appears and the actual traffic and terrain avoidance task starts. The setup of each scenario is described below. After clearing the terrain and traffic, the scenario ends at the next waypoint.

Scenario 1 In this scenario, the pilot has to climb 2,000 *ft* to clear the terrain with opposing traffic starting 8,000 *ft* above the ownship and descending. This results in the conflict zone geometry shown in Figure 5.6. This scenario leaves a very narrow gap between both conflict zones. Due to the geometry of this conflict, the boundaries of both conflict zones do not move much. The narrow gap remains for the majority of the time, while the intruder is still in front of the ownship. The causal conflict zone for the intruder, on the other hand, moves away steadily when the intruder approaches the ownship. This means that over time, a small excursion into the intentional traffic conflict zone will become less and less critical.

Scenario 2 In this scenario, the pilot has to climb 4,000 *ft* to clear the terrain with opposing traffic starting 2,000 *ft* above the ownship and climbing. This results in the conflict zone geometry shown in Figure 5.7. The climbing intruder leads to a much narrower conflict zone compared to Scenario 1. This leaves some room at the top of the performance envelope. When the ownship flight path vector stays below the intruder conflict zone, the top leg of this zone will move up and eventually there will be no more room left in the top part of the envelope. So, if the pilot immediately initiates a climb and puts the flight path above the conflict zone, it is possible to climb faster than the intruder and pass over it. When the pilot waits, however, it will be impossible to sustain a climb speed large enough to climb over the intruder.

Scenario 3 In this scenario, the pilot has to climb 4,500 *ft* to clear the terrain with opposing traffic starting 10,000 *ft* above the ownship flying straight and level. This results in the conflict zone geometry shown in Figure 5.8. The main difference with the other scenarios is that the intruder is not changing altitude.




Figure 5.6: Scenario 1




Figure 5.7: Scenario 2




Figure 5.8: Scenario 3




Figure 5.9: Scenario 4

This means that since the required altitude for the flight plan is well below the altitude at which the intruder is flying, the ownship will always level off well before getting close to the protected zone. As a result, in this scenario, the pilot can choose to completely ignore the traffic conflict and simply climb to the assigned altitude.

Scenario 4 In this scenario, the pilot has to climb 4,500 *ft* to clear the terrain, the traffic is flying in the same direction as the ownship, just ahead and 7,000 *ft* above the ownship but descending. This results in the conflict zone geometry shown in Figure 5.9. The intruder starts off flying slightly faster, which results in the large open space at the “slow side” of the performance envelope. A small decrease in speed would be sufficient to stay out of the intruder’s protected zone, but pilots might be tempted to slow down much more because of the large open space.

5.3.4 Independent Variables

A between-subjects design was chosen because the main focus of this experiment is on the strategy employed by the subjects. A within-subjects experiment would expose subjects to the same scenario more than once, possibly creating foresight that could influence their decision making.

The experiment used two between-subjects variables: flying experience and display configuration. For the flying experience, half of the population had no (non-simulated) flying experience [NOV]. The other half had at least a PPL [EXP]. Because of the focus on the strategic nature of the task, no further requirements were set for the subjects.

The display configuration was either a baseline display, which corresponded to the EVSD display of Rijnveld et al. (2010) [BASE], or the new display with the explicit visualization of causal and intentional constraints [IVSD]. An overview of the configuration of the experiment is shown in Table 5.2.

Table 5.2: Experiment Configuration

Total Population				
16 subjects				
Experience	NOV		EXP	
	8 subjects		8 subjects	
Display Type	BASE	IVSD	BASE	IVSD
	4 subjects	4 subjects	4 subjects	4 subjects

5.3.5 Procedure

The experiment started with a briefing of approximately half an hour, included in Appendix C, to introduce the concepts used in this display to the subjects. After this, a training phase was started to familiarize the subjects with the dynamics of the simulated aircraft and the various display features. Eight training runs were used to gradually introduce all concepts. This took close to an hour. Once these training runs were completed, and subjects were proficient with the display, a small break was held.

After the break, eight runs were performed of which four were measurement runs, and four were filler scenarios. The order of these runs was random, based on a Latin square matrix to compensate for possible learning effects based on the scenario order.

Subjects were instructed to clear the terrain and traffic, while following the flight plan at their own discretion, in a way *they* considered safe and comfortable. They were not given any minimum altitude or clearance instructions, but were asked to rely on their own best judgment. Once they were sufficiently clear of terrain and traffic the run was ended and subjects were asked how they experienced the scenario and for a short description of their strategy. The individual runs lasted between two and four minutes.

At the end of the experiment there was a debriefing where subjects provided feedback on the overall experiment and the specifics of the displays.

5.3.6 Dependent Measures

A qualitative analysis was used to investigate the strategies used by the subjects. This analysis will be based on the angular difference between the velocity vector and the constraint under investigation. For terrain constraints, this is straightforward since the velocity vector and the terrain angle have the same origin, the own aircraft, which means they can be compared directly. With the traffic constraints, however, there is no direct relationship. To calculate the angular difference for traffic constraints, the angle of the vector from the conflict zone tip to the velocity vector is calculated. This angle can be directly compared to the angles of the conflict zone legs. Figure 5.10 shows a schematic overview of how both angles are defined.

Three objective measures were used to quantify possible changes in pilot behavior: (i) The minimum terrain clearance (in [*ft*]), (ii) the distance between ownship and intruder at the Closest Point of Approach (CPA) (in [*m*]) and (iii) the number of protected zone violations.

Table 5.3 shows an overview of the four different kinds of possible protected zone violations that can occur where MSA stands for Minimum Safe Altitude, 1,000 *ft* in this experiment.




Figure 5.10: Definition of the angular deltas of the velocity vector with respect to the terrain conflict zone $[\delta_{terr}]$ and the traffic conflict zone $[\delta_{traf}]$.

Table 5.3: Possible protect zone violations

	Terrain	Traffic
Causal	Terrain collision	Traffic collision Severe wake turbulence
Intentional	MSA violation	Loss of separation

5.3.7 Hypotheses

It was hypothesized that:

1. Terrain clearance will increase with the enhanced visualization of the intentional constraints in the IVSD (Figure 5.4 versus Figure 5.5).
2. Protected zone violations will increase with the IVSD. It is expected that subjects will use more of the control space allowed by traffic constraints (relative to the experiment conducted by Rijnveld et al. (2010)) as causal and intentional constraints can be more clearly distinguished.
3. Experienced subjects are expected to use more conservative strategies relative to the novices without flying experience.

5.4 Results

5.4.1 Strategy Observations

In this section, a qualitative investigation of the experiment runs will be presented. As indicated in Section 5.3.6, the main variable of interest is the angular delta between the ownship velocity vector and the closest edge of the conflict zone under investigation.

The upcoming figures in this section show the angular delta with respect to the four constraints under investigation: the causal traffic constraint, the intentional traffic constraint, the causal terrain constraint, and the intentional terrain constraint (Table 5.3).

The horizontal axis of all graphs shows the Along Track Distance (ATD) between the ownship and the traffic. Using the along track distance means that we lose absolute time information, but it relates much better with the geometry of the situation, making it possible to compare different subjects. The start of each graph corresponds to the situation at waypoint 1 one where the intruder aircraft appears. The graph ends when the intruder has passed and the ATD is 5 *NM*.

The vertical axis shows the angular delta between the tip of the ownship velocity vector and the constraint. For the terrain constraint, shown in the bottom part of the plot, the delta is calculated by subtracting the terrain angle from the flight path vector. When the delta is positive, the velocity vector lies outside the conflict zone. When it is negative, the velocity vector lies inside the conflict zone. The solid line represents the delta with the intentional terrain conflict zone, the dotted line shows the delta with the causal terrain conflict zone.

The top part of each graph shows the angular delta with respect to the traffic constraint as shown in Figure 5.10. As with the traffic constraint, a positive delta indicates that the velocity vector lies outside the conflict zone, a negative value indicates the velocity vector lies inside the conflict zone. The solid lines show the delta with the intentional traffic constraint, the dotted lines show the delta with the causal traffic constraint. Since for the majority of the situations, the traffic conflict zones lie above the velocity vector, the vertical axis has been reversed. As a result, the traces above the dotted line indicate when the velocity vector is inside the conflict zone.

Scenario 1

In the first scenario, shown in Figure 5.6, the ownship is lower than the upcoming terrain and an intruder is descending toward the ownship. The combination of both conflict zones results in a very narrow control space that requires precise control inputs to stay clear of traffic and terrain conflicts. An overview of the strategies employed in this scenario is given in Figure 5.11 and Figure 5.12.

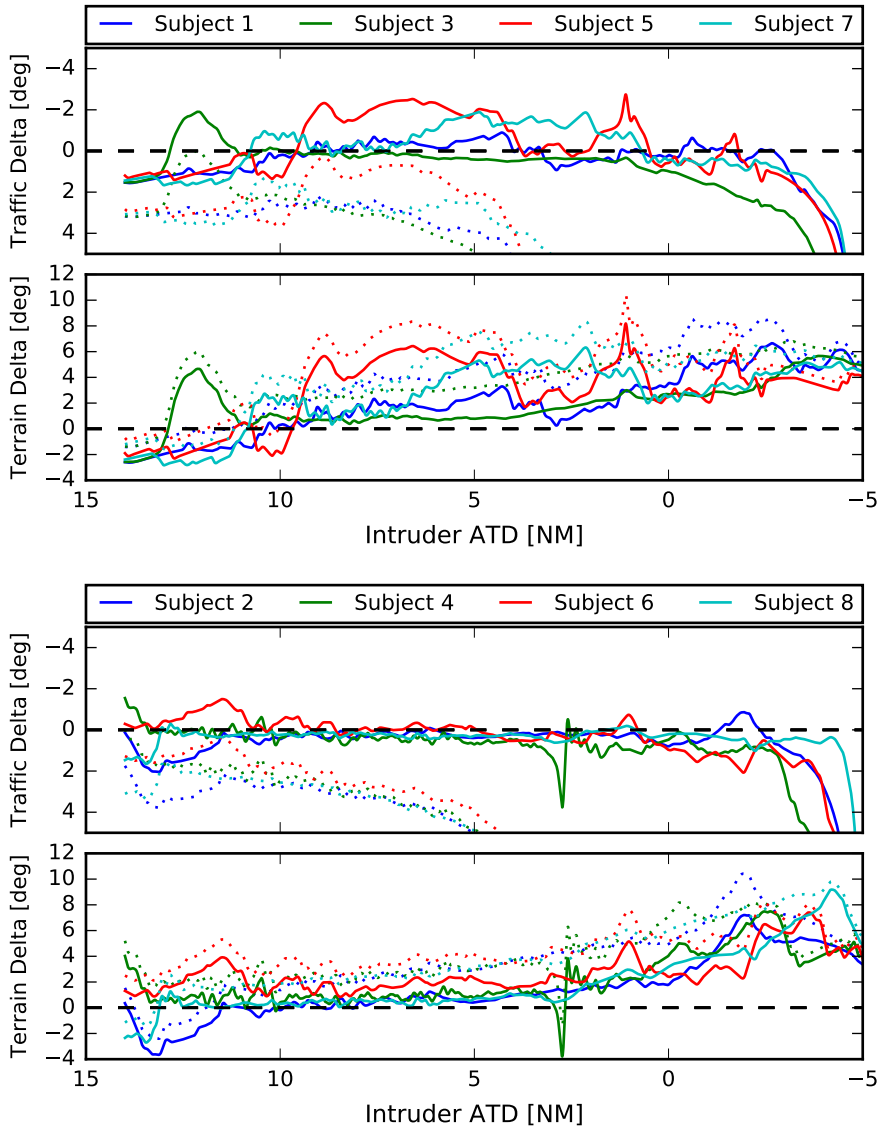


Figure 5.11: Scenario 1 Phase Plots, Experienced Subjects. Baseline display [top] IVSD Display [bottom]. Solid lines represent intentional delta, dotted lines represent causal delta.




Figure 5.12: Scenario 1 Phase Plots, Novice Subjects. Baseline display [top] IVSD Display [bottom]. Solid lines represent intentional delta, dotted lines represent causal delta.

Shown in the top part of Figure 5.11, in the group of experienced subjects [EXP] using the baseline display [BASE] all subjects, except one, follow the flight plan until reaching waypoint 2, indicated by the terrain delta below zero until approximately 10 *NM* ATD. Only Subject 3 initiates an immediate climb, initially putting the velocity vector into the intentional conflict zone. The dotted green line almost reaches zero temporarily, indicating the velocity vector was on the edge of the causal traffic conflict zone. It is important to note that with the baseline display used, the causal traffic conflict zone was not visible, so this was not a deliberate maneuver. After the initial excursion, Subject 3 stays out of the traffic conflict zone.

Subject 1 employs a different strategy. After pitching up, leaving the terrain conflict zone, he keeps his velocity vector on the boundary of the intentional traffic conflict zone. Subjects 5 and 7 also make active use of the intentional traffic conflict zone during their climb towards waypoint 3. Subject 5 even gets close to entering the (invisible) causal traffic conflict zone.

The bottom part of Figure 5.11 shows the results for the experienced subjects [EXP] using the IVSD display [IVSD]. Compared to the experienced subjects using the baseline display it can be noted that the spread in the traces is much smaller.

Initially, three out of four subjects had already moved their velocity vector out of the traffic conflict zone. Subject 6 waited until passing waypoint one, and then moved his velocity vector out. Subject two had his velocity vector outside of the conflict zone, but decided to move it back in to stay closer to the flight plan. After passing waypoint 2, all subjects have their velocity vector outside of the terrain conflict zone.

With the exception of Subject 6, all subjects keep their velocity vector on the boundary of the intentional traffic conflict zone. Subject 6 decided to climb faster, putting the velocity vector into the intentional traffic conflict zone, approaching but not entering the causal traffic conflict zone, indicated by the red dotted line.

Around 2.5 *NM* ATD it can be seen that Subject 4 was temporarily distracted and let the velocity vector drop back into the causal terrain conflict zone. He immediately recovered going back to the boundary of the intentional traffic conflict zone.

The novice subjects [NOV] using the baseline display [BASE] use a different strategy than their experienced counterparts. All subjects, except Subject 9, initiated a climb before reaching waypoint 1. This can be seen in The top part of Figure 5.12 where three traces start with a terrain delta around four degrees. As a result of this initial climb, the traffic delta shows that these three subjects have their velocity vector inside the traffic conflict zone when the intruder appears. Subjects 11 and 13 in the intentional part, just outside the causal part. Subject 15 is just inside the causal part, not shown in the baseline display.

The three subjects immediately respond to the intruder, Subject 13 and 15 immediately lower their flight path vector and put it on the boundary of the intentional traffic constraint. During this maneuver Subject 13 briefly drops the veloc-

ity vector into the intentional terrain conflict zone. Subject 11 responds differently, instead of lowering the velocity vector, he raises it into the causal traffic conflict zone (invisible to him) and keeps it there until leveling off at the assigned altitude around 6 *NM* ATD.

Subject 9, who started with his velocity vector below the terrain line, continues to fly level until reaching waypoint 2 where he raises the velocity vector and puts it on the boundary of the intentional traffic conflict zone.

Looking at the novice subjects [NOV] using the IVSD display [IVSD] in the bottom part of Figure 5.12, the same trend as with the experienced subjects is visible. The spread between the traces is much smaller. Two subjects start above the terrain conflict zone, two subjects are still following the flight plan when the intruder appears. Subject 14 immediately climbs, Subject 10 follows the flight plan until waypoint 2. After this initial response, all subjects keep their flight path on the boundary of the intentional traffic conflict zone. Around 4 *NM* ATD, when the traffic is close and starts to move overhead, Subject 10 moves the flight path vector up into the intentional traffic conflict zone to further increase altitude. At the end, he deliberately clips the protected zone, which is indicated by the solid blue line moving up.

Scenario 2

In the second scenario, shown in Figure 5.7, the ownship is again lower than the terrain. The intruder is climbing. This results in a narrow conflict zone for the intruder that quickly widens when the distance decreases. The open space above the traffic conflict zone allows subjects to climb over the traffic. An overview of the strategies employed in this scenario is given in Figure 5.13 and Figure 5.14.

In this scenario, with the group of experienced subjects [EXP], there is virtually no difference in strategy. All subjects put their velocity vector on the boundary of the intentional traffic conflict zone. Subject 4 chose to put the velocity vector above the traffic conflict zone, climbing over the intruder as a result. This can be seen by the large offset of the green terrain delta trace in the bottom part of Figure 5.13. The spike shown in the traffic delta trace of Subject 6 shows a brief excursion into the intentional protected zone. This was not a deliberate action, the subject assumed he already cleared the traffic and did not realize raising the velocity a small distance into the conflict zone would result in an immediate violation.

With the novice subjects [NOV], the same trend is visible. All subjects put their flight path vector on the boundary of the intentional traffic conflict zone. Three subjects—two with a baseline display, one with the IVSD display—climbed over the traffic. From the three terrain delta traces it can be seen that the three subjects climbing over the intruder already initiated a steep climb before reaching waypoint 1 and before the intruder appears. This steep climb already puts them above or in the top part of the traffic conflict zone.




Figure 5.13: Scenario 2 Phase Plots, Experienced Subjects. Baseline display [top] IVSD Display [bottom]. Solid lines represent intentional delta, dotted lines represent causal delta.




Figure 5.14: Scenario 2 Phase Plots, Novice Subjects. Baseline display [top] IVSD Display [bottom]. Solid lines represent intentional delta, dotted lines represent causal delta.

Scenario 3

In the third scenario, shown in Figure 5.8, the ownship needs to clear terrain ahead while the intruder is flying towards it at a higher but constant altitude. This results in the top part of the performance envelope being blocked. In this scenario, reducing the velocity and following the flight path is an option. But slightly deviating from the flight path and using the open space is also possible. An overview of the strategies employed in this scenario is given in Figure 5.15 and Figure 5.16.

At waypoint one, two of the experienced subjects using the baseline display have their velocity vector below the terrain angle. Their initial response is to raise the velocity vector and put it just below the boundary of the intentional traffic conflict zone. The other two subjects were already close to the intentional traffic conflict zone and also keep their velocity close to the boundary. Subject 5 in particular stays out of the intentional traffic conflict zone but decides to follow the flight plan because the traffic is flying much higher.

The experienced subjects using the IVSD used a similar strategy. The main difference is that all subjects already raised their velocity vector out of the causal and intentional terrain conflict zone. All subjects already had their velocity vector outside of the terrain conflict zone and put their velocity vector close to the boundary of the intentional traffic conflict zone. Subject 6 initially kept his velocity vector out of the intentional traffic conflict zone, but upon realizing the traffic was not descending decided to increase his climb rate by putting his velocity vector into the intentional traffic conflict zone. Subject 2 accidentally introduced a large pitch spike around 6 *NM* ATD.

Three out of four novice subjects using the baseline display have their velocity vector above the intentional terrain conflict zone, Subject 9 moves his velocity vector out as soon as the intruder appears. Two subjects put their velocity vector into the intentional traffic conflict zone, Subject 11 briefly enters the causal traffic conflict zone (not shown in this display).

The novice subjects using the IVSD also start with three subjects above the terrain angle and one subject below. After the intruder appears, all subjects move their velocity vector close to the intentional traffic conflict zone boundary. Only Subject 16 uses small excursions into the conflict zone to clear the traffic.

Scenario 4

In the fourth scenario, shown in Figure 5.9, the ownship needs to clear terrain ahead while a descending intruder comes from behind. This results in the front part of the performance envelope being blocked. An overview of the strategies employed in this scenario is given in Figure 5.17 and Figure 5.18.

The traces shown are slightly different from the ones presented for the previous scenarios. In the previous scenarios, the ATD started with a large value and continuously decreased during the run. In this scenario the ownship is trailing the intruder, and the ATD starts off small and slightly increases. Because of this,




Figure 5.15: Scenario 3 Phase Plots, Experienced Subjects. Baseline display [top] IVSD Display [bottom]. Solid lines represent intentional delta, dotted lines represent causal delta.




Figure 5.16: Scenario 3 Phase Plots, Novice Subjects. Baseline display [top] IVSD Display [bottom]. Solid lines represent intentional delta, dotted lines represent causal delta.




Figure 5.17: Scenario 4 Phase Plots, Experienced Subjects. Baseline display [top] IVSD Display [bottom]. Solid lines represent intentional delta, dotted lines represent causal delta.




Figure 5.18: Scenario 4 Phase Plots, Novice Subjects. Baseline display [top] IVSD Display [bottom]. Solid lines represent intentional delta, dotted lines represent causal delta.

the horizontal axes of the graphs are reversed with respect to the horizontal axes of the previous graphs. Furthermore, looking at Figure 5.9 it can be seen that the intentional portion of the conflict zone is large, and that the causal part is far away. As a result, the angular delta for the causal traffic constraints is much larger than 90 degrees. Including this in the graphs would make the vertical resolution of the graphs too small and therefore the causal angular delta is not shown for the traffic constraint.

The experienced subjects [EXP] using the baseline display [BASE] all follow the flight plan and stay below the terrain angle until reaching the second waypoint. Subject 1 and Subject 3 immediately reduce their velocity to remove the velocity vector from the intentional traffic constraint. Subjects 5 and 7 stay inside the intentional traffic conflict zone until reaching the second waypoint.

All experienced subjects [EXP] using the IVSD display [IVSD] immediately reduce their velocity to leave the intentional traffic conflict zone. Subjects 4 and 6 stay close to the border, Subjects 2 and 8 reduce their velocity significantly. Half of the subjects initially kept their velocity vector in the intentional terrain conflict zone.

All novice subjects [NOV] immediately move their velocity vector out of the intentional traffic conflict zone. The subjects using the baseline display [BASE] stay closer to the intentional terrain conflict zone than the subjects using the IVSD display. Subject 9 frequently drops his velocity vector into the causal terrain conflict zone in first part of the scenario. The subjects using the IVSD display stay further from the intentional terrain conflict zone. At the end of the scenario, Subject 16 enters the intentional traffic conflict zone again due to his velocity increase as a result of decreasing his flight path angle.

5.4.2 Terrain Clearance

The minimum terrain clearance for the four measurement scenarios is shown in Figure 5.19. The dashed line indicates the 1,000 *ft* intentional constraint shown in the display. For the novice subjects, Minimum Terrain Clearance (MTC) increases in each scenario with the IVSD display. For the experienced subjects, there is no such relationship. Note that the number of measurements is low, however, and it is difficult to draw conclusions.

With the first scenario, the differences are small. This can be attributed to the small open space between the terrain and traffic constraints. Only a few subjects deliberately put their velocity vector in the intentional traffic conflict zone, and even in those cases the excursions were small. As a result all subjects flew quite similar trajectories, which results in only minor variations in MTC for Scenario 1. The spread in MTC does decrease when comparing the data from subjects using the IVSD compared to the ones using the baseline display.

In the second scenario, the differences in MTC between both subject experience and display configuration are large. As explained in Section 5.3.3, subjects had




Figure 5.19: Minimum terrain clearance.

two distinct options to clear the terrain: stay low and have the traffic pass over them, or immediately start a climb and pass over the oncoming traffic. Subjects choosing the former option are not able to climb much higher than 2,500 ft above the terrain, while subjects opting for the climb will end up at least 5,000 ft above the terrain because they are forced to keep climbing because of traffic. The large spread in MTC for this scenario is a direct result of the different subjects choosing different options.

The MTC for Scenario 3 shows a clear difference between the BASE and IVSD display in both experience groups. In both cases the spread in MTC is much smaller with the IVSD. This difference, however, is not a result of the difference in display type. Because of the available space in this scenario, subjects were able to arrive at their designated altitude just before reaching the highest point in the terrain. In this scenario, subjects using the IVSD display were more precise in following this altitude. Subjects in the BASE display group deviated more from this altitude, resulting in a slightly larger spread.

The horizontal distance to the upcoming terrain was larger in the fourth scenario than in the other scenarios. This results in an MTC that is higher on average for the whole scenario, compared to the others. There is also a difference between

the BASE and IVSD display conditions both for NOV and EXP subjects. The distance to the terrain results in a terrain angle that is smaller than the other three scenarios. Every scenario starts out with the velocity vector in the terrain constraint, but for the fourth scenario, this intrusion is less pronounced. Examining the runs, it appears that subjects using the BASE display were more inclined to follow the flight plan initially, which required them to fly straight for some time. Most subjects using the IVSD display immediately initiated a shallow climb to keep their speed out of the terrain constraint.

5.4.3 Minimum Distance to Intruder

The minimum distance between ownship and intruder, or the distance at Closest Point of Approach (dCPA), for the four scenarios is shown in Figure 5.20. It follows the same trends as found for the terrain clearance.

For the first and third scenario, because of the geometry, when the terrain clearance increases the dCPA decreases. Traffic is always passing overhead, so increasing terrain clearance always decreases dCPA.

In the second scenario, this relationship still holds but is less straightforward since some subjects chose to climb over the traffic. The spread of the dCPA is less than the MTC.

Only in scenario four this relationship does not hold because of the geometry of this scenario. Here, the ownship starts out trailing the intruder with a higher velocity, thus overtaking. In response, all subjects reduced their speed to a value close to the speed of the intruder. Because of this, the distance between ownship and intruder remained fairly constant at the initial separation distance.

5.4.4 Protected Zone Violations

Since the altitude of the ownship at the beginning of each scenario is lower than the terrain ahead, without action from the pilot all scenarios would result in a violation of both the intentional and causal terrain constraint. Furthermore, all flight plans were defined in such a way that, when followed, they would lead to an intentional traffic constraint violation for Scenarios one, three, and four; in Scenario 2, an “impossible” flight plan was defined that would lead to an intentional and causal terrain violation. Besides this, because of the freedom pilots had to choose their strategy, in every scenario there was the possibility of violating any of the four constraints.

The protected zone violations are summarized in Table 5.4. Only two violations occurred, one in the group of novice subjects and one in the group of experienced subjects. Both of them were a violation of an intentional traffic constraint.

The first one occurred in Scenario 1 when Subject 10—a novice subject using the IVSD display—purposely clipped the intentional traffic protected zone at the end of the run, see Figure 5.12. The second violation happened in Scenario 2 when Subject 6—an experienced subject using the IVSD display—accidentally

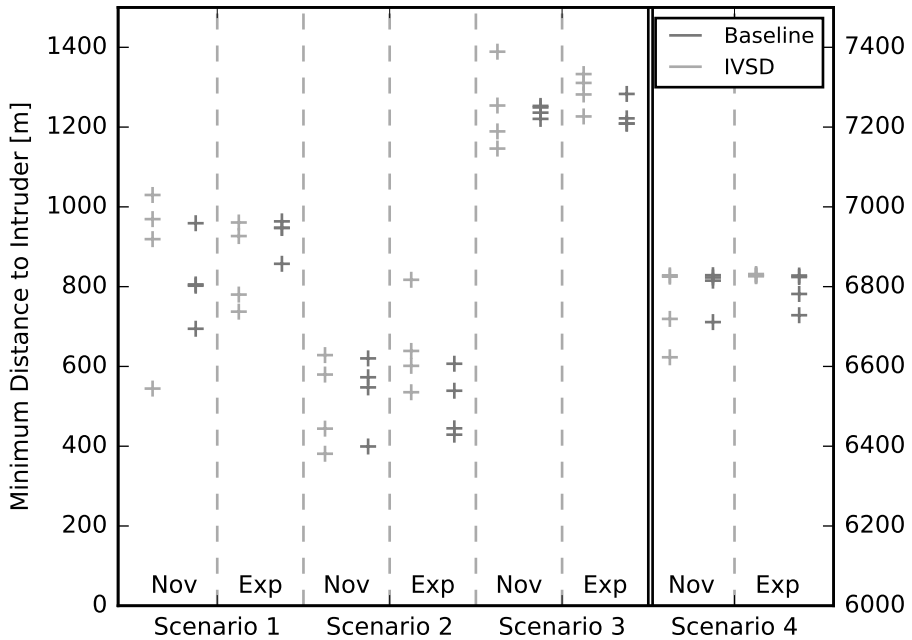


Figure 5.20: Minimum distance to Intruder, dCPA.

entered the intentional traffic protected zone after he mentally marked the traffic as ‘cleared’ and decided to climb further, not realizing at that point that the smallest excursion of the velocity vector into the intentional traffic conflict zone would lead to an immediate protected zone violation, see the bottom part of Figure 5.13.

Table 5.4: Protected Zone violations per subject type, based on Table 5.3.

	Novice				Expert			
	Terrain		Traffic		Terrain		Traffic	
	BASE	IVSD	BASE	IVSD	BASE	IVSD	BASE	IVSD
Causal	0	0	0	0	0	0	0	0
Intentional	0	0	0	1	0	0	0	1

5.4.5 Debriefing

During the debriefing, subjects were asked to comment on their overall impression of the display. Since display type was a between-subjects variable, they were not able to comment on the differences between the BASE and IVSD display.

All subjects indicated that they liked the concept of the display (both BASE and IVSD). The novice subjects described the displays as intuitive and easy to use. The expert subjects saw the added value in the conflict zone representation to resolve these types of conflicts.

A large number of subjects stated that one aspect they missed is timing information. With the representation used in the experiment, all velocities leading to a conflict were shown, even though it might take minutes or even hours for an actual protected zone violation to occur. A number of subjects clearly understood how time affected the situation, but could not find the proper cues in the display to support them in their understanding.

Looking at the novice subjects, the main feedback is that they approach the display in a game-like fashion where their main objective is to keep the velocity vector in the open space.

The experienced subjects were able to relate the display and the tasks to the real world situations they are expected to deal with. It was noted that both the lack of time information and the exaggerated vertical scale of the display make it look like the situations are more threatening than they are in reality.

Experienced subjects working with the baseline display also noted that there was an imbalance between the massive gray conflict zone and the orange line indicating the terrain angle.

A general sentiment among experienced subjects was that a better understanding of the conflict zone dynamics would greatly improve their ability to make better decisions. One of them indicated that he realized that if you understand the geometry, flying into a conflict zone is not necessarily a bad thing which indicated more training might be needed.

Finally, one of the experienced subjects noted that the experiment forced him to change airspeed a lot while in actual practice, pilots are reluctant to change velocity and prefer vertical or horizontal maneuvers to resolve conflicts.

5.5 Discussion

The goal of this experiment was to conduct a preliminary investigation to evaluate whether using an explicit split between causal and intentional constraints improves pilots' understanding of mixed terrain and traffic conflicts.

Overall, the total number of subjects available for this experiment is too low to formulate statistically significant conclusions. However, looking back at the hypotheses formulated in Section 5.3.7, a number of observations can be made.

Looking at the first hypothesis, terrain clearance will increase when using the IVSD, the clearance data presented in Figure 5.19 needs to be interpreted on a per scenario basis. In the first scenario there was a very narrow open space for the velocity vector. Most subjects kept their velocity within this space resulting in a narrow spread in MTC and minor differences between display and expertise groups. Only a few subjects used deliberate excursions into the traffic conflict zone, as we expected. The rest focused on always keeping the velocity vector in the open space.

For scenario two, the large clearance differences are the result of the fact that subjects had two options to resolve the conflict. Choosing the option to climb over the traffic automatically results in a larger clearance, which explains the large spread in the MTC for scenario two. If the group of subjects would have been larger, it could have been possible to make a better distinction between the subjects based on the solution they chose, but with the current sample size this would be meaningless.

The results from the third scenario can be misinterpreted as showing a significant effect of display type on minimum terrain clearance. But as explained in Section 5.4.2, this is a result of the geometry of the scenario. The spread is due to some subjects not flying at the waypoint altitude, but a little bit above or below.

The fourth scenario is different than the previous ones when considering MTC. In this case, from the strategy observations it became clear that the subjects using the baseline display favored following the flight path. This kept their velocity just inside the intentional terrain conflict zone. Which is indicated by a single orange line in the baseline display. For the subjects using the IVSD, the scenario also starts with their velocity vector inside the terrain conflict zone, but represented by a filled polygon. In this case, almost all subjects immediately pitched up to get their velocity vector out of the terrain conflict zone. For some subjects this was a deliberate strategy, but other subjects did not mention it as a strategy, which supports the assertion that the difference in constraint visualization has an effect on the experiment.

The second hypothesis was that protected zone violations would increase with respect to the experiment by Rijnveld et al. (2010). Comparing Table 5.1 with Table 5.4 it is clear that this does not hold. First of all, the results for the baseline display do not match the results by Rijnveld et al. (2010). Two major differences can be identified between both experiments.

First and foremost, we did not fully replicate the experiment, the scope was too broad for this research, so the focus was only on a number of specific terrain and traffic conflicts. From the paper (Rijnveld et al., 2010), it is not possible to determine which scenarios resulted in the intrusions shown in Table 5.1.

Next to this, subjects in Rijnveld et al. (2010)'s experiment were instructed to solve conflicts in a safe way, with *minimum deviations* from the intended flight plan. The rationale given was that the display is meant to sport decision making rather than being a command display. While sharing the same philosophy,

the current experiment only emphasized perceived safety, and not minimum deviations from the flight plan. This could explain why subjects in this experiment focused on keeping their velocity vector in the open space as much as possible, resulting in no terrain or traffic zone intrusions with the baseline display.

Looking at the differences between the two display types, there are only two traffic protected zone violations with the IVSD. These are the only two protected zone violations in the experiment. One violation was intended, namely in order to create more terrain clearance, the other one was accidental, based on a misjudgment when the conflict was almost over.

Referring back to the second hypothesis, it is clear that it does not hold for this experiment. Although unexpected, in hindsight it is not surprising. One of the main goals of this display design philosophy is to keep pilots away from dangerous situations. In this case that means keeping them out of the protected zones. Even with the baseline display this proved to be straightforward. The two intrusions are no exception to this, although the unintentional one can be related to a time information issue that will be discussed below.

Interpreting the results of the experiment with respect to the third hypothesis, experienced subjects apply more conservative strategies, is not straightforward. Looking at the observations and the debriefing it appears that pilots use more of their flying expertise when looking for solutions in the scenarios. They are generally more aware of their speed and altitude while most novice subjects use a game-like attitude where they mainly focused on keeping the velocity vector out of the conflict zones. This did not, however, result in different strategies employed between novice and experienced subjects.

The fact that there are no differing strategies does not need to be problematic. First of all, the limited number of subjects makes it hard to truly identify patterns. But more importantly, it does speak to the fact that the display provides an intuitive way of driving pilot behavior, for experts and novices. Even without in-depth knowledge, the display drives the strategies employed. Using the IVSD considerably reduced the spread in the traces recorded during the experiment for all subjects, suggesting that all subjects were more aware of the constraints and could better deal with them.

Looking back at the hypotheses, the experiment does not appear to show a large benefit of the IVSD over the baseline display. A number of reasons for this can be identified.

Since this experiment was an initial investigation into the advantages of using an explicit split between causal and intentional constraints, a limited number of subjects was used. More subjects would result in more statistical power, but especially in preliminary studies, using more than ten to twenty subjects is not feasible. Especially not when trained pilots are required, whose availability is typically limited.

Training of subjects also complicates experiments with Ecological Interface Design (EID) displays. While the basic features in EID displays are usually intuitive,

the finer details are usually complex and require a thorough understanding of the underlying concepts. Most of this understanding comes from experience and exposure, which is difficult to transfer in the limited amount of time available during an experiment. In this experiment, training focused on providing subjects with an overview of the display features that they could use to solve the conflicts. In essence they were shown *what* the display shows them, but not *how* they should use it. This was a deliberate choice to try to not bias subjects to employ a specific predefined strategy. The disadvantage of this approach, as seen in the experiment, is that subjects mainly fall back on using the most basic strategy of avoiding conflict zones altogether. Ideally, subjects would have had much more exposure to the concepts used in the display. They could, for example, be provided with a simple simulator of a generic non-experiment display that they can use to familiarize themselves with and explore the features. This would improve the subjects' understanding of the concepts, but is logistically not feasible and would also complicate the experiment since the amount of training then becomes a variable (and possibly a confounding factor) as well.

Finally, perhaps the most important reason for the lack of clear results can be tied to the decision to attempt to at least partially repeat the experiment by Rijnveld et al. (2010). During the experiment, it became clear that the scenarios chosen were perhaps not the most suitable scenarios. While the scenarios were probably suited to compare an EID display to a baseline display not using EID, better scenarios could be developed for the purpose of the current experiment. This would lose the relation with previous work, but could probably result in clearer results.

Decision Support The introduction discussed how current terrain and traffic warning systems are last resort warning systems for short-term collision avoidance. The VSD and IVSD are intended to provide collision avoidance support with a larger time horizon. Ideally, they provide long-term collision avoidance. Obstacles in the form of traffic and terrain, and the constraints they impose on the ownship can be seen well in advance. Pilots can avoid conflicts by periodically monitoring the situation and adjusting well before an obstacle becomes problematic.

Even if pilots would not be monitoring potential conflict situations, this experiment shows that both the VSD and the IVSD support mid-term collision avoidance. In all of the experiments, the subjects were within minutes of a potential collision with traffic or terrain. In these situations, all subjects had ample time to analyze the conflict geometry and devise a strategy to resolve potential protected zone violations or collisions.

Time-to-conflict A number of subjects indicated a difficulty in relating the urgency of a conflict to the conflict zones in the display. This is an important observation, and has been investigated before (Ellerbroek, Brantegem, van Paassen,

de Gelder, & Mulder, 2013; Mercado Velasco, Borst, Ellerbroek, van Paassen, & Mulder, 2015). For this experiment it was decided to use the same representation as Rijnveld et al. (2010), i.e., without time information. The fundamental issue is that, when looking at a conflict zone, all velocities inside the conflict zone will *eventually* lead to a conflict. The time it takes to reach this conflict, however, is not visualized when working with the default triangles used in this experiment.

The full explanation on how time is embedded in the conflict zone is outside the scope of this work (Ellerbroek, Brantegem, van Paassen, & Mulder, 2013), but it can be summarized as follows. The tip of the conflict zone represents the velocity of the intruder, while the vector from the tip to the ownship velocity vector represents the relative velocity of the ownship with respect to the intruder. When the ownship velocity vector is put on top of the tip, both intruder and ownship velocity vectors are equal and the resulting relative velocity is zero. As a result, the time taken to reach the intruder is infinite. Moving the ownship velocity vector away from the conflict zone tip—but keeping it inside the conflict zone—increases the relative velocity of the ownship, and thus decreases the time taken to reach the intruder or the protected zone around the intruder. A proper way to visualize this gradual decrease in time-to-conflict has not been demonstrated yet. But it is very well possible to “cut off” the part of the conflict zone that results in conflict times above a specific threshold.

Applying a conflict zone cut-off strategy could potentially change how subjects solve conflicts, the results would change, but new problems would probably arise. Cutting-off the conflict zones introduces a new dynamic element to the conflict zones. One advantage of the full triangle is that when the velocity vector is outside a conflict zone, it will always stay out of the conflict zone unless the intruder changes his own velocity vector. With the cut-off, this does not hold anymore. The cut-off point will gradually move towards the tip of the conflict zone.

Summarizing the above, adding time-to-conflict information to the display would be an important next step, but the representation needs to be chosen with care.

Training in EID Displays Related to the previous discussion there is also a challenge with the subjects understanding of the displays. This is not only important for this experiment, but is relevant to all EID related experiments. Despite aiming to yield intuitive displays for general use, EID displays require a significant amount of training for optimal use. Keeping the velocity vector out of the conflict zone is the straightforward feature that is easily picked up and used, but subjects new to the display have a difficult time estimating the intruder’s position and velocity even though this is also deducible from the conflict zone geometry.

This issue highlights the importance of training in this experiment, and in other experiments with EID displays. The limited amount of time available in an experiment is usually not sufficient to expose subjects to the display for a long time, making them intimately familiar with the details. Even if this time was

available, training remains non-trivial. Since a major focus of the current aviation research on EID focuses on conflict resolution there is a large possibility of training subjects for a specific task, biasing them towards the researcher's view on conflict resolution, while the purpose of EID is to leave as much freedom as possible to the subject to determine the best approach.

One way to tackle this issue could be to provide subjects up-front with a simulation of the display with some training situations and a way to explore all capabilities in self-made scenarios. This could improve subjects understanding without bias, but would also put a lot of responsibility and a large burden with the subjects.

A counter argument to the above was provided by one of the experienced subjects. Since the display is aimed at supporting conflict resolution in rare events, expecting real life pilots to remember training on a feature they almost never use would be optimistic. While this is a valid point for the real life implementation, in an experimental setting, a better understanding of the display would be preferred.

Uncertainties This experiment models an idealized world where all parameters are known perfectly. Outside an experimental setting, measurement and modeling uncertainties have a significant effect on avionics capabilities. As an example, the ownship velocity and performance constraints are represented by the flight path vector and the performance envelope in the display. In the experiment, they are represented by a well defined arrow and lines.

The performance envelope can not be directly measured. It is the result of a number of performance calculations based on measured flight parameters. As a result, there are two sources of inaccuracies. The first source is the measurement errors from the sensors. The second source are modeling imperfections. The performance calculations are typically based on simplifying assumptions, wind tunnel measurements, flight test data, etc. This also introduces small errors.

The inaccuracies obscure the actual parameters. Instead of knowing the exact value, the value lies somewhere in an interval with a statistical distribution. This gives rise to an intentional constraint, that needs to be visualized slightly different from the intentional constraints used in this experiment. Instead of representing a buffer around a causal constraint, the intentional constraint representing the uncertainty *replaces* the causal constraint since it is impossible to know its actual value.

This could, for example, result in a performance envelope that has thicker lines with a gradient representing the statistical distribution of the parameter. In this way the pilot will be aware of the uncertainties in the system and will be able to use this knowledge in a similar way as with the other intentional constraints.

5.6 Conclusion

The goal of this experiment was to conduct a preliminary investigation to evaluate if using an explicit split between causal and intentional constraints improves pilots' understanding of mixed terrain and traffic conflicts.

It was hypothesized that explicitly visualizing intentional constraints would increase terrain clearance. The experiment failed to show any significant difference due to the display. This can be attributed to a combination of scenario design, subject training, and amount of subjects.

Only two protected zone violations occurred. One was a deliberate violation and was part of that subject's resolution strategy. The other one was an accidental violation after the subject felt the situation had been resolved. This result indicates that the hypothesis that more protected zone violations would occur with the IVSD does not hold.

Although experienced subjects used more of their flying expertise, exhibiting more awareness of airspeed and altitude, the experiment did not show a clear difference in strategies employed by novice and experienced subjects.

Even though none of the hypotheses were confirmed, the experiment did show that using the IVSD resulted in a reduced spread in the traces recorded during the experiment runs. This indicates that subjects were more aware of the constraints and were able to fine-tune their strategy.

Discussion & Conclusions

This thesis addressed the role of rules, regulations, and procedures in the safety of cockpit displays using Ecological Interface Design (EID). This final chapter will reflect on the work presented in the previous chapters, present the conclusions, and provide some recommendations for future work.

6.1 Discussion

The first part of this discussion will focus on topics related to the research questions. The second part will address related topics that stood out during the research for and writing of this thesis.

6.1.1 On the Research Questions

Ever increasing automation has played a key role in the rapid development of the aviation industry in the last decades, and it will continue to play a key role in the development of future aviation systems. Safety and efficiency have been enhanced greatly, but in the mean time, new opportunities for mistakes and accidents are introduced (Parasuraman & Riley, 1997; Sarter, Woods, & Billings, 1997). The archetypical ‘ironies of automation’, such as designer errors and leaving tasks that can not be foreseen to the operator, are omnipresent in aviation (Bainbridge, 1983).

The use of EID can help human operators cope with these highly automated future aviation systems. Ecological interfaces aim to visualize the possibilities and constraints on a system. This makes them well suited to help the operator in handling complex tasks, but they have also been shown to lead operators to operate close to constraint boundaries (Rasmussen, 1997). In certain aviation contexts—like terrain or traffic avoidance—this can pose severe risks. Typically, rules, regu-

lations, and procedures are in place to prevent human controllers from operating close to risk boundaries.

Rather than attributing the limit-seeking behavior to the EID method itself, this thesis attempted to investigate whether the proper visualization of rules, regulations, and procedures can mitigate or prevent this behavior. This led to the following problem statement: *Can we, by clearly visualizing rules, regulations, and procedures, create ecological interfaces that lead to safer overall human-machine system performance?*

This problem statement led to three main research questions:

1. *How do rules, regulations, and procedures fit into the EID framework?*
2. *Will pilots be able to distinguish between physical constraints, and constraints introduced by procedures, rules, regulations, and procedures when they are visualized in the interface?*
3. *Will pilots make better decisions based on the additional information?*

Chapter 2 and Chapter 3 primarily dealt with the first research question and their outcomes will be discussed first in this chapter. Next, the two experiments described in Chapter 4 and Chapter 5 will be discussed as they attempt to answer the two remaining research questions. Finally, some general observations and lessons learned will be discussed.

Incorporating rules, regulations, and procedures To answer the first research question, Chapter 2 started by looking at the Work Domain Analysis (WDA) underlying EID displays. Rasmussen et al. (1994) already characterized work domains in terms on the relative degree of causal and intentional constraints.




Figure 6.1: A constraint-based view on the work domain, comparing the purely causal view (left) with a view including intentional constraints (right).

Figure 6.1 illustrates the conceptual difference between causal and intentional constraints. The left hand side shows a purely causal system, which is only governed by physical processes. The causal constraint boundary delimits the space of action possibilities, an abstract representation of all control strategies that an operator can employ. Physical viability obviously is the baseline for process control, because at a bare minimum, the operators should be aware of the physical limitations of the system, especially when violation of these limitations can have catastrophic consequences. In practical applications, physical viability is not sufficient. When economic factors are at play, for example, a cost effective and labor efficient operating procedure is also important. These types of constraints are called intentional constraints, and are shown on the right hand side of Figure 6.1. Intentional constraints form a subspace of the space of action possibilities. This subspace reflects the actors' intentions, rules, and practices (Hajdukiewicz et al., 1999). The main difference with causal constraints is that intentional constraints can be violated without directly or inevitably resulting in system failure.

To ensure safety in aviation systems, rules, regulations, and procedures are in place that constrain the physical space of action possibilities. For example, there is no physical reason why an airplane would not be able to fly over a city at 200 *ft* when obstacles are avoided, but rules are in place that prohibit this because it could be unsafe in case of an engine failure. The rule, however, is not set in stone. Under exceptional circumstances, pilots can deviate from the rule if they deem this necessary to ensure the safety of the flight. From a WDA perspective, these rules, regulations, and procedures are intentional constraints, they represent actors' intentions and can be violated under exceptional circumstances.

In most ecological interfaces for the aviation domain that have been developed so far, there has been no explicit emphasis on intentional constraints. Either because the scope of the WDA only accounted for the causal aspects of the system, or because they were implicitly included and not identified as such. Chapter 3 investigated two different ecological displays which, for some participants, showed limit-seeking behavior during their evaluation, in an attempt to identify whether intentional constraints and their visualization could be a factor in this limit-seeking behavior. This provided some interesting insights.

Based on the above two categories, displays can be distinguished in the context of intentional constraints: purely causal interfaces and interfaces that lump together causal and intentional constraints. The Ecological Synthetic Vision Display (ESVD) of Borst et al. (2010) and the Enhanced Vertical Situation Display (EVSD) of Rijnveld et al. (2010) each represent one category of displays. The ESVD only presents the physical terrain constraints combined with the climb power constraints of the aircraft itself. The EVSD presents a Protected Zone (PZ) around nearby aircraft as a 'no-go' area, without any reference to the fact that only a small part of this area represents a causal constraint. This was combined with a terrain constraint that was mainly represented as a causal constraint. Both categories of displays have their own issues.

Purely causal constraints do not necessarily fit well with practical aviation systems. They do provide clear guidance on the physical aspects of flying, but they do not map properly to the often highly procedural nature of modern day aviation. During most normal operations, pilots rarely run into the physical limitations of their aircraft. Among other things, rules, regulations, and procedures keep them away from these physical boundaries. Therefore, they should be part of the interface they use in their decision making.

The fact that the majority of the time the intentional constraints are the most important feature, might lead to a decision to exclusively visualize intentional constraints to avoid distractions. However, this would be misguided. While causal constraints might not be immediately relevant, they still provide important context to the intentional constraints. Furthermore, they are indispensable in unforeseen emergency situations. If the causal constraints are not visible by default, they would need to be introduced when a need for them is present. This again requires a priori knowledge by the display designers that is not always available. The inability to design for unknown, unexpected circumstances is one of the shortcomings of current day display design that EID attempts to overcome (Rasmussen & Vicente, 1989).

The second category of displays can be described as an incomplete implementation of causal and intentional constraints. A number of existing aviation EID displays that use the Protected Zone (PZ) concept—including the EVSD under study in this thesis—falls into this category. In contrast with the purely causal displays, these displays do include procedural aspects of traffic avoidance by not focusing on the exact physical location and conflict geometry of intruding aircraft, but instead they incorporate the safety buffers put in place by rules, regulations, and procedures. They have been shown to work well for medium- to long-term conflict avoidance, even in the presence of multiple intruders (Ellerbroek, Mulder, & van Paassen, 2011). They can, however, also result in situations where pilots virtually have no choice but to use the intentional space of one of the intruders to safely resolve a conflict. The problem in these situations is that typically, the PZ is presented as a uniform region where the exact physical location of the intruder aircraft is not clear. This can make it difficult for pilots to make clear trade-offs during conflict resolution.

The EVSD experiment by Rijnveld et al. (2010) did not introduce multiple intruders, and therefore did not suffer from the issue described in the previous paragraph. It did, however, combine the PZ intruder representation in the vertical plane with a terrain representation on a Vertical Situation Display. In this configuration, there was a mismatch between the representation of the traffic and terrain constraints. This is another pitfall when not properly distinguishing between causal and intentional constraints. When using a fully intentional boundary for one type of constraint, and only a causal boundary for the other type of constraint, there is a risk of over-emphasizing one constraint over others. This will inadvertently bias decision making towards the most emphasized one. Clearly

distinguishing between causal and intentional constraints early on in the design, and for EID displays this means in setting-up the Abstraction Hierarchy (AH), can help to avoid this and can assure proper balance between different constraints.

Distinguishing between causal and intentional constraints Based on the theoretical foundations in Chapter 2, Chapter 3 presented two attempts to improve the ESVD and the EVSD. Two experiments were designed based on these improved interfaces. One of the goals of these experiments was to investigate whether pilots can indeed clearly distinguish between causal and intentional constraints. From a visualization perspective this distinction was clear. The importance lies in whether pilots actually understand the distinction between the constraints, whether they understand the implications on their action space, and whether they act on them accordingly. In other words, will they make better decisions.

The Intentional Synthetic Vision Display (ISVD) experiment indicated that subjects are indeed able to clearly distinguish between causal constraints and intentional constraints. The intuitive nature of the visualization ensured that subjects clearly understood the implications of the causal terrain constraint and the intentional safety margin above the terrain. Different subjects used the intentional layer in different ways. Some used it to validate their predetermined strategy, others adapted their strategy based on the extra information. These differences in usage of the intentional constraint strongly indicate that subjects were clearly aware of the implications of the intentional constraints and were also able to incorporate this information to improve their decision making.

In the Intentional Vertical Situation Display (IVSD) experiment, the results were less clear. While the visual distinction between the causal and intentional constraints was clear, only a few subjects showed an understanding of their meaning and significance. Probably the most important reason for this is that subjects were unfamiliar with the underlying concepts of the display. While the training before the experiment was sufficient for the subjects to provide a basic grasp of conflict resolution with the IVSD, it was not sufficient to pass on the finer details required to be able to clearly reason about the situations they found themselves in. This mainly resulted in subjects employing a strategy where they simply kept their velocity vector out of the Conflict Zones. While this strategy was sufficient to solve the conflicts in the experiment, it did not trigger pilots to explore the additional possibilities offered by the visualization of the intentional constraints.

The ISVD experiment shows that pilots were able to benefit from a clear visualization of intentional constraints. The IVSD experiment, on the other hand, makes it difficult to claim a clear benefit of an explicit visualization of intentional constraints. However, this is more likely a result of insufficient training and scenario design than of the visualization itself. Therefore, the ISVD results indicate that explicitly visualizing intentional constraints allow pilots to distinguish between physical constraints and constraints originating from rules, regulations, and procedures, but more validation is still required.

Decision-making based on intentional constraints Besides being able to distinguish between causal and intentional constraints, the ultimate goal of visualizing them is to improve operator decision making. The previous section already indicated that with the ISVD pilots were able to clearly differentiate between causal and intentional constraints. They were also able to process their implications and either adjust their terrain avoidance strategy, improve their strategy or simply validate their predetermined strategy. This led to improved decision making, which in the context of the ISVD experiment meant an improved terrain clearance.




Figure 6.2: Minimum clearance per condition and repetition. The whiskers indicate the lowest and highest datum within the 1.5 IQR of the lower and upper quartile.

Figure 6.2 clearly shows the differences in behavior between the baseline display and the display with intentional constraints. Mainly in the ‘Hard’ scenarios, where there was little margin for errors, the minimum terrain clearance significantly increased. This indicates that the additional information provided by the intentional constraints helps pilots in their decision making during a terrain avoidance task.

As indicated in the previous section, subjects in the IVSD experiment were

not able to benefit from the addition of intentional constraints to the interface. Although there were some improvements in the performance of the subjects when using the IVSD, they do not indicate a strong benefit over the baseline display. Most likely, the reason for this is insufficient training and scenario design. None of the scenarios truly forced pilots into exploring strategies that made use of the ‘intentional space’ offered by intentional constraints. This in itself might not be an issue when subjects have an in-depth understanding of the display, and would use the intentional constraints for better performance or increased safety.

Optimal scenarios would put pilots into a situation that has a physical resolution, but not one within the rules, regulations, and procedures. In these scenarios, ‘breaking the rules’ would be the only viable option. By doing this, subjects would be forced into the constrained space which would more likely lead to a clearer difference between conditions the baseline condition and conditions using an explicit visualization of the intentional constraints.

The ISVD experiment indicated that when pilots are able to distinguish between causal and intentional constraints, their decision making improves, while the IVSD—due to training and scenario issues—failed to show a clear improvement. A single experiment is not sufficient to conclusively prove a hypothesis, but it does show that future validation experiments could be worthwhile.

6.1.2 On Related Topics

Naturally, during research new insights are gathered that are not necessarily directly related to the main research question but are still relevant and useful. This section attempts to address some of those insights.

Intuitive versus expert displays A common misconception about EID is that the resulting displays should be simple or natural. Borst et al. (2014) argue that EID is a methodology to design expert displays and that they should therefore not necessarily be simple or intuitive. Complex tasks typically require complex interfaces, and an over-simplification of the problem space is in fact quite contrary to the idea of EID.

The work in this thesis is certainly not at odds with this view, but it does highlight a nuance that applies to relatively short-term conflict resolution displays. When, for whatever reason, pilots have only a short time to respond to an event, they need quick, unambiguous support in their decision making. In conversations with pilots during the experiments in this thesis, they typically expressed favorable opinions with respect to the support they get from EID displays in terms of tactical and strategical planning. In quick response situations, however, they usually favor command-like interfaces that would simply instruct them to quickly and unambiguously perform a certain mitigating action.

At a first glance, this might seem to contradict the statement that EID interfaces are not supposed to be simple or intuitive. It might even be tempting to consider

a mode switching display that presents a rich EID interface under nominal conditions and switches to a command style interface under specific conditions. The obvious problem that arises with this is that the switching logic again depends on a priori knowledge of these specific conditions. This is at odds with the idea that it is impossible to predict all possible failures of a system, a fundamental underpinning of EID.

Looking at both experiments, however, there is already a certain level of intuitiveness present in the displays. Despite the richness of the information presented (especially in the IVSD experiment) pilots can always revert to a binary interpretation of the displayed elements. Simply moving the velocity vector into 'open' space will always resolve a conflict. Although this is not yet the same as a command style display, it does require limited mental processing or training and provides an intuitive way of resolving difficult situations.

When we want to leverage this feature, it is necessary to be able to define the specific subset of information required to support it. In essence, this results in a narrowing of the WDA scope. In a practical display, the WDA scope would be relatively broad in order to capture all relevant aspects of the system. This includes both physical aspects as well as safety, strategic, and efficiency aspects. When shifting from normal operations to short term conflict resolution, the scope only needs to include what is necessary for this resolution. Typically this will mainly include vehicle energy management and locomotion. Properly identifying this particular subset in the interface can help pilots to quickly, and intuitively, identify short-term resolution strategies.

In addition to a clear visualization of all relevant types of constraints, a command style overlay can always be added. When done properly, this approach can combine the best of both worlds. The EID features would provide a full overview of the situation, while the command style cues would highlight a solution that is deemed most appropriate for a given situation. In this way pilots can be quickly guided to a resolution strategy, while still having a complete picture to evaluate the appropriateness of the proposed solution for the actual situation at hand. For this approach to work, it would be important that both the representation on the interface and the command cues work based on a shared representation of the situation (Parasuraman, Sheridan, & Wickens, 2008). If this were not the case, it could be difficult or even impossible for pilots to evaluate the proposed solution because it might not correlate properly with the underlying representation.

In summary, complex systems do need complex interfaces, but under exceptional circumstances the interface complexity should often be reduced to relatively simple systems. When this simple system is considered in the analysis, it can be used to identify the prominent features that could support pilots in short-term conflict resolution. In order to further assist pilot's decision making, command style cues can be added to highlight potentially efficient resolutions.

Naturalistic experiments In this thesis, a choice was made to perform two experiments in a more naturalistic fashion than most experiments with EID cockpit displays that have been performed before. One of the reasons for this was the realization that under simulated conditions, pilots respond differently to deterministically introduced emergency events compared to truly unexpected events (Casner, Geven, & Williams, 2013).

This does not imply that previous experiments were wrong or inadequate. In previous experiments pilots were given a clear task and instructions on how to complete this task. Obviously, a certain amount of instruction is needed for any experiment. However, it is easy to bias subjects into a specific strategy. If, for example, subjects are asked to fly straight as much as possible, they will be less inclined to look for a solution that includes turns. Depending on the experiment this might be a deliberate choice by the researcher, or it might introduce bias unknowingly.

With task-oriented displays, this is typically not a problem since they are designed to support a specific task, and they require subjects to perform that specific task. In a landing guidance display, for example, instructions to follow the cues presented are obviously appropriate. This is different for most EID displays. EID displays attempt to visualize the systems' control space, and provides the operators with a certain amount of freedom to chose a control strategy suitable for the situation they face. With these types of interfaces, providing instructions that are too specific can introduce an additional set of constraints (which are not necessarily a part of the interface).

In the conflict resolution displays studied in this thesis, there is much emphasis on the ability of operators to come up with creative problem solving strategies supported by an EID interface. For this reason, a choice was made to design *naturalistic* experiments. The goal was to design experiments that reflect tasks that would have to be performed in actual flight operations. When unexpected situations (similar to the ones in the experiments) appear, there are no specific instructions for pilots beyond "staying clear of the terrain or the traffic".

The downside of choosing for this type of experiments is that it makes them much harder to analyze after the data have been collected, especially with a relatively low number of subjects. When subjects are asked to follow a specific strategy, there is usually quite some opportunity to find clear metrics to compare different subjects and different conditions with each other. When subjects have more freedom, in case of a naturalistic experiment, they will usually employ a number of different strategies that are hard to compare to each other. Two very specific problems arose in the two experiments in this thesis.

In the ISVD experiment, it was relatively easy to identify the general strategies employed by the subjects. This was mainly the result of a combination using proficient pilots and a realistic scenario. While they can employ different strategies (e.g., immediately climbing vs. starting with a turn), they all performed the basic flight maneuvers required to execute their strategy in a consistent manner. The difficulty in this experiment is mainly the low number of subjects. While general

observations about which strategies are used ‘more or less’ when changing experiment conditions can be made, there were not enough subjects to detect any significant trends.

The Intentional Vertical Situation Display (IVSD) experiment suffered from a different problem. The experiment scenario, avoiding terrain and traffic in the vertical plane, was more abstract and further removed from actual practice. This resulted in a much greater variability in both strategic and tactical control. Most subjects focused on keeping their flight path vector out of the conflict zones, but the way in which they accomplished this varied significantly between subjects. Combined with the low number of participating subjects, this again lead to results that were difficult to analyze and generalize to a higher level.

These shortcomings do not invalidate the use of naturalistic experiments, but they do show that they require a lot of attention to detail in the design phase. For the ISVD this would have meant using a larger number of subjects. The IVSD experiment would probably have benefited from a different approach. With the current implementation, there was a significant emphasis on repeating the previous EVSD experiment (Rijneveld et al., 2010) as much as possible in order to have a reference to compare it to. In hindsight, this was probably not the right choice. A reboot of this experiment would benefit from improved conflict scenarios that would allow for multiple resolution strategies. This could provide crisper results which would make the analysis more comprehensive. Furthermore, better metrics matching naturalistic behavior should be defined to improve the analysis. Finally, for this experiment, too little subjects made use of the possibility of using ‘intentional’ conflict zones to resolve a conflict. This is partly due to the conflict geometry, but also due to the lacking familiarity of subjects with the display. This aspect will be discussed in the next subsection.

Training aspects of EID experiments As discussed before, EID displays are designed to be displays for experts. They require a significant amount of knowledge of the system under control by the operator. As a consequence, they also require a certain amount of training to become proficient with them.

When performing experiments with EID displays, training subjects for their tasks can quickly become difficult. Next to understanding the many visual cues on the display, subjects also need to be or become proficient in the underlying fundamental principles behind the information. In some cases, when showing an optimal climb angle for example, this is straightforward since the display elements refer to concepts already familiar to the subjects. In other cases, the conflict zones for example, most subjects would be working with unfamiliar concepts that take a significant amount of time to fully understand.

In the second case it is easy to see that significant training is required. But also in the first case, additional training is necessary. Even if the basic concepts would be known, merely visualizing them means that subjects will need to learn how it

interacts with other display features. So in any case, usually a lot of training is necessary in order to get meaningful results from an experiment.

Sufficient training in the context of EID experiments would imply that subjects receive a detailed briefing on the display concepts and ample hands-on experience with the display, where they can explore the full functionality and familiarize themselves with all the features.

Due to time and resource constraints, a typical experiment with a subject takes approximately four hours, of which two or three can be dedicated to training. For the ISVD experiment this was sufficient (although more training would still have been better) since the display addition was relatively intuitive and the task was straightforward. For the IVSD experiment, however, training was probably inadequate. Most subjects were unfamiliar with the concepts used in the display, and were unable to gather enough knowledge from the training to truly use the new features to their full advantage.

There is no straightforward solution for this problem. Limited resources will always be a constraint to performing experiments with human operators. However, putting more emphasis on the quality of the briefing would be a first step. Adding videos or interactive training aides, next to a clear description of the concepts could improve understanding. Even occasional tests or quizzes could be added to get a feel for the knowledge level of the subjects could be beneficial. Furthermore, finding subjects with an inherent motivation to participate in the experiment can also help. Even though all subjects participate on a voluntary basis, there is always a difference in eagerness to really understand the experiment, which could correlate for instance with how well subjects study the briefing beforehand.

Finally, a different approach to EID experiments would be to use a dedicated pool of test subjects that are thoroughly trained on both the concepts and the display features. By regularly participating in experiments or training sessions, they could maintain a basic level of 'EID-proficiency' which would facilitate training for specific new experiments. The downside, of course, is again the resources required to create and maintain such a pool of qualified subjects.

The human element in aviation safety The role of human operators in aviation safety typically comes to light in the aftermath of serious incidents or accidents. Mechanical failures have become rare, so it is not surprising that the human element is almost always a significant factor in analyzing mishaps. This is not surprising, pilots are ultimately responsible for the safety of a flight. A simplistic view on human error, combined with hindsight bias too often leads to the view that humans are a hazard to safety (Roesel & Vohs, 2012). This can, and has, lead people to believe that introducing more automation and further reducing the role of the human operator is the right course of action to improve aviation safety.

This view, however, glances over the importance of the human operator in this highly complex domain. Almost two decades after writing his seminal book on

human error (Reason, 1990), James Reason went on and wrote another book highlighting how human heroic acts have the potential to bring troubled systems back from the brink of disaster (Reason, 2008). The realization that humans actively contribute to the safety of complex systems is an important underpinning of the work in this thesis.

A set of rules, regulations, and procedures are in place to maintain a high levels of safety in aviation systems. Pilots are expected to operate within these rules, regulations, and procedures as much as possible, except when faced with off-normal conditions where the pilot needs to use his own experience and judgment to resolve the situation. The amount of creativity and ingenuity required to deal with rare, unique, unanticipated events is one of the most important reasons to keep pilots in the cockpit of a commercial airliner.

There is, however, a catch. The decision on when to ignore rules, regulations, and procedures is non-trivial. This was highlighted by the Royal Air Maroc Go-Around incident presented at the beginning of Chapter 2. In this incident, pilots erroneously felt they had to deviate from the Go-Around procedure and with that decision put the aircraft and its passengers at significant risk.

Being able to support pilots' decision making and leveraging pilots' creativity and problem solving skills in these unanticipated situations will greatly improve the overall safety of aviation systems. This means that rather than increasing the level of automation, automation is needed that cooperates more closely with the human operator (van Paassen et al., 2013; Christoffersen & Woods, 2002).

Beyond safety applications The focus in this thesis has been on intentional constraints related to safety. Obviously, intentional constraints are not limited to representing safety related to rules, regulations, and procedures. Intentional constraints can also be used among other things to shape behavior for performance, environmental, and efficiency reasons. As an example, consider noise abatement procedures. Since noise pollution has a high impact on the quality of living, especially in close proximity of airports, it is important that under normal circumstances pilots adhere strictly to these procedures in order to minimize noise pollution. From an interface point of view, these procedures are intentional constraints and they can be represented as such on an interface. By explicitly marking them as intentional constraints, they can be visualized in conjunction with similar, but causal constraints such as towers, tall buildings, and large cranes. In this way, when an emergency would occur, pilots will be able to quickly assess the risk associated with violating these noise abatement related intentional constraints and can avoid obstacles while potentially deciding to temporarily ignore the noise concerns.

6.2 Conclusions

This thesis showed how rules, regulations, and procedures can be fitted into the EID by properly identifying them and representing them in the AH. This thesis proposes a distinct split in the AH between the causal domain and the intentional domain. In this way, there is a clear distinction between the rigid causal constraints and the more flexible intentional constraints. This in turn allows for the design of an interface that presents the same distinction which can help pilots in prioritizing constraints, especially in unexpected circumstances.

The ISVD experiments shows that pilots were able to benefit from a clear visualization of intentional constraints. The IVSD experiment, on the other hand, makes it difficult to claim a clear benefit of an explicit visualization of intentional constraints. However, this is more likely a result of insufficient training and scenario design than of the visualization itself. Therefore, while the ISVD results indicate that explicitly visualizing intentional constraints allows pilots to distinguish between physical constraints and constraints originating from rules, regulations, and procedures, more validation is still required.

The ISVD experiment indicated that when pilots are able to distinguish between causal and intentional constraints, their decision making improves, while the IVSD—due to training and scenario issues—failed to show a clear improvement. A single experiment is not sufficient to conclusively confirm a hypothesis, but it does show that future validation experiments could be worthwhile.

A number of potential causes were identified as to why the IVSD experiment failed to show a clear difference between the baseline condition and the addition of explicit intentional constraints. First and foremost, subjects lacked familiarity with the basic principles of the interface due to insufficient training. Next to this, the scenarios failed to put subjects into the kinds of situations that would truly benefit from the improved visualization. Finally, a greater number of subjects would be required.

6.3 Recommendations

This thesis will be concluded by some recommendations for future work.

Subject training Future EID experiments in general, and naturalistic experiments in particular can benefit from an improved attention to subject training before experiments. As shown in the IVSD experiment, the non-trivial information-rich EID interfaces require a thorough understanding of participants of the underlying concepts for them to be able to make effective use of the display.

This requires a more elaborate briefing that would not only focuses on the experiment-specific features, but also puts strong emphasis on the foundation of

the display. The hands-on training session should aim to gradually build familiarity and understanding with all features of the display. In designing training scenarios, it is important to ensure that no implicit bias in control strategy is induced by the training.

Realistic scenarios The discussion on naturalistic experiments did already address the reasons for favoring experiments that aim to replicate realistic scenarios, where subjects have a large operational freedom in choosing the appropriate resolution strategy. Future experiments can certainly benefit from this approach. The ISVD experiment showed how subjects were able to relate to the scenario, which resulted in decision making that could better compare to real-life decision making in similar situations. They do, however, require well designed scenarios that clearly map to realistic situations.

Improved visualization Since this thesis presented an initial exploration of the inclusion of intentional constraints in EID-based displays, there was limited emphasis on the design aspect of the actual display features. The representations were suitable for experiments regarding the fundamental implications of intentional constraints, but they lack the refinement required for actual operational settings.

The simple terrain overlay used for the ISVD can be further improved to better support pilots when they fly below the Minimum Safe Altitude (MSA). When this happens, it is not obvious how low the aircraft actually is. As a result, pilots can not easily differentiate between a minor dip below MSA and a significant incursion that requires an immediate response to avoid a terrain collision.

With the current display, pilots only know that they are below MSA but have no feedback about how severe this violation is or what the most efficient escape maneuver would be.

Experience with the IVSD highlights a different problem. In this display, multiple types of constraints are visualized and multiple actors can be present at the same time. In these types of displays clutter quickly becomes an issue. A careful design of the display features is required to maintain display clarity without losing the advantage of visualizing intentional constraints.

Improved Metrics for Naturalistic Experiments The operational freedom subjects have in naturalistic experiments can result in a larger variation in the strategies employed by subjects. This variability makes it harder to use typical objective measurements used in traditional experiments. To define metrics for naturalistic experiments, a careful review of the experiment setting and its corresponding expected naturalistic behavior is required.

ISVD Briefing

A.1 Background

Ecological Interface Design (EID) has been successfully used in previous research to develop new displays that support pilots and air traffic controllers. One of the strengths of this design methodology is that it provides operators with the boundaries and limitations of the system. Based on this information, operators can reason about their preferred approach to solve a problem or complete a task. As an example, instead of only providing a 'climb' instruction, an ecological display can show the climb performance with respect to the terrain. Based on this information, the pilot can decide which evasive maneuver fits best in his situation. Past research has shown that EID displays can enhance terrain awareness and lead to less terrain incidents.

One limitation of the current EID displays is that they are mainly based on physical constraints. This is sufficient to avoid overstepping physical boundaries, but is not sufficient in a complex domain like aviation where rules, regulations and procedures are necessary to achieve an acceptable level of safety. This experiment is part of a research project attempting to incorporate procedural constraints in EID displays.

Procedural constraints differ from physical constraints in the fact that they are soft constraints. They indicate desired behavior, but can be violated if the pilot decides it is necessary. In the context of a terrain awareness display, like the one used in this experiment, this would mean both indicating the physical edge of the terrain and an indication of a safe margin above the terrain.

A.2 Interface

A.2.1 Basic Instruments

The basic Synthetic Vision Display (SVD) is shown in Figure A.1. The display design is based on the Garmin G1000 found in Cessna Nav III equipped aircraft operating in revisionary mode. All indicators work like their real life counterparts with a few exceptions.

- There is no side slip indicator. The rudder controls in the lab are out of service, preventing the pilot from directly controlling side slip.
- The airspeed indicator does not show speed ranges and speed trend. Only the maximum angle of climb speed (V_x) is shown as a reference speed next to the airspeed tape.
- The comm and nav radio controls, together with the button bar at the bottom of the screen have no function and are only included for visual conformance.

Next to the traditional instruments, there is also a Flight Path Vector (FPV) indication. This shows the current flight path angle (γ) and the current track angle (χ). In the absence of wind and with the side slip controller used in the experiment, the track angle will be equal to the heading most of the time and the FPV will be centered horizontally. The flight path angle indicates the actual climb angle, basically the pitch angle minus the angle of attack. When combined with a synthetic vision system, the FPV shows where the aircraft will end up in relation to the surrounding terrain. If the FPV is above the terrain, the aircraft will climb over it. If the FPV is somewhere on the terrain, the aircraft will eventually fly into the terrain at the spot indicated by the FPV.

A.2.2 EID Additions

The final addition to the basic display is the climb performance indicator. This indicator shows two horizontal green brackets. The outer brackets show the maximum sustained climb angle at full power. The IAS required to achieve this climb angle is V_x and is indicated on the airspeed tape. It is of course possible to climb steeper than the indicated maximum angle, but this will eventually bleed airspeed until the stall speed is reached or the nose is lowered.

The inner bracket is similar to the outer bracket, but it shows the maximum sustained climb angle achievable with the current power setting. At full power, both brackets will indicate the same climb angle. It is important to note that both climb angles are calculated based on the current altitude. For short climbs, this will not have a big influence, but for longer climbs, climb performance will deteriorate significantly with increasing altitude.

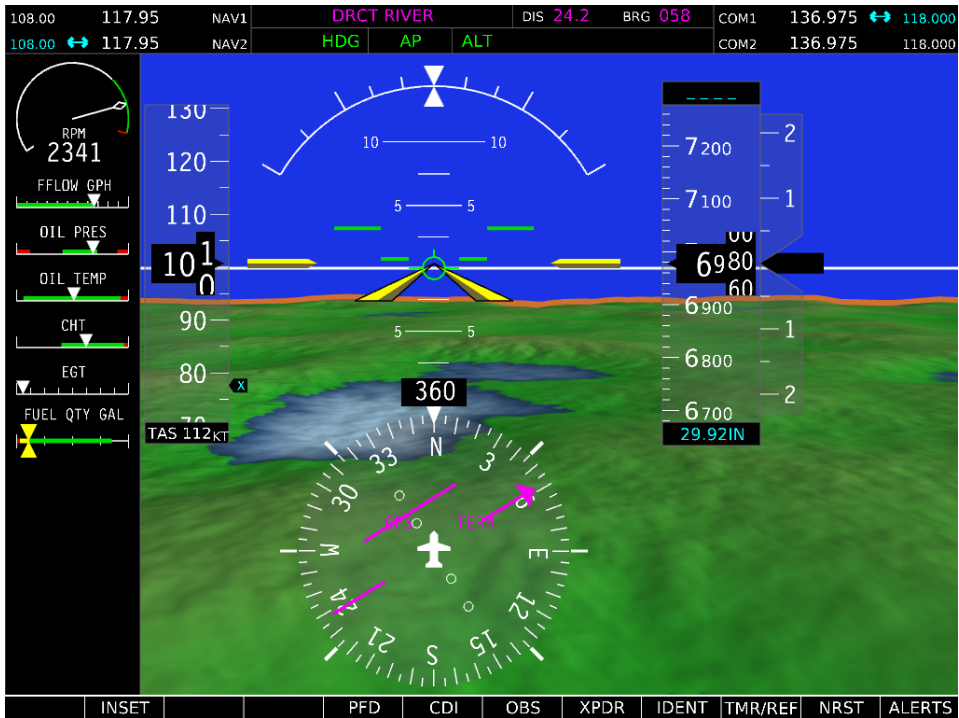


Figure A.1: The basic Synthetic Vision Display

A.2.3 Procedural Additions

The procedural constraint modeled in this display is the minimum altitude required above the terrain. It is presented as an orange layer over the terrain. This layer provides immediate feedback on the required flight path angle to fly over the terrain with the expected clearance. Keeping the flight path angle in the orange region will still provide terrain clearance, but will violate the minimum altitude requirement.

A.3 Aircraft Model

The aircraft model used in the simulation is a non-linear Cessna 172 model. Each run the model starts in a trimmed state with the autopilot maintaining heading and altitude. After the autopilot disconnects the pilot can control pitch, roll and engine power. The rudder pedals are inoperative, side slip and roll moments from engine torque are controlled by a roll and yaw controller. During the run, the pitch trim can be used to decrease pitch control forces.

A.4 Experiment

A.4.1 Goal

The goal of this experiment is to evaluate the effect of adding procedural constraints to a display on pilot's decision making and flight safety.

A.4.2 Design

The experiment is divided in two blocks, one with the procedural constraints and one without them. Within these blocks, a number of runs will be performed with differing difficulty and performance settings.

A.4.3 Scenario

All experiment runs will take place in a fictitious fjord like terrain. You are low on fuel and by mistake entered the wrong valley. Instead of leading to the airport, this valley is a dead end. To make things even worse, outside visibility has dropped to near zero. Your fuel will not allow you to turn around and take the original route, so your best bet is to climb over the fjord towards the waypoint selected in your HSI.

A.4.4 Task

During each experiment run, you will have to complete five main tasks:

1. Set the throttle to the trim position indicated on the screen
2. Observe situation until the autopilot disconnects
3. Silence the autopilot disconnect horn
4. Navigate towards the waypoint in a way that you perceive as safe, no specific strategy is required, use your own judgment
5. Provide feedback about strategy and performance

A.4.5 Apparatus

The experiment will be conducted in the Human-Machine Interface Laboratory located at the Faculty of Aerospace Engineering of the Delft University of Technology. The pilot will be seated in the right seat of the fixed-base simulator. A 18 inch monitor in front of the pilot will show the experiment display. The hydraulic side stick located on the right of the seat is used to control pitch and roll. The throttle on the left controls engine power.

On the throttle, there are two switches that will be used during the experiment. The first one is the autopilot disconnect switch, number 6 in Figure A.2. This button needs to be pressed to silence the autopilot disconnect horn. The second switch is the china hat switch, number 4 in Figure A.2, used to control the pitch trim. Pushing the switch forward will trim the aircraft nose down, pulling the switch backward trims the aircraft nose up.

The mode control panel located on top of the main instrument panel can be used to control the selected course on the HSI. Rotating the course control knob, shown in Figure A.3, will adjust the selected course on the MCP and the HSI.




Figure A.2: Throttle button layout



Figure A.3: MCP

Four projectors will project an outside visual on the walls surrounding the simulator. Only some flares of clouds will be shown to provide some sensation of motion, but no task related information will be present on the outside visual.

A.4.6 Procedure

The experiment will start with a *briefing and training* phase during which the pilot can get familiar with the experiment and the control and dynamics of the simulation. Once the pilot feels comfortable with the display and control of the aircraft, the *measurement* phase starts. In this phase, the different measurement runs will be performed. The runs are split in two blocks of 8 runs, with each run taking between 3 and 5 minutes each. An estimated timing is shown in Table A.1.

30 min	Briefing
60 min	Training
15 min	Break
60 min	Measurement Block 1 (8 runs)
15 min	Break
60 min	Measurement Block 2 (8 runs)

Table A.1: Experiment timing

Appendix B

ISVD Questionnaire

General

Name:

Age:

Experience:

Experiment Questions

Does the procedural addition make the task clearer?

Do the procedural additions change your strategy?

Does the level of safety you perceive change with the procedural additions?

Are the procedural additions useful?

Other remarks

IVSD Briefing

C.1 Background

This experiment is part of an on going research topic in constraint based display design. Constraint based displays try to better assist pilots in understanding complex situations by showing opportunities to solve a problem instead of providing one predetermined solution. This can not only help increase situational awareness for pilots, but also allows pilots to take into account circumstances that were not known or anticipated during the design of the display. The goal of this experiment is to evaluate a new constraint based display.

C.2 Interface

C.2.1 Basic Instruments

Two displays are used in this experiment. The first one, a standard PFD is not part of the experiment and is only used to support in the flying task. It is based on a standard Garmin G1000 display, without the bottom part and using modified engine instruments.

At the center of the PFD there is a classical attitude indicator ① combined with a flight path vector indicator ②. The bottom part shows a heading indicator ③ which will not be used since the heading is fixed in this experiment. On the right side an altitude tape ④ is shown with the required altitude ⑤ from the flight plan shown on top and a caret ⑥ indicating this altitude reference on the tape.

The speed tape ⑦, indicating calibrated airspeed is shown on the left of the attitude indicator. The true airspeed ⑧ is shown in the box below the tape, the required airspeed ⑨ from the flight plan is shown above the tape with a caret ⑩ indicating this airspeed on the tape. A green band indicates the normal speed




Figure C.1: The Primary Flight Display

range with two red bands showing the minimum and maximum airspeed. Finally, the speed resulting in the best climb angle V_x (11) is shown with a blue flag with an x.

The engine instruments are shown on the far left of the display. The top part shows fan RPM N1 in both bar (12) form and digital readout (13). The bottom part shows turbine RPM N2 in digital form (14).

Between the different runs, a throttle position indicator (15) will be visible in the display to aide in putting the throttle in the trim position. The field will indicate either up, down, or ok. Before starting the next run, the throttle should be moved up or down according to the indication until ok is indicated.

The second display is the actual experiment display. This display is a vertical situation display, or VSD. It shows a schematic side view of the aircraft together with the flight plan and terrain and traffic obstacles in the vertical plane. There is quite a bit of information in this display.

The display is centered around the ownship aircraft symbol shown in yellow (1). This symbol has a fixed position and will never move. The aircraft's altitude is indicated with an altitude tape on the left side (2) in combination with the horizontal grid lines fixed altitude increments. When the altitude increases both the altitude tape and horizontal grid lines will move down. The altitude associated with the active waypoint is indicated by a magenta bracket (3).

The horizontal axis shows the remaining along track distance (4) at which obstacles and waypoints are currently located. It is important to note that the horizontal and vertical scale are different. With an aspect ratio of approximately five, vertical distances are shown to be about five times larger than they actually are.

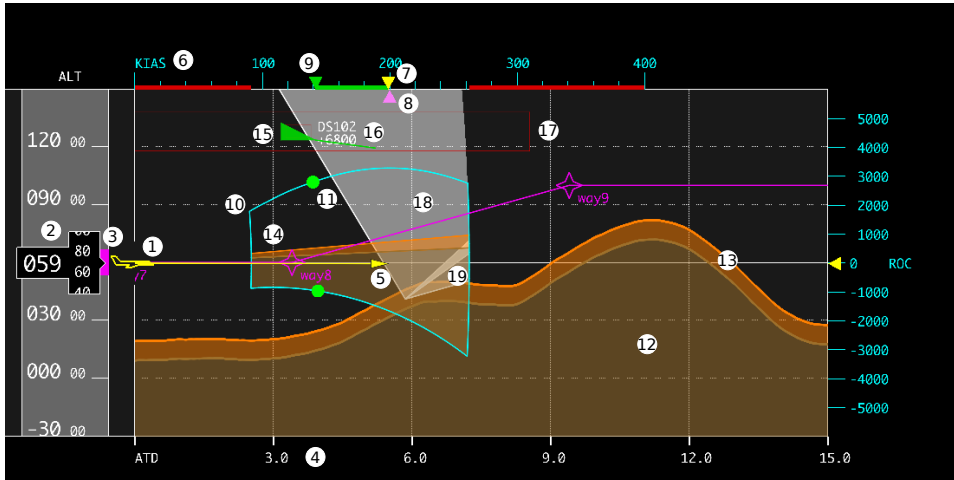


Figure C.2: The Vertical Situation Display

The advantage is that obstacles become more pronounced, but the disadvantage is that some correlation, especially for angles, is lost with respect to the PFD.

A yellow arrow represents the flight path vector (5). Its length scales with true airspeed and the angle corresponds to the flight path angle compensated for the display aspect ratio. The right side of the display shows the vertical projection of the velocity vector, resulting in the actual rate of climb.

The top part of the display shows the indicated airspeed (6) associated with the current flight path vector. The yellow indicator shows the current indicated airspeed (7). In horizontal flight, the yellow indicator will line up with the tip of the flight path vector. Since the display always shows true airspeed, and the conversion to indicated airspeed depends on altitude, the scale will expand or contract with increasing or decreasing altitude respectively. A magenta triangle indicates the required speed from the flight plan (8) and a green triangle indicates the best angle of climb speed (9). A green line connects the current speed with the best angle of climb speed to visualize the remaining speed that can be traded in for altitude. Finally, the minimum and maximum speed are indicated by red lines.

C.2.2 EID Cues

The performance envelope of the aircraft is shown in cyan (10). This envelope shows the limits of (steady optimal) performance on the flight path vector. The left and right side represent the minimum and maximum velocities respectively. The top part shows the maximum sustained climb performance. At full power,

when the velocity vector is above this line, the speed will always decrease as long as it is above this line. The bottom part indicates the best glide angles possible at idle power. At idle power, speeds above this line cannot be sustained. Two green dots indicate the optimal climb and optimal glide speed/angle (11).

The terrain is drawn at the correct distance and height, indicating its physical position (12). Since the flight path vector is drawn using true airspeed, the climb angle can be directly related to the terrain. If the vector points above the highest point, the aircraft will clear the terrain. The terrain is drawn with two different colors. The brown color is used to indicate the physical terrain, the orange layer shows a 1,000 *ft* safety margin (13). As an additional cue, two terrain conflict zones are drawn (14) indicating which combinations of flight path angle and velocity lead to a terrain collision or a violation of the safety margin with a lookahead of 15 *NM*.

Traffic is also drawn at the location where it appears (15). It consists of a triangle with a stem that indicates the direction of travel of the intruding aircraft (16). Two red boxes are drawn around the intruder (17). The large box shows the protected zone, the area that should be avoided to maintain the necessary safety separation. The smaller box represents the area occupied by the physical aircraft. Entering this area will result in a collision or severe damage to the aircraft.

Since the intruders are moving, the relationship between the flight path vector and the intruder are not as straightforward as with the terrain. To visualize this relationship, two conflict zones are drawn. The lighter one represents the protected zone (18), the darker one represents the physical aircraft (19). The conflict zones show which combinations of flight path vector and speed will result in a trajectory into the protected zone of or into the intruder. Keeping the flight path vector out of this conflict zone will ensure a safe distance to the intruding aircraft.

C.3 Experiment

C.3.1 Goal

The goal of this experiment is to evaluate a new constraint based display.

C.3.2 Scenario

All scenarios take place in a fictitious mountainous terrain. You will be placed at an intermediate altitude somewhere in this terrain. The only option is to fly straight climbing over the terrain if necessary. At some point traffic will pop-up giving you an extra constraint to take into account.

C.3.3 Task

During the experiment you will have to complete 4 tasks:

1. Set the throttle to the trim position indicated on the PFD
2. Track the altitude and speed of the flight plan as much as possible
3. Avoid both traffic and terrain in a way that feels safe
4. Provide feedback about strategy and performance

C.3.4 Apparatus

The experiment will be conducted in the Human-Machine Interface Laboratory located at the Faculty of Aerospace Engineering of the Delft University of Technology. The pilot will be seated in the right seat of the fixed-base simulator. A 18 inch monitor in front of the pilot will show the experiment display. The hydraulic side stick located on the right of the seat is used to control pitch. This is a 2 dimensional experiment so the roll and yaw can not be controlled. A trim switch is located on the top of the stick. The trim mechanism works like a conventional joystick trim, so it is important to move the stick towards the center position while trimming. The throttle on the left controls engine power.

C.3.5 Procedure

The experiment will start with a *briefing and training* phase during which the pilot can get familiar with the experiment and the control and dynamics of the simulation. Once the pilot feels comfortable with the display and control of the aircraft, the *measurement* phase starts. In this phase, the different measurement runs will be performed.

30 min	Briefing
60 min	Training
15 min	Break
30 min	Measurements (8 runs)
15 min	Debriefing

Bibliography

- Amelink, M. H. J., van Paassen, M. M., & Mulder, M. (2003). Total Energy-Based Perspective Flight Path Display for Aircraft Guidance along Complex Approach Trajectories. In *Proceedings 12th international symposium on aviation psychology*. Dayton (OH), USA.
- Bainbridge, L. (1983, nov). Ironies of automation. *Automatica*, 19(6), 775–779. doi: 10.1016/0005-1098(83)90046-8
- Bennett, K. B., & Flach, J. M. (2011). *Display and Interface Design : Subtle Science, Exact Art*. CRC Press.
- Billings, C. E. (1996a). *Aviation Automation – The Search for a Human-Centered Approach*. Lawrence Erlbaum Associates, Inc.
- Billings, C. E. (1996b). *Human-Centered Aviation Automation: Principles and Guidelines* (Tech. Rep.). Moffet Field (CA): NASA Ames Research Center.
- Bliss, J. P. (2003, jul). Investigation of Alarm-Related Accidents and Incidents in Aviation. *The International Journal of Aviation Psychology*, 13(3), 249–268.
- Boeing Commercial Airplanes. (2016). *Statistical Summary of Commercial Jet Airplane Accidents* (Tech. Rep.).
- Bolton, M. L., & Bass, E. J. (2008). Using Relative Position and Temporal Judgements to Identify Biases in Spatial Awareness for Synthetic Vision Systems. *The International Journal of Aviation Psychology*, 18(2), 183–206.
- Borst, C., Flach, J. M., & Ellerbroek, J. (2014, apr). Beyond Ecological Interface Design: Lessons From Concerns and Misconceptions. *IEEE Transactions on*

- Borst, C., Mulder, M., & van Paassen, M. M. (2010). Design and Simulator Evaluation of an Ecological Synthetic Vision Display. *Journal of guidance, control, and dynamics*, 33(5), 1577–1591.
- Borst, C., Suijkerbuijk, H. C. H., Mulder, M., & van Paassen, M. M. (2006, oct). Ecological Interface Design for Terrain Awareness. *The International Journal of Aviation Psychology*, 16(4), 375–400.
- Burian, B. K., Barshi, I., & Dismukes, K. (2005). *The Challenge of Aviation Emergency and Abnormal Situations* (Tech. Rep.). Moffet Field (CA): NASA Ames Research Center.
- Burns, C. M., & Hajdukiewicz, J. R. (2004). *Ecological Interface Design*. CRC Press.
- Casner, S. M., Geven, R. W., & Williams, K. T. (2013, jun). The Effectiveness of Airline Pilot Training for Abnormal Events. *Human Factors*, 55(3), 477–485. doi: 10.1177/0018720812466893
- Christoffersen, K., & Woods, D. D. (2002). How to make automated systems team players. In E. Salas (Ed.), *Advances in human performance and cognitive engineering research* (Vol. 2, pp. 1–12). Elsevier.
- Ellerbroek, J., Brantegem, K. C. R., van Paassen, M. M., de Gelder, N., & Mulder, M. (2013, may). Experimental Evaluation of a Coplanar Airborne Separation Display. *IEEE Transactions on Human-Machine Systems*, 43(3), 290–301. doi: 10.1109/TSMC.2013.2238925
- Ellerbroek, J., Brantegem, K. C. R., van Paassen, M. M., & Mulder, M. (2013, may). Design of a Coplanar Airborne Separation Display. *IEEE Transactions on Human-Machine Systems*, 43(3), 277–289. doi: 10.1109/TSMC.2013.2242888
- Ellerbroek, J., Mulder, M., & van Paassen, M. M. (2011, may). Evaluation of a Separation Assistance Display in a Multi-Actor Experiment. In *Proceedings of the 16th international symposium on aviation psychology*.
- Ellerbroek, J., Visser, M., Van Dam, S. B. J., Mulder, M., & van Paassen, M. M. (2011, sep). Design of an Airborne Three-Dimensional Separation Assistance Display. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 41(5), 863–875. doi: 10.1109/TSMCA.2010.2093890
- EUROCONTROL. (2006). *Air Transport Framework: The Current Situation* (Tech. Rep.). doi: SESAR TR DLM-0602-001-03-00

Bibliography

- EUROCONTROL. (2010). *Long-Term Forecast - Flight movements 2010-2030* (Tech. Rep.). doi: CND/STATFOR Doc415
- European Transport Safety Council. (2001). *Transport Accident and Incident Investigation in the European Union* (Tech. Rep.). Brussels.
- FAA. (2011). *Introduction to TCASII Version 7.1* (Tech. Rep.).
- Gerz, T., Holzäpfel, F., & Darracq, D. (2002). Commercial aircraft wake vortices. *Progress in Aerospace Sciences*, 38(3), 181–208. doi: 10.1016/S0376-0421(02)00004-0
- Hajdukiewicz, J. R., Burns, C. M., Vicente, K. J., & Eggleston, R. G. (1999, sep). Work Domain Analysis for Intentional Systems. In *Proceedings of the human factors and ergonomics society annual meeting* (Vol. 43, pp. 333–337). doi: 10.1177/154193129904300343
- Hill, D. (1994). Aircraft configuration management using constraint satisfaction techniques. In *International conference on control '94* (Vol. 1994, pp. 1278–1283). IEE. doi: 10.1049/cp:19940321
- Hollnagel, E., & Woods, D. D. (1983). Cognitive systems engineering: New wine in new bottles. *International Journal of Man-Machine Studies*, 18, 583–600.
- ICAO. (2016). *Doc 4444, Procedures for Air Navigation Services - Air Traffic Management* (Tech. Rep.). Montreal, Quebec: Author.
- JPDO. (2011). *Concept of Operations for NEXtGen Air Transportation System Version 3.2* (Tech. Rep.).
- Klomp, R. E., van Paassen, M. M., & Mulder, M. (2011). Air Traffic Control Interface For Creating 4D Inbound Trajectories. In J. Flach, M. Vidulich, & P. Tsang (Eds.), *Proceedings of the 16th international symposium on aviation psychology* (pp. 263–268). Dayton (OH), USA.
- Langewiesche, W. (1944). *Stick and Rudder: An Explanation of the Art of Flying*. New York: McGraw-Hill.
- Lovesey, E. J. (1977). The instrument explosion - a study of aircraft cockpit instruments. *Applied Ergonomics*, 8(1), 23–30. doi: 10.1016/0003-6870(77)90113-2
- Mercado Velasco, G. A., Borst, C., Ellerbroek, J., van Paassen, M. M., & Mulder, M. (2015, aug). The Use of Intent Information in Conflict Detection and Resolution Models Based on Dynamic Velocity Obstacles. *IEEE*

Transactions on Intelligent Transportation Systems, 16(4), 2297–2302. doi: 10.1109/TITS.2014.2376031

Naikar, N., Hopcroft, R., & Moylan, A. (2005). *Work domain analysis: Theoretical concepts and methodology* (Tech. Rep.). DSTO Air Operations Division.

Naikar, N., & Sanderson, P. M. (1999). Work Domain Analysis for Training-System Definition and Acquisition. *The International Journal of Aviation Psychology*, 9(3), 271–290.

NATS. (2014). *Aeronautical Information Circular P 3/2014 United Kingdom - Wake Turbulence* (Tech. Rep.). Middlesex, UK: Author.

Nolan, M. S. (2010). *Fundamentals of Air Traffic Control*. Cengage Learning.

NTSB. (1985). *China Airlines Boeing 747-SP* (Tech. Rep.).

NTSB. (1989). *United Airlines Flight 232* (Tech. Rep.).

NTSB. (2009). *Loss of Thrust in Both Engines After Encountering a Flock of Birds and Subsequent Ditching on the Hudson River*.

Parasuraman, R., & Riley, V. (1997, jun). Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 39(2), 230–253. doi: 10.1518/001872097778543886

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2008, jun). Situation Awareness, Mental Workload, and Trust in Automation: Viable, Empirically Supported Cognitive Engineering Constructs. *Journal of Cognitive Engineering and Decision Making*, 2(2), 140–160. doi: 10.1518/155534308X284417

Perrow, C. (1984). *Normal Accidents*. NY: Basic Books.

Rasmussen, J. (1983). Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and Other Distinctions in Human Performance Models. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-13(3), 257–266. doi: 10.1109/TSMC.1983.6313160

Rasmussen, J. (1985, mar). The role of hierarchical knowledge representation in decisionmaking and system management. *IEEE Transactions on Systems, Man, and Cybernetics*, 15(2), 234–243. doi: 10.1109/TSMC.1985.6313353

Rasmussen, J. (1997). Risk Management In A Dynamic Society : A Modelling Problem. *Safety Science*, 27(2), 183–213.

Bibliography

- Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive Systems Engineering*. Wiley.
- Rasmussen, J., & Vicente, K. J. (1989, nov). Coping with human errors through system design: implications for ecological interface design. *International Journal of Man-Machine Studies*, 31(5), 517–534. doi: 10.1016/0020-7373(89)90014-X
- Reason, J. T. (1990). *Human error*. Cambridge University Press.
- Reason, J. T. (2008). *The human contribution : unsafe acts, accidents and heroic recoveries*. Ashgate.
- Rijneveld, P. (2010). *Towards Integrating Traffic and Terrain Constraints into a Vertical Situation Display*.
- Rijneveld, P., Borst, C., Mulder, M., & van Paassen, M. M. (2010). Towards Integrating Traffic and Terrain Constraints into a Vertical Situation Display. In *Proceedings of the aiaa guidance, navigation, and control conference*.
- Roesel, N. J., & Vohs, K. D. (2012). Hindsight Bias. *Perspectives on Psychological Science*, 7(5), 411–426. doi: 10.1177/1745691612454303
- Ruijgrok, G. J. J. (2009). *Elements of Airplane Performance*. Delft: VSSD.
- Sarter, N. B., & Woods, D. D. (1995). How in the World Did We Ever Get into That Mode? Mode Error and Awareness in Supervisory Control. *Human Factors*, 37(1), 5–19. doi: 10.1518/001872095779049516
- Sarter, N. B., Woods, D. D., & Billings, C. E. (1997). Automation Surprises. *Handbook of Human Factors & Ergonomics*, 1926–1943.
- The Dutch Safety Board. (2011). *Emergency landing after bird strike* (Tech. Rep. No. June). The Hague.
- Van Dam, S. B. J., Mulder, M., & van Paassen, M. M. (2008). Ecological Interface Design of a Tactical Airborne Separation Assistance Tool. *IEEE Transactions on Systems, Man & Cybernetics, Part A*, 38(6), 1221–1233.
- van Paassen, M. M. (1999). Functions of Space and Travellers. In *Proceedings of 18th european annual conference on human decision making and manual control*. Delft.
- van Paassen, M. M., Borst, C., Klomp, R. E., Mulder, M., van Leeuwen, P., & Mooij, M. (2013). Designing for shared cognition in air traffic management. *Journal*

of *Aerospace Operations*, 2(1-2), 39–51. doi: 10.3233/AOP-130030

van Paassen, M. M., Gernaey, J., Veld, A. C. i. t., & Mulder, M. (2007). Using Ecological Interface Design for Energy Management During Idle-Thrust Approaches. In R. Jensen (Ed.), *Isap-07* (pp. 727–732). Dayton (OH), USA: Wright State University.

Vicente, K. J. (1999). *Cognitive Work Analysis*. Lawrence Erlbaum Associates, Inc.

Vicente, K. J., & Rasmussen, J. (1990). The Ecology of Human-Machine Systems II: Mediating Direct-Perception in Complex Work Domains. *Ecological Psychology*, 2(3), 207–249.

Vicente, K. J., & Rasmussen, J. (1992). Ecological interface design: theoretical foundations. *Systems, Man and Cybernetics, IEEE Transactions on*, 22(4), 589–606. doi: 10.1109/21.156574

Withey, P. A. (1997, jun). Fatigue failure of the de Havilland comet I. *Engineering Failure Analysis*, 4(2), 147–154. doi: 10.1016/S1350-6307(97)00005-8

Woods, D. D., Johannesen, L. J., Cook, R. I., & Sarter, N. B. (1994). *Behind Human Error: Cognitive Systems, Computers, and Hindsight* (Tech. Rep. No. December). CSERIAC.

Glossary

AH Abstraction Hierarchy

APU Auxiliary Power Unit

ATC Air Traffic Control

ATCo Air Traffic Controller

ATD Along Track Distance

CD&R Conflict Detection & Resolution

CFIT Controlled Flight Into Terrain

CPA Closest Point of Approach

CSE Cognitive Systems Engineering

CWA Cognitive Work Analysis

CZ Conflict Zone

dCPA distance at Closest Point of Approach

EGPWS Enhanced Ground Proximity Warning System

EID Ecological Interface Design

ESVD Ecological Synthetic Vision Display

EVSD Enhanced Vertical Situation Display

HMS Human-Machine System

ISVD Intentional Synthetic Vision Display

IVSD Intentional Vertical Situation Display

KBB Knowledge-Based Behavior

MSA Minimum Safe Altitude

MTC Minimum Terrain Clearance

ND Navigation Display

NTSB National Transportation Safety Board

PFD Primary Flight Display

PPL Private Pilot License

PZ Protected Zone

RBB Rule-Based Behavior

SBB Skill-Based Behavior

SRK Skills, Rules, Knowledge

SVD Synthetic Vision Display

TAS True Airspeed

TCAS Traffic Collision Avoidance System

VSD Vertical Situation Display

WDA Work Domain Analysis

Samenvatting

De toename in automatisering heeft een belangrijke rol gespeeld in de snelle ontwikkeling van de luchtvaart gedurende de laatste decennia. Dit heeft geleid tot een sterke toename van veiligheid en efficiëntie. Er bestaat consensus dat automatisering een prominente rol zal blijven spelen in de luchtvaart. Niet uitsluitend als het gaat om het verhogen van de mate van automatisering, maar ook in de ontwikkeling van automatisering waarbij de mens als operator een centrale rol blijft spelen.

Één van de belangrijkste lessen die, in de luchtvaart en daarbuiten, is geleerd, is dat bij het verplaatsen van taken van de mens naar geautomatiseerde systemen, fouten en ongevallen kunnen worden vermeden. Hierbij ontstaan echter ook nieuwe mogelijkheden voor andere soorten fouten en ongevallen. Deze archetypische 'ironieën van automatisering' zijn overal aanwezig in de luchtvaart.

Inschikkelijkheid en verveling zijn slechts enkele van de problemen die naar voren komen bij het gebruik van sterk geautomatiseerde systemen. Één van de uitdagingen in toekomstige ontwerpen van automatisering ligt in het vinden van de juiste balans in de verdeling van taken tussen automatisering en mensen. Een groot deel van deze uitdaging ligt in het domein van de mens-machine interactie en mensgerichte automatisering.

Ecological Interface Design (EID) heeft als doel de mens als bestuurder te ondersteunen in complexe socio-technische systeem domeinen die, soms, uitgebreide operationele procedures hebben. Experimenten met een aantal experimentele human-in-the-loop interfaces voor de luchtvaart hebben laten zien dat ecologische interfaces een betere ondersteuning kunnen bieden aan piloten in complexe situaties. Ecologische interfaces hebben als doel de mogelijkheden en beperkingen van het systeem in beeld te brengen. Dit kan echter ook soms leiden tot situaties waarbij te dicht op de grenzen van het systeem wordt geopereerd.

Het opzoeken van deze grenzen kan risico's met zich mee brengen wanneer de operator zich dichtbij fysieke beperkingen begeeft waarvan overschrijding kan leiden tot zware incidenten of ongevallen. Deze kans op risico migratie is niet

noodzakelijk het gevolg van de EID methodiek zoals soms aangenomen wordt. Het visualiseren van regels, voorschriften, en procedures kan mogelijk het opzoeken van grenzen beperken. Dit leidt tot de probleemformulering: *Kunnen we, door het duidelijk visualiseren van regels, voorschriften, en procedures, ecologische interfaces ontwerpen die leiden tot veiligere mens-machine systeem prestaties?*

Deze probleemformulering geeft aan dat de EID methodologie expliciet rekening moet houden met regels, voorschriften, en procedures. Daarnaast is een analyse nodig van de impact van het invoeren van een dergelijke methode. Dit leidt tot de volgende drie onderzoeksvragen van dit proefschrift:

1. *Hoe passen regels, voorschriften, en procedures in het EID raamwerk?*
2. *Kunnen piloten een onderscheid maken tussen fysieke beperkingen en beperkingen die voortvloeien uit regels, voorschriften, en procedures wanneer deze in beeld worden gebracht op de interface?*
3. *Zullen piloten betere beslissingen nemen op basis van deze extra informatie?*

De eerste onderzoeksvraag behandelt de basis van een EID interface, de Work Domain Analysis (WDA), die leidt tot de Abstraction Hierarchy (AH). De AH is een weergave van het werk domein op verschillende niveaus van abstractie. Functies op aangrenzende niveaus zijn verbonden door middel van middelen-doelen verbindingen. De AH laat de *functionele* structuur zien, gekenmerkt door de doelen en de beperkingen die het dynamisch gedrag limiteren. De beperkingen op het systeem kunnen in twee categorieën verdeeld worden.

De eerste categorie omvat de *causale* beperkingen. Deze beperkingen komen voort uit de natuurwetten en de fysieke processen die van toepassing zijn. Deze beperkingen vormen de fysieke grenzen van het systeem. Causale beperkingen zijn overschrijdbaar dan wel niet-overschrijdbaar. Overschrijdbare beperkingen vertegenwoordigen grenzen die overschreden kunnen worden en daarbij vervolgens leiden tot een ernstige degradatie van het systeem, of zelfs tot ongevallen. Een voorbeeld van een overschrijdbare beperking is een obstructie op een weg waar een auto tegenaan kan rijden. Niet-overschrijdbare beperkingen vertegenwoordigen beperkingen met asymptotisch gedrag. Deze resulteren niet in systeem degradatie, maar limiteren het gedrag van het systeem. Een voorbeeld hiervan is de maximale snelheid van een auto. Het is onmogelijk sneller te rijden, maar dit heeft geen invloed op de integriteit van het systeem.

De tweede categorie beperkingen wordt gestuurd door de intenties, waarden, regels, enz. van de actoren. Deze worden *intentionele* beperkingen genoemd. Intentionele beperkingen vormen geen fysieke beperking. Hun doel is het sturen van gedrag. Diverse redenen kunnen een drijfveer zijn voor dergelijke beperkingen, zoals bijvoorbeeld: veiligheid, efficiëntie, winst, enz. In dit proefschrift ligt de focus op intentionele beperkingen die het gevolg zijn van regels, voorschriften, en procedures gerelateerd aan veiligheid. Een voorbeeld van een intentionele beperking is de snelheidslimiet waar auto's zich aan moeten houden. Een auto kan

perfect in staat zijn om deze limiet te overschrijden, maar het in acht nemen van deze limiet zal de veiligheid verhogen voor alle verkeersdeelnemers.

De implicaties voortvloeiende uit beide categorieën beperkingen zijn echter zeer verschillend. Causale beperkingen bepalen hoe men *kan* handelen, intentionele beperkingen bepalen hoe men *moet* of *zou willen* handelen. Dit onderscheid kan cruciaal zijn in onverwachte situaties en noodsituaties. Bij het vliegen bijvoorbeeld is het belangrijk om te allen tijde hoge bouwwerken te vermijden. Het verbod om over een gebied te vliegen vanwege geluidsnormen kan echter genegeerd worden wanneer dit zou helpen een noodsituatie veiliger op te lossen. In de huidige EID interfaces wordt geen expliciet onderscheid gemaakt tussen causale en intentionele beperkingen. Sommige interfaces visualiseren enkel causale beperkingen, andere visualiseren beide, zonder een duidelijk onderscheid te maken.

Op basis van het bovengenoemde onderscheid zijn in dit proefschrift twee bestaande experimentele cockpit interfaces geanalyseerd, met als doel te bepalen welke types van beperkingen gevisualiseerd werden, om vervolgens te onderzoeken hoe deze interfaces aangepast konden worden om een expliciet onderscheid tussen causale en intentionele beperkingen aan te brengen.

Een Ecological Synthetic Vision Display (ESVD) werd gebruikt als referentie interface voor de eerste case study. Deze interface bestaat uit een traditioneel Primary Flight Display met daarop een drie-dimensionale synthetische weergave van het terrein, gecombineerd met een indicatie van de beschikbare klim capaciteit. De visualisatie van de klim capaciteit kan eenvoudig aan het terrein gerelateerd worden, waardoor een snelle inschatting gemaakt kan worden of het mogelijk is om over het terrein te klimmen.

De belangrijkste elementen voor een terrein vermindings interface zoals het ESVD zijn: het virtuele terrein, de indicatie van de manoeuvreer afstand, en de minimaal benodigde hellingshoek. Deze drie elementen zijn allen gebaseerd op fysieke beperkingen. De focus van de case study lag op het combineren van zowel de horizontale als de verticale separatie minima. Op basis van deze criteria werd een nieuwe interface ontworpen die een visualisatie van deze separatie buffers bevat waarmee piloten zowel de vereiste veiligheidsmarges alsook de onderliggende fysieke obstakels kunnen zien.

De tweede case study gebruikte een Enhanced Vertical Situation Display (EVSD) als referentie interface. Het EVSD toont een zijaanzicht van het eigen vliegtuig, het aankomend verkeer en terrein, samen met de operationele begrenzings van het vliegtuig. Binnen deze begrenzings worden driehoekige conflict zones getoond die laten zien welke snelheden en klimhoeken leiden tot een conflict met aankomend verkeer, samen met een lijn die de benodigde klimhoek aangeeft om het terrein te vermijden met een bepaalde marge.

De analyse liet zien dat de huidige EVSD interface een combinatie van causale en intentionele beperkingen bevat, zonder een expliciet onderscheid te maken tussen de twee. Beperkingen vanwege het verkeer worden gevisualiseerd

door middel van een gevulde driehoek die een buffer om het aankomende vliegtuig aangeeft. Dit resulteert in een zuiver intentionele beperking, die de onderliggende fysieke beperking verbergt. De beperking vanwege het terrein wordt weergegeven door een visualisatie van het terrein in combinatie met een lijn die de minimaal vereiste hoogte boven het terrein aangeeft. Dit combineert beide soorten beperkingen, maar in twee verschillende representaties die moeilijk aan elkaar te relateren zijn. Daarbij komt dat de visualisatie van de verkeersgerelateerde beperking veel prominenter aanwezig is dan de beperking vanwege het terrein.

Op basis van deze analyse werd een interface uitgewerkt die een duidelijke scheiding aanbrengt tussen causale en intentionele beperkingen samen met een meer uniforme visualisatie voor beide beperkingen.

Na deze theoretische analyse gaat het proefschrift verder met twee hoofdstukken die de experimentele evaluatie van de twee concept interfaces beschrijft.

In het eerste experiment werden de toevoegingen in een nieuw Intentional Synthetic Vision Display (ISVD) vergeleken met het referentie ESVD. Na een trainingsperiode vlogen zestien professionele piloten vier meetscenario's waarbij ze in een terrein conflict gebracht werden. Dit deden ze zowel met de referentie interface als de verbeterde ISVD interface.

Een analyse van de resultaten liet een toename zien in minimale afstand tot het terrein tijdens hun ontwikkelingsmanoeuvres. Visualisatie van intentionele beperkingen resulteerde in een betere naleving van de intentionele beperkingen. Bovendien waren de piloten in staat om de extra informatie intuïtief te gebruiken. In situaties die eenvoudig op te lossen waren werden de toevoegingen gebruikt om hun prestaties te finetunen, om zo afwijkingen van de route te beperken en brandstof te besparen terwijl ze nog steeds aan de minimum hoogte eisen voldeden. In moeilijker scenario's met minder mogelijkheden konden piloten gebruik maken van de representatie om hun gekozen strategie nauwkeurig uit te voeren. Bijvoorbeeld, wanneer de piloot besloot een bocht te maken, kon deze onmiddellijk zien in welke richting deze het meest efficiënt uitgevoerd kan worden en wanneer hij de bocht kon stoppen.

In het tweede experiment werd het referentie EVSD vergeleken met het verbeterde Intentional Vertical Situation Display (IVSD) dat zowel causale als intentionele beperkingen laat zien. Zestien proefpersonen, van wie acht piloten en acht studenten, werden verdeeld in twee groepen. Na training gebruikten één groep de referentie interface om vier scenario's met terrein en verkeer op te lossen. De andere groep gebruikte de IVSD interface om dezelfde scenario's op te lossen.

Dit experiment liet niet de verwachte toename in hoogte boven het terrein zien, noch waren er minder overschrijdingen van de beschermde zone. Daarnaast liet het experiment ook niet zien dat er een duidelijk verschil was in strategie tussen de ervaren en onervaren proefpersonen. Er was wel een beperktere spreiding van de gevlogen trajecten zichtbaar bij het vergelijken van de referentie interface met het IVSD. Dit geeft aan dat de proefpersonen die het IVSD gebruikte

zich meer bewust waren van de beperkingen en dus in staat waren hun strategie nauwkeuriger uit te voeren.

Samenvattend, dit proefschrift liet zien hoe regels, voorschriften, en procedures passen in het EID raamwerk door ze correct te identificeren en weer te geven in de AH. Dit proefschrift stelt voor de AH te splitsen in een causaal en intentioneel domein. Op deze manier kan een duidelijk onderscheid aangebracht worden tussen de rigide causale beperkingen en de meer flexibele intentionele beperkingen. Dit zorgt er voor dat hetzelfde onderscheid aangebracht kan worden tijdens het ontwerp van de interface, wat op zijn beurt piloten kan helpen in het stellen van prioriteiten, vooral in uitzonderlijke situaties.

Het experiment met het ISVD liet zien dat piloten in staat waren gebruik te maken van een duidelijke visualisatie van intentionele beperkingen. Op basis van het IVSD experiment is het moeilijk om duidelijke verbetering vast te stellen als gevolg van de visualisatie van intentionele beperkingen. Dit is echter waarschijnlijk eerder het gevolg van ontoereikende training en scenario ontwerp dan van de visualisatie zelf. Het ISVD experiment laat zien dat een expliciete visualisatie van intentionele beperkingen piloten in staat stelt een onderscheid te maken tussen causale beperkingen en intentionele beperkingen door regels, voorschriften, en procedures, maar nadere validatie hiervan is echter vereist.

Het ISVD experiment liet zien dat het besluitvormingsproces van piloten verbetert wanneer ze in staat zijn een onderscheid te maken tussen causale en intentionele beperkingen. Het IVSD experiment liet niet dezelfde verbetering zien, hoofdzakelijk vanwege gebreken in de training en het ontwerp van scenario's. Een enkel experiment volstaat niet om een hypothese te bevestigen, maar de resultaten geven aan dat verder onderzoek de moeite waard kan zijn.

Er kunnen een aantal mogelijke redenen genoemd worden waarom het IVSD experiment het verwachte verschil tussen de referentie interface en de toevoegingen niet kon laten zien. Vooral de beperkte training zorgde er voor dat de proefpersonen de vertrouwdheid met de complexe interface misten die nodig was voor dit experiment. Daarnaast zorgden de gekozen scenario's er niet voor dat de proefpersonen in die situaties terecht kwamen waarbij de nieuwe interface een voordeel zou kunnen opleveren. Ten laatste zou een grotere groep proefpersonen nodig zijn.

Acknowledgements

What a flight this has been, even though the landing runway wasn't always in close view, having crossed the finish line I can only be grateful for everything I have learned along the way. This was definitely no solo flight. I would not have been able to do this without a great group of people surrounding and supporting me throughout these years.

First and foremost, I would like to thank my parents, Raf and Kristien, for being the most supportive parents anyone can wish for. For as long as I can remember you have allowed and encouraged me to pursue all of my passions, from my first flights from Kiewit airfield to my (one-way) journey to the Netherlands, even when the latter involved a brisk fade of my Flemish accent. You are truly the foundation of who I am today. Furthermore, I also have to thank my sister Evva, and my brother-in-law Chris, who were always there for me. Despite the ocean separating us, you always had my back, and always welcomed me to unwind during my holidays.

My PhD adventure started when Max gave me the opportunity to pursue a PhD after getting a foothold on the Control and Simulation grounds during my MSc. I am thankful for all your support and feedback, and keeping faith in me through the past years. Furthermore, I also want to thank our section guru René for providing key insights at crucial moments. Finally, I also have to thank Clark. Your help and guidance have been extremely important throughout all these years. I truly enjoyed all of our serious and less serious conversations.

I would like to thank all my colleagues and friends that I have worked with over all these years. The first person I need to mention is Paul, who would have expected that being in the starting blocks of our PhDs together would lead to getting you and Renske married some years after, thank you for being my partner in crime all those years. A shout out to Daan and Rita, my basketball buddies. From the basketball field to that Portuguese wedding, being able to share this experience with you was a great joy. A big thank you also goes out to Erik-Jan,

Coen, Herman, and Joost, for all the seriousness and lack thereof at coffee breaks and Friday evening drinks. Liguó, Deniz, Jia, Laurens, Wouter, Yazdi, Sophie and Jaime, you made my life at the office at any of its assigned room numbers a pleasurable experience, for which I want to thank you. The same goes out to my colleagues in the other offices, Rolf, Maarten, João, Immanuel, Tomasso, Dia, Jan, and Tao.

Also many thanks to Olaf, for allowing me to occasionally play with Simona, even though that wasn't strictly required for my research. I also need to thank Hans and Xander, for bringing some (quite literal) lightness along this journey by giving me the opportunity to play a role in the flight practical, and for all the time I could spend in our research aircraft.

During my time the C&S department, I saw the group grow from a small cozy group to a steadily-growing research group that attracts more students than ever. To all of those I have worked with the past years but haven't been mentioned before, thank you for being there along my journey.

Ben and Anneleen, thank you for being great supportive friends throughout this whole experience and for being a piece of home away from home in Delft. Processing all our PhD experiences over dinners and game nights definitely kept me going.

Finally, I would like to thank my girlfriend, Madelon. You swiped your way into my life at what was probably the toughest part of my PhD. Without your support and encouragement this would have never worked. Thank you. I love you.

Curriculum Vitae

Jan Comans was born on June 25, 1984 in Hasselt, Belgium. He attended secondary school at Humaniora Kindsheid Jesu in Hasselt from which he graduated in 2002. After the summer, he moved to Delft where he started studying at the Aerospace Engineering Faculty of the Delft University of Technology. He obtained his Bachelor of Science in 2006 and then started his graduate studies at the Control and Simulation Department. During this period, he spent half a year at Boeing Research and Technology Europe in Madrid for an internship. This was followed by a graduation project on visual delay reduction in the Simona research simulator. He obtained his Master of Science degree in July 2009. After presenting his graduation work at the AIAA Modeling and Simulation Technologies Conference in Chicago, IL, he started working towards his PhD at the Aerospace Software and Technologies Institute.

He started his work within the Cognitive Adaptive Man-Machine Interfaces (CAMMI) project for which he represented the TUD at the project meetings and for the technical work. After CAMMI, he shifted his focus to his main research topic on Ecological Interface Design (EID). Besides his research, Jan was involved in the supervision of a number of master students and was heavily involved in the third year flight practical.

In December 2014, Jan started working at Sim-Industries in Sassenheim where he worked as a software engineer on the flight dynamics model and QTG's of a full-flight Boeing 787 simulator. In September 2016, he joined TASS International as a software engineer working in the real-time team on PreScan, a simulation framework for the automotive industry.

Publications

- *Comans, J, Stroosma, O, Paassen, MM van & Mulder, M (2009). Optimizing the Simona Research Simulator Visual Display System. In John D. Schier-*

man (Ed.), Proceedings of 2009 AIAA GNC/AFM/MST (pp. 1-19). Reston (VA): AIAA.

- *Comans, J, Paassen, MM van & Mulder, M (2010). Context Aware Human-Centered Adaptive Automation in Aviation. In C Jonker, V Evers & B Riemsdijk (Eds.), Proceedings of the D-CIS Human Factors Event (pp. 14-15). Delft: D-CIS Lab.*
- *Comans, J, Paassen, MM van & Mulder, M (2010). Pilot Workload Monitoring and Adaptive Aviation Automation - A solution Space-Based Approach. In Willem.Paul Brinkman & Max Neerincx (Eds.), Proceedings of the European Conference on Cognitive Ergonomics 2010 (pp. 245-250). Delft: European Association of Cognitive Ergonomics.*
- *Lodder, J, Comans, J, Paassen, MM van & Mulder, M (2011). Altitude- Extended Solution Space Diagram for Air Traffic Controllers. In J Flach (Ed.), Proceedings of the International Symposium on Aviation Psychology (pp. 345-350). Dayton, OH: Wright State University.*
- *Comans, J, Borst, C, Paassen, MM van & Mulder, M (2011). Solution Space-Based Complexity Analysis for Context Aware Automation. In J Flach (Ed.), Proceedings of the International Symposium on Aviation Psychology (pp. 475-480). Dayton, OH: Wright State University.*
- *Paassen, MM van, Ellerbroek, J, Comans, J, Borst, C & Mulder, M (2012). Ecological Interface Design for Aircraft Guidance and Conflict Avoidance. In A Bernard & F Chinesta (Eds.), Proceedings of the ASME 2012 11th Biennial Conference on Engineering Systems Design and Analysis (ESDA 2012) (pp. 1-8). American Society for Mechanical Engineers.*
- *Comans, J, Borst, C, Paassen, MM van & Mulder, M (2013). Risk Perception in Ecological Information Systems. In PS Tsang (Ed.), Proceedings of the 17th International Symposium on Aviation Psychology (pp. 436-441). Dayton, OH, USA: Write State University.*
- *Comans, J, Borst, C, Mulder, M & Paassen, MM van (2014). Risk Perception in Ecological Information Systems. In MA Vidulich, PA Tsang & JM Flach (Eds.), Advances in Aviation Psychology (pp. 121-138). Farnham: Ashgate.*

Workshops

- *Paassen, MM van, Borst, C, Ellerbroek, J & Comans, J (2011, May 2). Applied Cognitive Systems Engineering in Aviation. Dayton, Ohio, USA, Workshop cognitive systems engineering.*

Curriculum Vitae

- Borst, C, Paassen, MM van, Ellerbroek, J & Comans, J (2013, May 6). Applied Cognitive Systems Engineering in Aviation. Dayton OH, Full-day Workshop at Wright State University as part of the 17th International Symposium on Aviation Psychology.

