Advanced Interactive Display Formats for Terminal Area Traffic Control

A Joint Research Project Between NASA Ames Research Center,
Moffett Field, CA., and The Faculty of Aerospace Engineering at the Technion,
Israel Institute of Technology, Haifa, 32000, ISRAEL

Final Report for the period: February 1, 1998 - January 31, 1999

Cooperative Grant No.: NCC 2 - 1023

Investigators:

Dr. Arthur J. Grunwald[1] and Dr. G. E. Shaviv[2]
Faculty of Aerospace Engineering
Technion, Israel Institute of Technology
Haifa, 32000, ISRAEL

Contract Technical Monitor:

Dr. Stephen R. Ellis
Mail Stop 262-2
NASA Ames Research Center
Moffett Field, CA., 94035
USA

June 15, 1999

---

1. Principal investigator
2. Presently NRC Post Doctoral Research Associate at NASA Ames Research Center. Parts of this final report have been published as G. E. Shaviv's Ph.D thesis, submitted to the Faculty of Aerospace Engineering of the Technion

Preface

This report summarizes the activities of the last three years of a multi-year research cooperation between NASA Ames Research Center and the Faculty of Aerospace Engineering, Technion, Haifa, Israel, in the area of advanced pictorial displays for Aerospace applications. The main drive for carrying out this research project was to find the most suitable way of communicating complex dynamic spatial information to a human observer. The air traffic control display has been chosen as an interesting application in which the key elements are excessive operator workload, the need for simultaneously managing large volumes of complex dynamic data, and stringent requirements for safety.

Although spatial data sets are most naturally viewed in a perspective projection that integrates the horizontal and vertical information in one *natural* format, these formats are not necessarily the most suitable ones for the typical air traffic control task. The 'classical' air traffic control display shows a 'north-up' map of

the area in which the aircraft are plotted as symbols, furnished with numerical information about its identification, its intended destination, altitude and airspeed. Many researchers have made comparisons between these classical formats and more advanced perspective displays. Although their conclusions are often clear-cut and their findings indicate that one display type is performing better than another one, these findings can rarely be applied to the fully fledged air traffic control task in a busy TRACON area. The same applies to our results. Therefore, the merits of our work are not in the design of a novel ATC display, but rather in the developments of tools and the establishment of a sound methodology for understanding and optimizing the transfer of spatial information to the human operator through perspective formats.

Air traffic control consists of a number of diverse and complex activities, geared towards very specific objectives, ranging from maintaining a steady well-separated traffic flow and solving unforeseen conflicts or handling emergency situations, to planning national traffic patterns many hours ahead of time. Since perspective formats, by nature, represent the spatial situation *as seen from a given vantage point*, objects that are close by and viewed centrally are given more attention than objects that are far away and at the edge of the visual field. Hence, the separation between far away aircraft in the peripheral vision might be difficult to judge. In a flow control task this might pose a problem, since separation violations might occur anywhere and at any time. Therefore, terminal area controllers dealing with dense and interlacing traffic flows, might favor the constancy of north-up, plan view displays.

In contrast to the flow control task, the traffic planning task demands insight in three-dimensional traffic patterns and terrain lay-outs. In a re-routing task, for example, the three-dimensional shape of a weather front might have to be considered. Properly designed perspective formats should be geared towards providing the operator with a clear and unambiguous understanding of the spatial situation.

The basic principle underlying our research activities is the notion that the viewing parameters of perspective displays should be optimally adjusted to the spatial lay-out that is viewed and to the task that the operator is performing.

This adjustment could be done manually as well as automatically. Early stages of the research dealt with manual viewing parameter schemes. 'Manipulation handles' were naturally integrated in the three-dimensional scene and enabled the operator a straight-forward and intuitive way of manipulating the viewing parameters. Special motion patterns were designed, resembling the motions of a well-designed servo system, that allowed gradual transitions from one parameter setting to another.

The method by which the viewing parameters are optimally adjusted is quite unique, and has led to the formulation of a concept to which we refer to as *an active display system*. The term *active* refers to the fact that the display no longer functions as a passive window frame, in which spatial data are projected, but acts as an *intelligent* system, that continuously and actively analyses, selects, processes and presents the spatial information in such a way, that the operator is able to derive from it the best understanding of the dynamic spatial situation. The interested reader will discover that this *intelligent* process involves four basic steps: *analyzing, optimizing, deciding and executing*, of which the first two steps are the most complex ones. *Analyzing* is complex, since it involves deciding which aspects of a viewed spatial lay-out are important. *Optimizing* is complex because it requires the definition of a relevant cost function, which should include a model of the human spatial perception process. *Deciding* and *executing* deal with the question whether, when and how the changes, suggested by the two previous steps should be realized.

Apart from the practical issue of how to numerically find the optimum of a complex multi modal cost surface, the main issue with any optimizer, remains to prove that the selected optimum viewing parameter setting indeed represents the best way of presenting the spatial information to the operator. A great deal of efforts were devoted to proving this last issue. Two large experimental studies were carried out, each of which lasting for one year, in which active air traffic controllers participated. These studies considered both static as well as dynamic aspects of interpreting spatial layouts. Although the tasks were quite abstracted and structured to evaluate specific aspects of the human response, the results

clearly proved that properly optimized scenes yield superior performance, in particular when complex spatial data sets are involved.

After three years of work, we feel that we are at the beginning of an exciting venture, rather than at the end if it. Rather than providing a ready product, the merits of our work are in the definition of a novel and unique concept, which has clearly proved its potential in a series of structured laboratory experiments. Now the time has come to evaluate the concept in realistic air traffic scenarios. Here, the *analyzing and optimizing* steps will demand a great deal of involvement of experienced air traffic control personnel, who will be able to give the vital clues on how the essential information, necessary for performing the task, is hidden in the layout. We hope that the reader will enjoy this report and might find the inspiration in it to try out new horizons in this exciting field of spatial information transfer.

Arthur Grunwald, Haifa, June, 1999

Table of Contents

List of Figures

List of Tables

x

Abstract

This research project deals with an on-line dynamic method for automated viewing parameter management in perspective displays. Perspective images are optimized such that a human observer will perceive relevant spatial geometrical features with minimal errors. In order to compute the errors at which observers reconstruct spatial features from perspective images, a visual spatial-perception model was formulated. The model was employed as the basis of an optimization scheme aimed at seeking the optimal projection parameter setting. These ideas are implemented in the context of an air traffic control (ATC) application. A concept, referred to as an active display system, was developed. This system uses heuristic rules to identify relevant geometrical features of the three-dimensional air traffic situation. Agile, on-line optimization was achieved by a specially developed and custom-tailored genetic algorithm (GA), which was to deal with the multi-modal characteristics of the objective function and exploit its time-evolving nature.

Two series of part-task experiments were conducted, in which active air traffic controllers participated, as well as Aerospace Engineering students, who acted as a control group. Abstracted air traffic scenarios were presented to the subjects, who were required to perform spatial tasks that were designed to measure the level of their spatial awareness and comprehension. In the first experiment, still images of static scenarios were presented and the experiment task demanded comprehension of the current geometrical state. In the second experiment, moving images of dynamic scenarios were presented and the task demanded comprehension of both present and future geometrical states. A novel approach was used to evaluate the subjects' judgement capabilities using an interactive routing chart, presented on a data tablet.

The results clearly show that the spatial awareness of the operator improves with the level of optimization of the viewing parameters, with complex scenes benefiting most from the optimization. Moreover, the performance of the subjects with the optimally chosen viewpoints was found to be significantly better then the one with their own manually chosen viewpoints or with two-dimensional plan-view displays. Enhanced graphics in the form of velocity vectors and altitude predictors aided in the spatial comprehension of scenes, with the biggest improvement for the optimized perspective images. Subjective evaluations revealed that the judgment of spatial features in the optimized perspective displays was as easy to the users, as it was in the plan-view displays.

The experimental results prove that perspective displays with viewing parameter optimization are an effective means of delivering spatial information. This means that the optimal manner to transfer the required spatial information to the operator, is by means of an optimized perspective display, of which the objective function is configured such that it correctly reflects the operator's task.

1 Overview

1.1 Scope of the Research

This final report describes the development and experimental evaluation of an advanced perspective display format for enhanced spatial perception, and its implementation for Air Traffic Control (ATC). Pictorial perspective formats portray more naturally the three-dimensional world than other schematic display formats, e.g. planar displays, and hence they are potentially better suited to display three-dimensional information. In particular, perspective displays can more genuinely represent the air traffic information in an airspace than is possible in the conventional plan-view displays currently used in ATC. Regarding the human operator as an information processor, the research is focused at optimizing the information transfer and retrieval of three-dimensional data from human interfaces.

1.2 Background and Literature Survey

1.2.1 State of Air Traffic Control

While perspective displays have a broad range of applications, the concepts in this research are demonstrated for specific ATC applications. Since the ATC task demands the operator to have a great deal of understanding of the dynamic spatial situation, this task constitutes a challenge for the design of spatial information displays. A summary of the state of ATC is provided, with particular emphasis on the spatial aspects of the ATC task.

The Task: Safety and Efficiency

The air traffic control system is responsible for managing the air traffic, which consists of a complex mixture of commercial, general, corporate and military aviation. The system is required to maintain a high level of safety while efficiently providing the capacity to handle an ever increasing number of flights (NRC 1997). A projection made by the FAA (Perry 1997) states that by the year 2015 there could be a major aviation accident every 7 days. The projection is based on the anticipated growth of air traffic, combined with an accident rate that has been statistically flat at 10^{-7} accidents per flight hour in commercial flights for the past 15 years (Pelegrin 1998).

Guaranteeing minimum separation between aircraft is the primary means for ensuring safety. The minimum separation is defined in the horizontal and vertical dimensions, creating a cylinder of space around each aircraft. The separation criteria dimensions vary in different regions of the airspace and under different flight conditions. Increasing the safety level at the current technology would mean increasing the separation distances between aircraft. Alas, this would compromise the system's second goal, efficiency. The goal of efficiency is fourfold; maximizing the number of flights, minimizing delays, which are defined as the difference between a flight's scheduled time-of-arrival and it's estimated time-of-arrival had there been no other traffic; minimizing fuel consumption; and minimizing impact to the environment.

The process by which efficiency is met is more complex and constrained than the process by which safety is ensured. The maximum possible capacity is usually limited by the rate-of-arrivals at airports, particularly at large hubs. The number of gates, the number of runways and the speed at which an aircraft can

clear the runway, all dictate the constraints on an airport's capacity, which is defined by the number of flights it is able to receive per unit of time. The terminal area air traffic controller's goal is to optimize the traffic flow, by lining-up the aircraft for final approach and delivering them at regularly spaced time intervals, meeting airport capacity constraints. Several obstacles are to be overcome:

1) Aircraft cannot be "stacked" before the arrival airport, to be delivered as soon as a slot is available. To realize airport capacity, departures and upstream speed changes need to be strategically scheduled well in advance. Optimization is limited by the ability to predict uncertain factors that influence the flight schedule, e.g. head-winds.

2) Wake vortices, following the passage of an aircraft, force the controllers to maintain greater separations on the final approach, in particular for heavy aircraft.

3) Sudden changes in weather may force changes in airport configuration, e.g. closing or reversing runways, or force the rerouting of air traffic.

The air traffic system is continuously pressured for increased efficiency in its management of the spatial traffic flow. This system relies on the skills of its controllers and their ad-hock problem-solving ability, to address the sometimes-conflicting pressures for safety and efficiency.

Free Flight

Since today's air traffic follows fixed routes, controllers mainly monitor the route intersection points, and safely guide aircraft through them. This situation is depicted in Figure 1.1. Controllers receive aircraft data form the ATC radar and verbally communicate commands to pilots. The controllers mainly rely on their ability to identify troublesome flow patterns before they become hazardous and attempt to apply a standard set of procedures to resolve these situations. But this pattern may change with the introduction of "free flight" (RTCA 1995).

The free flight program is aimed at providing a safe and efficient flight operating capability while maximizing the airspace capacity. Under free flight, depicted in Figure 1.2, pilots are free to choose their path, altitude and speed, in real time. The only restrictions imposed on them are that separation is maintained, restricted air space (such as military space) is not entered, and airport capacity is not exceeded. Aircraft are to be equipped with modern navigation

Figure 1.1: The standard ATC scenario, as is practiced today.

equipment, e.g. GPS, and transmit their vehicle state and intention data, via data-link, to neighboring aircraft and to the ground controllers. This data is displayed on the flight-deck of each aircraft, and is used by the flight-crew to resolve conflicts by themselves without controller intervention. This concept, often referred to as *airborne separation*, poses new flight-deck related questions which will be discussed later. The ground based controllers must have the data of all aircraft in their controlled airspace displayed to them. Their task is to monitor the air traffic and to ensure that the limitations are not violated. Two spatial volumes are defined about each aircraft: the minimum separation volume, referred to as the *protected zone*, and a larger volume, enveloping the protected zone, referred to as the *alert zone*. Controllers intervene in case an aircraft's alert zone is violated, controlling the aircraft involved in the violation, until the situation is resolved.

Free flight has great potential for efficiency as each aircraft can choose its optimal flight conditions. To successfully monitor the complex traffic flow that free flight presents, controllers must, foremost, maintain a good situation awareness. Situation awareness is best defined as "the perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" (Endsley 1995). In ATC tasks, this amounts to an understanding of the current and future trajectories of all aircraft within a sector, some understanding of the representation of traffic about to flow into the sector, awareness of other relevant conditions such as weather, and an understanding of how all these factors affect the achievement of the ATC goals and constraints (NRC 1997).

For a controller to maintain situation awareness, a good mental model of where relevant events are likely to occur is necessary, so that selective attention can be allocated to sampling relevant parts of the display (Adams, Tenny and Pew 1995; Stein 1992). The predictive component of situation awareness is dependent on the spatial working memory's ability to "compute" likely trajectories based on current aircraft state, intended plans and aircraft dynamics. When multiple aircraft move in three-dimensions, and vary in air speed, the prediction task loads the controller's processing capacity to the utmost and limits the resolution with which the future state of traffic in the airspace can be visualized. Such a crowded, complex and heterogeneous airspace is envisioned under free flight. In order to support and maintain situation awareness, a 'good' display

Figure 1.2: The ATC scenario under free flight.

Figure 1.3: A TRACON controller station showing a radar scope, source (NRC 1997).

should aid the controller in anticipating future situations and support attention allocation to the right place at the right time (Sarter and Woods 1995).

Current Displays

The radar display, shown in Figure 1.3, constitutes the primary source of visual information for the controller. The radar detects a return signal from anything in the sky and displays it on the scope in a two-dimensional plan-view representation. Additional information is received from aircraft equipped with transponders; this information includes aircraft identity and possibly the barometric altitude, if the aircraft is quipped with a proper transponder.

When the ATC computer has a flight plan associated with the identity code transmitted by the aircraft, a data tag containing the aircraft's call sign, barometric altitude (if the aircraft is equipped) and ground speed will be displayed on the controller's radar display, as is depicted in Figure 1.4. A symbol at the end of the line

Figure 1.4: A data-tag.

extending from the data tag, represents the computer's estimation of the aircraft position. In addition to aircraft symbols, a terminal radar approach control (TRACON) display also contains landmarks, obstacles, feeder gates (which are "ports" through which traffic is received to the terminal approach area) and runways. Such a display is presented in Figure 1.5. The display is updated every 4 seconds, the radar's sweep time. In Figure 1.6 an En-Route control center display is presented. En-Route control centers are responsible for traffic between TRACON stations. This display differs from a TRACON display in that it does not contain raw radar returns and hence is more synthetic, and includes different auxiliary data such as radio beacons and navigational aids. On some En-Route displays, a set of equal-time cross-hatched lines behind the aircraft symbol represent its past trajectory.

The Pilot's Point-of-View

The pilot's objectives do not always coincide with the controller's objectives. Pilots are only concerned with one aircraft, whereas controllers are concerned with a population of aircraft that are spread over a large area. A pilot usually wants to fly the aircraft along the most efficient path, e.g. along a great circle, using tail winds, or at the most efficient altitude, e.g. "cruise climb".

Free flight will give the pilot the freedom to choose the optimal flight path, but the pilot must monitor the aircraft's surrounding airspace, a task which is new to the pilot. Free flight can be characterized as flying under "electronic visual flight rules" (Perry 1997), i.e. it confers on aircraft under instrument flight rules the same freedom of movement as those under visual flight rules today. This means that the pilot must be able to obtain from the electronic display at least the same level of situation awareness of the surrounding airspace as that obtained visually in the visible airspace. The challenge in designing a flight-deck display for free flight is that the display must present the information in a form the pilot can easily comprehend while being occupied with the primary task of

Figure 1.6: A traffic display in En-Route control centers, source (NRC 1997).

piloting the aircraft.

There currently exist no standards for free flight display. Nor do operational versions of this display exist. Several formats and technologies for a flight-deck display were examined in the literature (Azuma, Daily and Krozel 1996), including stereoscopic displays, alert zone and protected zone displays and various flight trajectory enhancements. O'Brien and Wickens (1997) examined two forms of a perspective format and two forms of a 2D format. Pilots were asked to steer clear of obstacles in the form of one intruding aircraft and a single weather front. No consideration was given to the selection of viewing parameters in the perspective displays. In these tests, the 2D displays exhibited superior performance. Van Gent *et al.* (1998) confronted a crew in a simulator with various free-flight scenarios containing multiple aircraft at high traffic loads. The crew was provided with a navigational plan-view display showing the surrounding traffic. The simulation results show that the pilots gazed at the navigational display 47% of the time. This figure reinforces the need for an effective flight-deck display in which the air traffic situation can be perceived quickly.

The quality of the solutions selected by pilots in resolving conflicts was
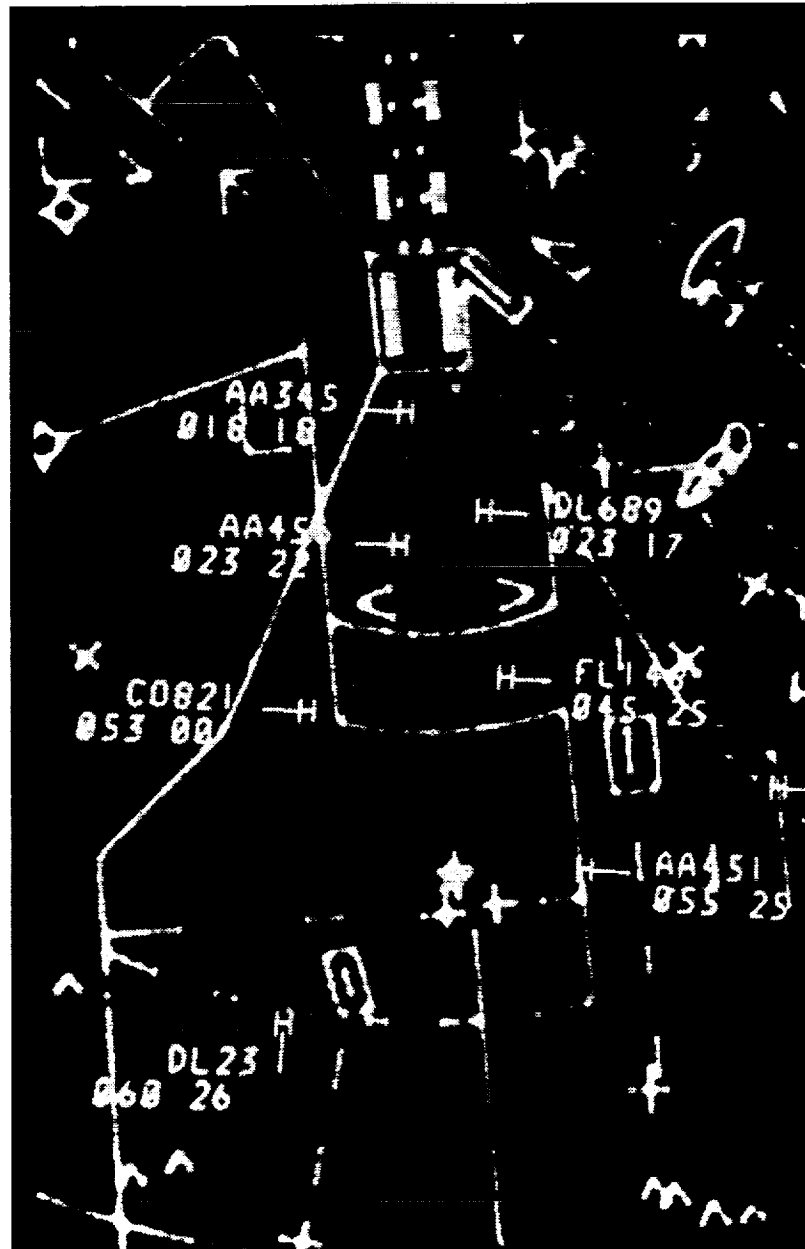
Figure 1.5: A plan-view ATC display at a TRACON center, source (NRC 1997).

examined by Ellis and McGreevy (1987). Multiple aircraft scenes were presented to pilots in a cockpit display of traffic information (CDTI), viewed perspectively from a viewpoint slightly behind and to the side of the pilot's own-ship. Pilots chose more maneuvers that involved both vertical and horizontal motion using this display than when viewing a plan-view display. This experiment demonstrated the advantage of the perspective display, which stems from its natural representation of the three-dimensional world.

1.2.2 The Impetus to Perspective Displays

Conventional ATC displays represent a plan-view of the air traffic area. The advantages are that the interpretation of the ground–referenced aircraft position is straight forward and all viewed areas are "weighted" with equal importance. In contrast, perspective displays represent the view of a 3-D scene from a well chosen vantage point, naturally integrating ground position and altitude of the aircraft.

Perspective displays are not without their pitfalls. Badly chosen viewing parameters might lead to disorientation of the operator, or impair the operator's spatial awareness. They may also result in clutter, exclusion of aircraft symbols from the view, or ambiguities in determining spatial positions. Perspective displays tend to bias the allocation of the operator's attention to nearby, centrally viewed areas on account of distant areas in the periphery.

Perspective displays have been attempted in many applications such as tele-operators (Chiruvolu, Hwang and Sheridan 1992), the presentation of a spatial flight-path (Grunwald 1984), spacecraft trajectory planning (Grunwald and Ellis 1988), air traffic control (Burnett and Barfield 1991) and data visualization (Wickens, Merwin and Lin 1994). The appeal of the perspective display format stems from its natural appearance when portraying the three-dimensional world. The appearance of perspective images is affected by the choice of viewing parameters and the choice of visual display enhancements, such as intensity depth cueing or the addition of reference lines. Selection of these factors has a profound impact on the observer's ability to comprehend the spatial layout, or to carry out the required task (Ellis *et al.* 1991; Kim *et al.* 1987). Often, a display designer needs to know how to present spatial data such that the user's understanding of certain geometrical features will be best. In interactive displays, the viewing parameter selection problem is commonly solved by handing over the

parameters selection task to the user by means of sliders or other manual controls for setting all or some of the viewing parameters; the initial viewing parameters are arbitrarily chosen. The main disadvantage of this approach is that the user is now burdened with an additional task, which may be critical in high workload situations (Jasek, Pioch and Zeltzer 1995). Additionally, to select the viewing parameters effectively, the user must have a good spatial understanding of the controlled parameters.

Several recent comparisons were made between conventional 2-D and 3-D display (O'Brien *et al.* 1997). Jasek *et al.* (1995) compared the ability of operators to predict collisions in three aircraft scenarios when using perspective displays and when using split plan-view/side-view 2-D displays. There was no systematic approach for setting the viewing parameters of the perspective displays. Generally, the 2-D displays had an advantage over the 3-D displays.

In an attempt to automatically select the "best" perspective viewing parameters for a teleoperator arm positioning and orientation task, Das *et al.* (1989) used a simple heuristic rule which was based on the arm's position relative to near-by objects. The operator could decide when to switch to new viewing parameters. The results of this experiment showed that the performance of well trained operators was best when they manually selected their viewpoints, more than with the automatically selected viewpoints or from a fixed viewpoint in space. The choice of viewpoints greatly affected the performance of the operators. This work demonstrated the importance of well chosen viewing parameters, but lacked a functional method for their selection. Such a method should consider human spatial perception mechanism and abilities.

1.2.3 Spatial Perception in Perspective Displays

Perception phenomena in perspective displays are well documented as early as La Gournerie (1859). When human observers reconstruct three-dimensional layouts from two-dimensional monoscopic displays certain judgement errors are inherent in the process. The Ames room (Ames Jr. 1925) strikingly demonstrates the human reconstruction mechanism by tricking it. Due to the special geometrical arrangement of the room, an observer falsely assumes he/she is viewing a rectangular room, with a rectangular tiled floor. This assumption is so dominant, that when a person walks across the room, the observer is forced to believe that the person is actually growing in size! Many works compared and

documented spatial judgements errors and their dependence on various perspective viewing parameters (Barfield and Rosenberg 1995; Ebert and Mac-Millan 1985; Ellis *et al.* 1991; Gregory and Ellis 1990; Kim *et al.* 1987; McGreevy and Ellis 1986).

One of the attempts to model the human reconstruction mechanism (Farber and Rosinski 1978) assumes the observer reconstructs a "virtual space" for the viewed image. According to this model, this virtual space is distorted when compared with the actual space since the observer falsely assumes he is viewing the image from the center of the perspective projection. Since the unique reconstruction of a three-dimensional layout from a single, monoscopic, two-dimensional perspective image is impossible, every model must assume the observer uses some additional knowledge which was not obtained from the image. The Farber and Rosinski model assumed the observer has knowledge of the depth of each point in the image and does not explain the source of this knowledge. Being deterministic, the model requires a basic imperfection in the viewing geometry and thus cannot explain judgement errors that occur when no such imperfection exist.

Additional models with similar characteristics are reviewed by Sedgwick (1991). These models differ from each other in the assumption of what prior knowledge the observer uses and what data is measured from the image, e.g. Lumsden (1983) assumed the observer is familiar with object size, Ellis *et al.* (1987) assumed the observer knows the true length of each line-of-sight and Purdy (1960) assumed texture gradients are being measured from the image. In spite of these differences, most of these models converge to the same results.

Different is the "familiarity cue" model (Grunwald, Ellis and Smith 1988). This model deals with "real-world" vision and not vision from a perspective image. The fundamental assumption of this model is that the observer uses *a priori* knowledge of certain attributes of the viewed objects, referred to as "familiarity cues". For example, an observer that views a tiled floor might assume the tiles are rectangular. The assumption that a priori knowledge is used in the reconstruction process is well established (Ames Jr. 1925; Wickens 1992). The observer perceives an object having the same familiar attributes and that would have generated lines-of-sight that best match the perceived lines-of-sight had it been the true object. Errors stem from noisy measurements of the angles of the perceived lines-of-sight. The statistics of the judgment errors were

obtained using a Monte-Carlo process.

The familiarity cue assumption is the basis for the spatial perception model developed in Chapter 2 which differs from the Grunwald *et al.* model in two aspects: (1) A reconstruction process from a two-dimensional image rather than from a real-world is modeled, (2) the error statistics are calculated directly, divesting the need for a lengthy Monte-Carlo simulation.

The perception of motion, both for ego-motion and exocentric object motion, has been the subject of separate research (Rogowitz and Allebach 1990). Most models for the spatial perception of motion agree that it is perceived from a sequence of static images (Sakaino and Sonehara 1996; Zacharias, Miao and Warren 1995), a notion which is enforced by the fixation movement behavior of the eye (Hacisalihzade, Stark and Allen 1992). In particular, Zacharias *et al.* (1985) successfully used Kalman filtering to model human perception of ego-motion by assuming it is estimated from noisy measurements of a spatially distributed optical flow-field. However, the model of Zacharias was only able to explain the perception of ego-motion, and did not suggest how to treat object motion.

1.3 Proposed Concept

To utilize the full potential of perspective displays, proper viewing parameters must be used, in conjunction with adequate graphical symbology, to present the "best" perspective image of the viewed geometry. The "best" image is the one in which the spatial information that is relevant to the operator is presented clearly and unambiguously, i.e. the operator can interpret the spatial information with the least errors, thus increasing the operator's ability to solve spatial problems. Since the airspace is constantly changing, the parameters of the "best" image should be changed in accordingly. Hence, an interactive display system is proposed that will monitor the airspace and update its display parameters in agreement with the evolving air traffic situation; such a system is hereafter referred to as an *active perspective display* system.

Several questions were addressed in order to design and implement an active perspective display:

1) *What needs operator attention?* What information is relevant and is extracted by the operator, and what is the relative priority of each information item?

2) *How can the quality of a perspective image be evaluated?* Can the human operator effectively comprehend the spatial information in the image and mentally reconstruct the 3D situation?

3) *How can the optimal perspective scene be created?* Given the evaluations of the second question and the priorities of the first, what is the best perspective setting?

Since the air traffic situation is dynamic, the answer to the first question will change with time as various parts of the airspace change their geometry, together with their relative importance. Knowing how to solve the second question, enables to rate a set of alternative perspective scenes and assign score to each one of them, which reflects their quality. The temporal nature of the first question leads to a conclusion that the score of each alternative is a function of the time. The third question can be restated as a time changing, multi-dimensional optimization problem, as the set of best scoring alternatives at each instant of time is sought.

The goal of the active display system is to improve the overall comprehensibly of the scene, and in an ATC implementation, to improve the performance of the controller. This is achieved by:

- Finding a compromise between the accuracy at which the different geometrical features can be estimated.

- Preventing clutter, or exclusion from view of any aircraft symbol and assuring all the relevant aircraft have a clear view.

- Directing the user's attention to most important features by allocating them to a "prime" area of the display.

- Relieving the user from the need to manually choose the viewing parameters and thus enabling the user to concentrate on the main task.
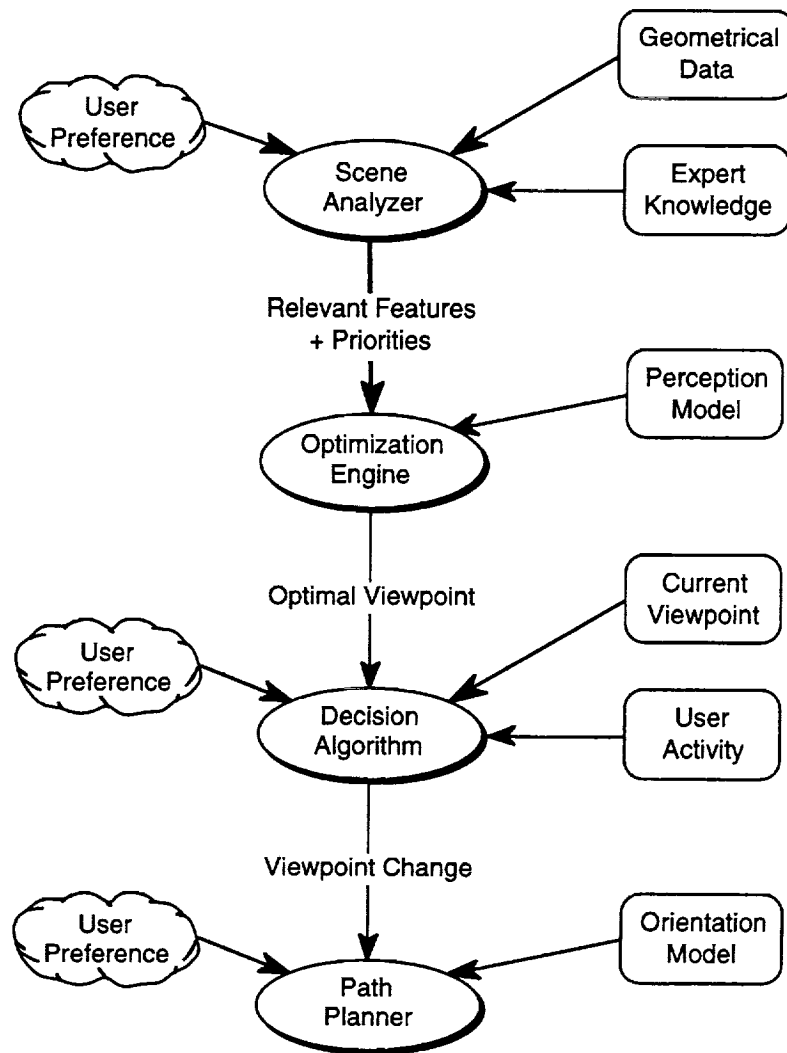
Figure 1.7: The structure and the data flow of the active display system.

1.4 Outline of the Active Perspective Display System

The active display system concept, as it was conceived and implemented, is composed of several components, depicted in Figure 1.7. A complete active display system for ATC should present the essential spatial information for routing air traffic flows, which consists of the aircraft positions, air routes, landmarks, radio beacons, geographical obstacles and possibly also weather-front information. Since this research project was focussed on evaluating the feasibility of the active display concept, rather than developing a fully fledged ATC display, a simplified system was implemented. This system, in which only geometrical aircraft data were displayed, was configured for the part-task experiments that were conducted.

1.4.1 Scene Analyzer

The scene analyzer examines the geometrical data to be displayed. Expert ATC knowledge, in the form of heuristic rules, is used to identify the relevant geometrical features and their relative priorities. The expert knowledge was extracted from studies which established the type of information controllers use when dealing with air traffic (Amaldi, Boudes and Cellier 1996; CourteixKherouf 1998).

1.4.2 Optimization Engine

The optimization engine finds the best viewing parameter setting for the information identified by the scene analyzer, such that a human observer can extract the relevant information from the perspective display with the least errors in perception. The optimization engine is the heart of the active display system and provides an answer to the latter two questions posed in Section 1.3, i.e. it dynamically finds the best image by rating alternative images in the process. The current implementation applies a genetic algorithm (Goldberg 1989) as it was found suitable for the multi-modal and time-evolving characteristics of the objective function (Shaviv and Grunwald 1997).

The focus of our research effort was on the development of the optimization engine, on the formulation of the static perception model and its extension to moving images, and on the development of the genetic algorithm used for the optimization.

1.4.3 Decision Algorithm

The decision algorithm continuously compares the optimal viewing parameters computed by the optimization engine with the current parameters in order to decide whether a change is appropriate. The decision is made by considering the score difference between the two parameter settings, the time elapsed since the last change and other factors such as the user activity and user preferences. The aim is to avoid agitation of the display and to reduce the impact of the change on the operator's situation awareness. A decision algorithm is required in particular when the viewed information is dynamic, as is the case in an ATC implementation.

The level of automation (Sheridan 1987), i.e. the mode of operation, can be set by user preferences. In a fully automatic mode, the system will initiate a

viewing parameter change whenever it finds it appropriate. This mode may disturb some operators, and is problematic if the operator is not gazing at the display all the time, e.g. in a flight-deck implementation. In a semi-automatic mode the system will request user authorization to perform a change, or wait a grace period enabling the operator to veto. In all modes, the user should be capable to manually override the change. The automatic changes must be reflected in the state of all manual viewing parameter controls.

1.4.4 Path Planner

The path planner determines the path, along which the viewing parameters are to be changed. Since discontinuities in the viewing parameters may impair the observers' spatial orientation, i.e. the observer's understanding of the viewpoint position in space and the viewing direction, the role of this component is to perform the requested change by sequencing gradual transitions, striving to maintain the highest level of spatial orientation during the change. The current implementation uses simple heuristic rules (Grunwald 1996). The gradual transitions are introduced by simulating as if the motions are controlled by a well-designed servo control system. The functions, realizing these motions were referred to as "slewing functions" and are extensively treated in the report of the previous year's research.

1.5 Organization of this report

As mentioned earlier, our research was focussed on the elements of the optimization engine. The heart of the optimization engine is the spatial perception model, presented in Chapter 2. Chapter 3 introduces genetic algorithms and describes the genetic algorithm that was used in the optimization. The active display system was evaluated in a part-task experiment, described in Chapter 4, demonstrating the basic feasibility of the concept and the validity of the image scoring method. Chapter 5 presents a spatial perception model for dynamic images, which is based on the static model of Chapter 2. This model enabled the second part-task experiment, described in Chapter 6, that evaluated the observer's ability to understand future situations in moving images and to extract dynamic quantities, such as aircraft velocities, thus providing a more complete examination of the observer's situation awareness. Chapter 7 concludes this report and provides final remarks.

2

Spatial Perception Model from Static Perspective Images

2.1 Introduction

Earlier work, in the framework of our cooperative efforts, has dealt with the development of mathematical models for the spatial perception of a human observer. These models analytically describe the process, by which the human observer reconstructs three-dimensional layouts from perceived two-dimensional images. The motivation for the development of these models was triggered by the need for understanding this process, when dealing with the development of spatial instruments and spatial displays. The design of a synthetic perspective display requires the choice

of the viewing parameters, i.e. viewpoint, viewing direction, field-of-view, as well as the geometrical representation used to portray each object. These choices have a profound impact on the observer's ability to comprehend the spatial layout, or to carry out the required task (Ellis *et al.* 1991; Kim *et al.* 1987). As has been mentioned in Section 1.2.3, the judgment errors made by human observers estimating spatial features from perspective images, strongly depend on the viewing parameters. The systematic judgement errors have been accounted for by several deterministic models known in literature, which explain the errors in a specific geometry, and hence are difficult to generalize or use as a design tool.

The model presented in this chapter deals with the random aspects (non-systematic) judgement errors. It is based on familiarity cues (Grunwald *et al.* 1988), and treats vision from computer generated, wire-frame, perspective images. An analytical expression for the statistics of the judgment errors is derived, making this model a suitable means to rate images in an on-line display application.

2.2 General Description of the Model

2.2.1 The geometry

The general geometry of the generation of a perspective image of a three-dimensional scene is described in Figure 2.1. The viewed objects consist of N vertices, of which the coordinates in the world coordinate system are $x_j^w = (x^w, y^w, z^w)^T; j = 1...N$. In case the image is produced by a camera, the objects are part of the real world, whereas in computer generated images, the objects are modeled in a database. The viewed objects consist of M edges. Each edge connects two vertices having the indices $[v_1(k), v_2(k)]; k = 1...M$, where each index $v_.(\cdot)$ is in the range $1...N$.

The image is generated by perspectively projecting the world coordinates x_j^w on the image plane P. This is done in two steps. First, the coordinates x_j^w are transformed into the coordinates x_j^s of the station point system. This system has its origin at the center of the perspective projection. The x-axis points towards a point-of-regard (POR) and defines the viewing direction. The POR is at a viewing distance D from the station point. Figure 2.1 shows the orientation of the station point system S with respect to the world system. The angles ψ and θ are the azimuth and inclination of the viewing direction, respectively. The angle

Figure 2.1: The geometry of a perspective image.

ϕ is the viewing twist, and is considered zero when the y-axis is parallel to the world x^w, y^w plane. The world-to-station point transformation is given by:

$$x_j^s = ([\phi][\theta][\psi])^T (x_j^w - r) + \begin{pmatrix} -D \\ 0 \\ 0 \end{pmatrix} \tag{2.1}$$

where r is the coordinate vector of the POR in the world system, $[\phi], [\theta]$ and $[\psi]$ are sequential rotation matrices about the x, y and z axes, respectively.

Second, the image plane vectors x_j^p are computed by intersecting the lines-of-sight, i.e. the lines connecting the station point with each vertex, with the image plane. The image plane is perpendicular to the viewing direction at a distance l from the station point, where the y^p and z^p axes are parallel to the y^s and z^s axis, respectively. The image plane coordinates of the perspective projection are then:

$$x_j^p = \frac{l}{h}\left(\frac{y_j^s}{x_j^s}, \frac{z_j^s}{x_j^s}\right)^T \tag{2.2}$$

Figure 2.2: The spatial perception process.

where these coordinates are given as fractions of half the screen height h. The parameters h and l determine the vertical field-of-view angle $\kappa = 2\tan^{-1}\left(\frac{l}{h}\right)$.

2.2.2 The Reconstruction Process

The basic assumption underlying the spatial perception model is that the observer has *a priori* knowledge of certain geometrical features of the viewed objects, referred to as familiarity cues. The familiarity cues essentially parametrize the object geometry, hence the world coordinates of the object vertices can be written as a function of the parameter vector c:

$$x_j^w = \mathcal{F}_j(c) \qquad\qquad (2.3)$$

where $\mathcal{F}_j(\bullet)$ are N familiarity cue functions, i.e. one function for each vertex, establishing a one-to-one mapping between c and x_j^w.

It is assumed that the observer deduces a correct set of familiarity cue functions from a first interpretation stage of the image. For example, if the observer is viewing the building of Figure 2.2, the observer might assume the building is composed of a rectangular parallelepiped and topped by a pyramid, the parameter vector c might consist of the width, height and depth of the parallelepiped and pyramid, and the familiarity cue functions will express the coordinates of

each vertex in terms of c using the knowledge that the parallelepiped is composed of perpendicular planar surfaces.

Since the geometry which produced the image is unknown to the observer, the geometry parameters that are required to reconstruct the image need to be estimated. It is assumed that the observer stipulates that the station point lies on an axis perpendicular to the image plane, passing through its center. This assumption is reasonable when the observer is aware of the image frame orientation (Cutting 1991). Hence, the parameters estimated by the observer can be written as the vector:

$$\chi = (\Psi, \Theta, \Phi, \kappa, D, r, c) \tag{2.4}$$

which is composed of the station point system orientation, the field-of-view angle, the viewing distance, the POR position and the familiarity cue parameter vector, respectively.

The vector χ is estimated by the observer by searching through all possible values for the one that would produce a layout, which is most likely to be the true one. The different layouts obtained by assigning a value to χ are hereafter referred to as *virtual layouts*. For each virtual layout, the observer constructs a *virtual image* by the perspective projection of Eqs. (2.1),(2.2). Since the virtual image is the perspective image that would have been seen had the virtual layout been the true layout, the model assumes that the observer selects the virtual layout with the virtual image which best matches the true image.

The errors in the perception process can be classified as deterministic errors and stochastic ones. Deterministic errors result in a bias on the observer's spatial judgments, and are caused by *misleading hypotheses* that the observer might make concerning the image geometry, or the displayed objects. For example, the observer's familiarization with an object similar to the viewed one, but of different size or proportions, may cause the observer to assume the viewed object is similar to the one the observer is familiar with. It has been reasoned that incorrect hypotheses due to size familiarization cause a higher frequency of rear-end collisions with small cars (Ebert *et al.* 1985). Incorrect hypothesis are modeled by assuming the observer has a preference for certain combinations in the values of χ.

Stochastic errors introduce a statistical distribution in the judgement errors

of different observers, as well as in the successive judgment trials of the same observer. This distribution in judgement errors is incorporated in the model by assuming it is caused by matching errors between the virtual images and the true image. This causes the search process for the best virtual layout to be imperfect such that, with a certain probability, an erroneous virtual layout may be selected. More precisely, it is assumed that for each virtual layout the observer calculates a *matching cost*, that rates how well its virtual image matches the true image. However, the observer has random errors in calculating the virtual layout's matching cost. Hence, a virtual image that does not have the lowest matching cost may be wrongfully selected. This is equivalent to a random selection process in which the probability of a virtual layout to be selected is inversely proportional to its matching cost.

2.3 Mathematical Formulation

2.3.1 The Matching Cost

The matching cost of a virtual image reflects the extent to which a virtual image differs from the true image. The matching cost, in case no misleading hypotheses exist, is defined as zero when the virtual image matches the true image, and it increases the less the virtual image resembles the true image.

It is assumed that the observer detects the edges in the image. Hence, the observer attempts to match the edges of the virtual image with the edges of the true image. It is assumed that the matching accuracy is proportional to the edge length. Accordingly, the matching cost of a virtual image is defined as the sum of the differences between each one of the virtual edges and its corresponding true edge, normalized by the length of the actual image edge. Assuming the correspondence problem is solved, i.e. the observer correctly established which edge of the virtual image and in the true image are corresponding, the matching cost may be written as:

$$\tilde{J}(\underline{\chi}) = \sum_{k=1}^{M} \frac{(\underline{e}_k - \underline{e}'_k)^T (\underline{e}_k - \underline{e}'_k)}{\left\| \underline{e}'_k \right\|} \tag{2.5}$$

where ε_k is the k^{th} edge vector in the 2-D virtual image and is equal to:

$$\varepsilon_k = x^p_{v_1(k)} - x^p_{v_2(k)},\tag{2.6}$$

and ε'_k denotes the corresponding actual edge seen in the true image. The various 2-D virtual image vertices $x^p_{(\bullet)}$ are computed from χ by first using the familiarity cue relations, Eq. (2.3), to obtain 3-D virtual world coordinates, and then calculating their perspective projection, Eqs. (2.1),(2.2).

By defining $\varepsilon_j = x^p_j - x'_j$ as the difference between the positions of the j^{th} virtual vertex projection and its corresponding vertex in the true image, and by using Eq. (2.6), Eq. (2.5) can be written as:

$$\tilde{J}(\chi) = \sum_{k=1}^{M} \frac{\left\| \varepsilon_{v_1(k)} \right\| + \left\| \varepsilon_{v_2(k)} \right\| - 2\varepsilon^T_{v_1(k)}\varepsilon_{v_2(k)}}{\left\| \varepsilon^a_k \right\|}.\tag{2.7}$$

Eq. (2.7) reflects the matching goal of an observer who does not have misleading hypotheses about the viewed geometry. The misleading hypotheses are expressed as a preference towards the hypothesized combination of values of χ, and are modeled by adding a term to the matching cost of Eq. (2.7) that increases the matching cost the further the virtual layout's parameters are from the hypothesized values:

$$J = \tilde{J} + \delta^T W \delta\tag{2.8}$$

where J is the corrected matching cost and δ expresses the combination of values of χ towards which there is a preferences. It is assumed δ can be written as a linear combination of χ:

$$\delta = \underline{B} + Q\chi.\tag{2.9}$$

2.3.2 The Error Model

The Bias

As stated earlier, the observer attempts to find the virtual layout which yields the minimum matching cost. Let χ_0 be the true parameters and χ_m the one that

minimizes the cost of Eq. (2.8). The difference $\chi_m - \chi_0$ constitutes the bias. Assuming the bias is small, we can expand the gradient of J to a Taylor series:

$$\frac{\partial J}{\partial \chi} = \frac{\partial J}{\partial \chi}\bigg|_{\chi_0} + (\chi_m - \chi_0)\frac{\partial^2 J}{\partial \chi^2}\bigg|_{\chi_0} + \mathcal{O}(\chi_m - \chi_0)^2. \tag{2.10}$$

The minimum is obtained by equating Eq. (2.10) to zero, and discarding the higher order terms, yielding the following approximation:

$$\chi_m - \chi_0 = \left[\frac{\partial J}{\partial \chi}\left(\frac{\partial^2 J}{\partial \chi^2}\right)^{-1}\right]_{\chi_0}. \tag{2.11}$$

In the design of perspective displays, it is desirable that the average judgment error, i.e. the bias, be small. To this end, the left hand side of Eq. (2.11) must be minimized.

The Variance

In order to obtain the variance of the errors, a shape for the probability density function of the matching cost J has to be assumed. Using the probability density function of J, the variance of χ can be calculated.

As was previously stated, a virtual layout has a probability to be selected which is proportional to its matching cost. If $J_d = J - J_m$ is defined as the difference between the virtual layout's matching cost and the minimal matching cost, then the probability that the selected layout has a matching cost difference less than a certain J_d increases with J_d, this is characteristic of an exponential probability density function. Hence, the probability density function that shall be assumed for J_d is

$$f_j(J_d) = \begin{cases} \frac{1}{\lambda}e^{-J_d/\lambda} & J_d > 0 \\ 0 & J_d < 0 \end{cases}, \tag{2.12}$$

where λ is a constant parameter of the distribution. If absolute judgement errors are required, λ could be evaluated from experiments with human observers. For the purpose of comparing the quality of several images, the exact value of λ is irrelevant.

To calculate the variance of the error in the estimation of the i^{th} element of

χ, the distribution of the i^{th} element of $\tilde{\chi} = \chi - \chi_m$, denoted as $f_{\chi_i}(\tilde{\chi}_i)$, has to be calculated. It is assumed that the observer is only attempting to estimate χ_i while all other elements of χ are known perfectly. This assumption is equivalent to assuming that errors due cross relations between the parameters may be neglected, in view of small estimation errors.

By definition, the variance $\sigma^2_{\chi_i}$ is

$$\sigma^2_{\chi_i} = \int\limits_{-\infty}^{\infty} \tilde{\chi}_i^2 f_{\chi_i}(\tilde{\chi}_i) d\chi_i \tag{2.13}$$

where $f_{\chi_i}(\tilde{\chi}_i)$ is the density function of $\tilde{\chi}_i$. Since $J_d = J(\chi) - J_m$, f_{χ_i} relates to f_j according to (Papoulis 1984):

$$f_{\chi_i} d\chi = f_j(J_d(\chi)) \left| \frac{\partial J}{\partial \chi_i} \right| d\chi_i . \tag{2.14}$$

By substituting Eq. (2.14) into Eq. (2.13), and using Eq. (2.12) we get:

$$\sigma^2_{\chi_i} = \frac{1}{\lambda} \int\limits_{-\infty}^{\infty} \tilde{\chi}_i^2 \left| \frac{\partial J}{\partial \chi_i} \right| e^{-\frac{J_d}{\lambda}} d\chi_i . \tag{2.15}$$

Eq. (2.15) can easily be evaluated by approximating the integrand as a Taylor series expanded around χ_m and neglecting higher order terms. Since J_d is at a minimum at χ_m, its Taylor expansion reduces to:

$$J_d \cong \tilde{\chi}_i^2 \frac{\partial^2 J}{\partial \chi_i^2} \Bigg|_{\chi_m} . \tag{2.16}$$

The Taylor expansion of $\left| \frac{\partial J}{\partial \chi_i} \right|$ reduces to:

$$\left| \frac{\partial J}{\partial \chi_i} \right| \cong |\tilde{\chi}| \frac{\partial^2 J}{\partial \chi_i^2} \Bigg|_{\chi_m} . \tag{2.17}$$

Substituting Eqs. (2.16) and (2.17) into Eq. (2.15) and integrating, yields:

$$\sigma^2_{\chi_i} = \lambda \left(\frac{\partial^2 J}{\partial \chi_i^2} \Bigg|_{\chi_m} \right)^{-1} . \tag{2.18}$$

It is interesting to note that a relation similar to Eq. (2.18) can also be obtained using a numerical minimization analogy. Assume the matching process is carried out in a fashion similar to a numerical minimum search algorithm. The search is discontinued once a certain precision has been reached, i.e. J is at a distance ΔJ from its minimal value. Since J is at a minimum at χ_m, using Taylor expansion, the distance to the true minimum $\Delta\chi$, can be shown to be $\Delta\chi^2 = \Delta J(\partial^2 J/\partial\chi^2)^{-1}$. This relation is reminiscent of Eq. (2.18) if one substitutes $\sigma_\chi^2 \rightarrow \Delta\chi^2, \lambda \rightarrow \Delta J$ by analogy. While this analogy is not mathematically rigorous, it provides some insight into the result of Eq. (2.18).

Eq. (2.18) states that the variance of the estimation error is proportional to the inverse of the curvature of the matching cost surface at the estimated parameter value. Using the numerical minimization analogy, this can be compared with finding the minimum of a sharp sink, i.e. a large curvature, versus that of a shallow depression, i.e. small curvature. In the first case, the minimum can be easily found, meaning the image is accurate and can only be interpreted with a unique spatial geometry. In the latter case, the exact location of the minimum is harder to find meaning many possible geometrical layouts can yield similar images.

2.4 Analytical Model Evaluation

To test the new model, a comparison was made with results obtained by using the Grunwald *et al.* model. The Grunwald *et al.* model was also tested against experimental human observer performance data. The test scene that was used is similar to that used in Grunwald *et al.* (1988) and is essentially an abstraction of scenes encountered in certain types of Air Traffic Control displays (McGreevy *et al.* 1986).

The scene consists of two small spheres which rest on poles extending from a ground grid. One sphere is 3 meters above the ground, and the other is 4 meters above the ground. The ground distance between the spheres is 4 meters. The size of the ground grid is 10 meters. It was assumed that the familiarity cues of the observer are that the poles are perpendicular to the ground plane, and the poles bottom edge lies in the ground plane.

The objectives of the observer are to judge the azimuth of one cube with respect to the other, the altitude difference between the cubes and the ground distance between them. Predictions to the expected judgement errors were pro-
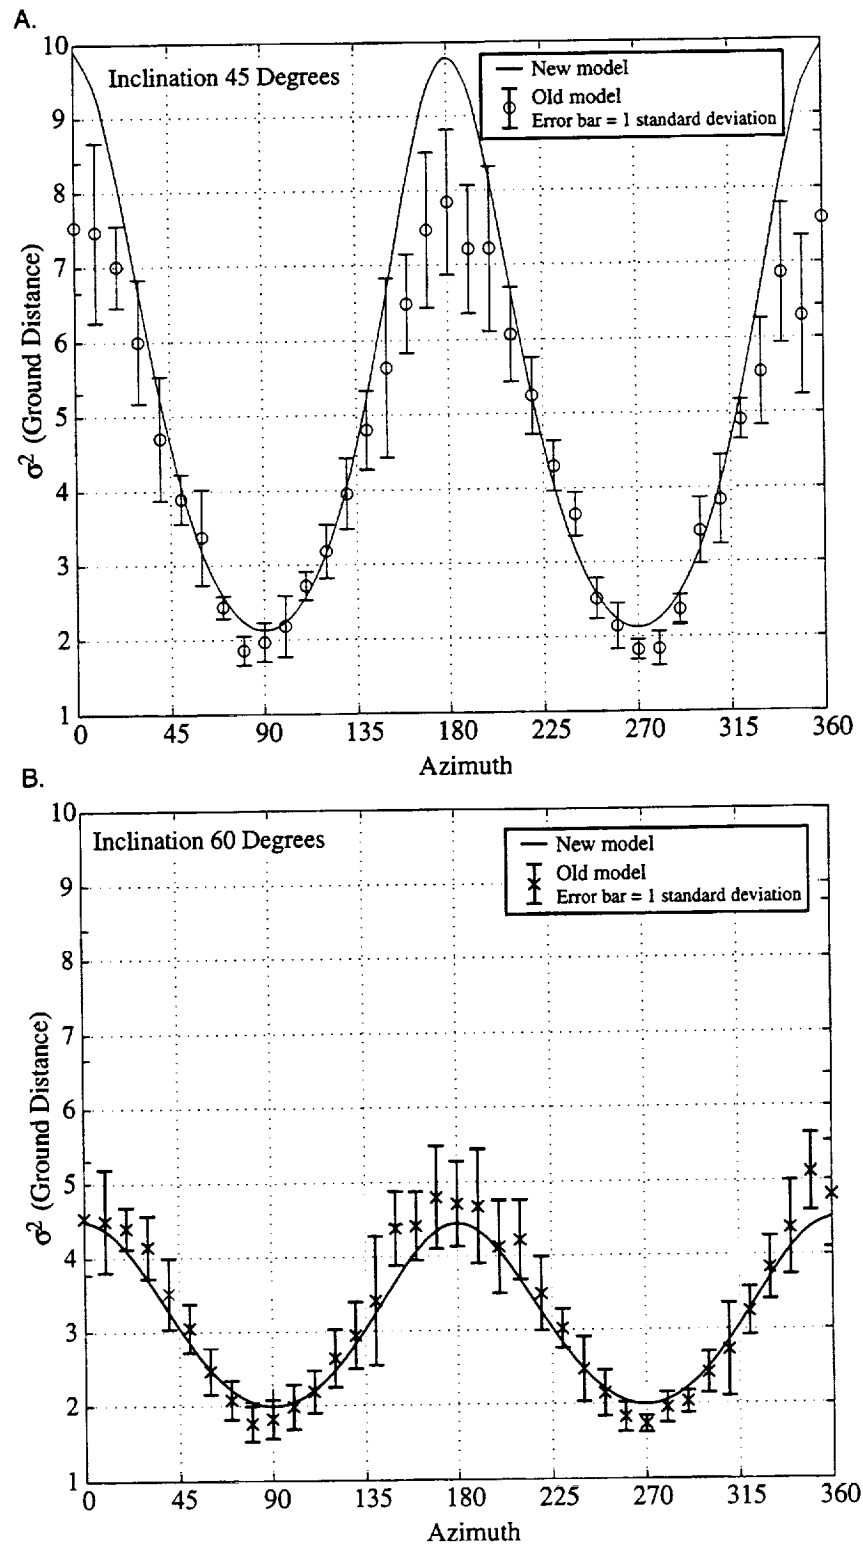
Figure 2.3: The error predictions of the new model vs. the Grunwald *et al.* model (old) as a function of the viewing azimuth angle.

duced using both models.

The comparison revealed that the two models yield close predictions. For example, Figure 2.3 shows the results for the variance of the ground separation judgement error at different viewing directions. The predictions of the Grunwald et al. model are represented as circles in part A of the figure when viewing at a 45° inclination, and as crosses in part B when viewing at a 60° inclination. The errors are plotted against the viewing azimuth angle. The error bars represent one σ of the prediction distribution as obtained from the Monte-Carlo run. The solid lines are the predictions obtained by the current model for the same viewing inclination.

The new model also reproduced the judgement error behavior observed in other works. For example, Figure 2.4 shows an expected judgement error behavior that is similar to the one observed by McGreevy and Ellis (McGreevy *et al.* 1986), which was referred to as a "braided sine wave". To explain this behavior, McGreevy and Ellis had to assume that the observer has knowledge of the length of the lines-of-sight emanating from the image. In the Grunwald et. al model, this effect was called the "telephoto lens" effect, referring to the fact that the observer has a wrong assumption about the zoom angle of the camera, used to generate the scene. It is well-know, that when observers view scenes through a telephoto lens, the depth of the object is underestimated. Also, in the new model, this behavior was reproduced, by assuming that the observer has a false notion of the field-of-view angle, i.e. the observer assumes a priori an incorrect field-of-view angle. In Figure 2.4, an azimuth angle of 45 degrees is overestimated (a positive error) for the narrow-angle field-of-view of 40 degrees and underestimated for the wide-angle field-of-view of 70 degrees.

2.5 Summary

The model presented offers a new and straightforward explanation to the spatial perception process, which is general and geometry independent. The model is based on the Grunwald *et al.* model which was formulated for real-world scenes; considered only vertices in the scene, and required a computation-intensive Monte-Carlo process to evaluate the variance of the judgement errors. The current model treats perspective images; considers both points and edges of the image, and provides a closed-form analytical solution. This makes the model suitable for an on-line application, such as the active display system. The closed-

Figure 2.4: The bias in the judgement error of the relative azimuth angle between the two spheres, as a function of the viewing azimuth angle, for different field-of-view angles.

form solution which was obtained has a fundamental implication; the variance of each geometrical feature is relative to the curvature of the cost function the observer uses to match the image. Using this result, the model can easily be extended to include other visual attributes like shaded light, lighting or textures, by modifying the matching function.

3 The Optimization Engine

3.1 Introduction

As mentioned in Chapter 1, the optimization engine, being at the heart of the active display system, is responsible for finding the optimal projection parameter setting. This optimal setting is found by optimizing an objective function, which rates the perspective image according to the expected judgment errors, which an observer would make when interpreting the image. These expected judgment errors, in turn, are predicted by the spatial perception model. This objective function, which is described in Appendix A, yields a complex, multi-modal optimization surface. Furthermore, since the ATC scene is inherently time-varying, this surface changes its shape with time. Finding the global optimum of a complex multi-modal and time-varying optimization surface is a non-trivial task!

We have considered a number of "classical" optimization techniques, such as gradient methods and second-order schemes. These schemes have the advan-

tage of quickly and accurately converging to the global optimum, provided the "initial guess" is in close vicinity of this global optimum. Since in our problem, is it is far from being guaranteed that the initial guess will be in the vicinity of the optimum, these classical optimization schemes have the tendency to "lock" on local minima and converge to incorrect solutions. Thus, these schemes lack robustness and the ability to deal with the time-varying nature of our problem.

To solve our multi-modal, time-varying optimization problem, we have chosen a genetic algorithm. These algorithms, with their close analogy to nature, are robust and flexible, and have the unique ability to adapt themselves naturally to changes, evolving in time. A general review of genetic algorithms is presented in Section 3.2, followed by the special adaptations and techniques that were implemented for the active display problem.

3.2 Synopsis of Genetic Algorithms

A brief overview of the workings of genetic algorithms (GAs) is presented. The purpose of this section is to provide the necessary background to comprehend the work, which is presented subsequently. For a more detailed description of genetic algorithms, the reader is referred to Goldberg (1989).

3.2.1 Mechanism of a Simple GA

The GA mechanism is based on a metaphor to genetic evolution, and will be explained using an illustrated example. Suppose one wants to find the maximum of the 2-D function:

$$f(x, y) = |x - y| \tag{3.1}$$

in the region:

$$\begin{aligned} 0 < x < 15 \\ 0 < y < 15 \end{aligned} \tag{3.2}$$

As a first step, possible solutions which are vectors in 2-D with components in the range specified by Eq. (3.2) have to be coded to a string representation. This can be done by coding the components of the solution vectors as binary strings, and concatenate these strings one after the other forming one long binary string, as illustrated in Figure 3.1. The string representation of the solu-

tion is called a chromosome. Each letter in the chromosome is called a gene and the possible values a gene can have are called alleles.

$$01001110$$
$$x=4 \ y=14$$

Figure 3.1: A Chromosome representing a possible solution.

GAs work with a population of solutions. Initially, the population is filled with random members; Figure 3.2 shows an initial population with four members. Each member in the population is assigned a fitness which reflects how the solution it represents compares with the rest of the population. In Figure 3.2 the fitness of each member was calculated by normalizing f, the function value of the solution from Eq. (3.1) by the average function value of the population $\bar{f}$.

	x,y	Chrom.	f	fitness
1.	4,14	01001110	10	1.48
2.	13,14	11011110	1	0.15
3.	0,11	00001011	11	1.63
4.	11,6	10110110	5	0.74

$$\bar{f}=6.75$$

Figure 3.2: The initial population and its fitness assignment according to the function value of each member f and the average function value $\bar{f}$.

As in nature, GAs use stochastic rules to spawn the next generation. First, members are randomly selected to pass to the next generation with a probability relative to their fitness. This process can be

Figure 3.3: The selection process.

thought of as spinning a weighted roulette wheel, as is illustrated in Figure 3.3. Selection applies pressure on the population towards convergence to a single solution.

Second, a preset proportion of members is randomly chosen to mate. Mating produces new members from existing ones by crossing-over the chromosomes of the two parents at a random cross-over point, illus-

Figure 3.4: A crossover.

trated in Figure 3.4. Cross-over applies pressure on the population towards divergence, by uniformly distributing alleles in the population.

Third, each bit in the population is randomly reversed with a low mutation probability. The probability of mutation is preset and is one of the parameters that has to be decided when implementing a GA, along with the population size and cross-over proportion. Selection eliminates alleles from the population which do not contribute to a higher fitness, the role of mutation is to preserve alleles from being lost by reinserting them to the population so their usefulness can be retested. Setting a high rate of mutation decreases the convergence rate of the GA, and decreases the probability it will converge to a local maximum.

Finally, a new generation is created. Analysis reveals (Holland 1975) the new population average function is guaranteed to be statistically higher than the average function value of the previous generation. An insight to the GA mechanism can be gained from the "two-armed bandit" problem. A finite number of coins is to be invested in a slot machine with two arms, each giving an unknown average payoff, with the intent of maximizing ones profit. A plausible strategy is to spend a proportion of coins to measure the payoff of each arm; then, play the rest of the coins on the arm that exhibited the higher payoff. Spending more coins to measure the payoff, decreases the probability the rest of the coins will be played on the wrong hand, but also decreases the number of coins for reaping the profit. Holland (1975) showed that GAs solve the two-armed bandit problem (more generally, a k-armed bandit problem) in an optimal manner, i.e. GAs will allocate the population members in an optimal compromise between exploring new knowledge and exploiting gained knowledge. This trait makes GAs less susceptible to local maxima. While most of the population members will concentrate on the region of the solution space exhibiting the highest objective value, some effort will be made to probe additional regions of the solution space.

Since their inception, GAs have been used in a variety of engineering applications, such as recursive adaptive filtering (Etter, Hicks and Cho 1982), structural optimization (Goldberg and Samtani 1986) or aircraft landing strut weight optimization (Minga 1986).

3.2.2 Advanced GA Techniques

Selection Revisited

The fitness in Figure 3.2 was calculated directly from the function value of each member. This makes selection pressure dependent on the scaling of the objective function, e.g. the fitness of members in the population will be different if their function values are between 0–10 or between 100–110. An alternative is to assign fitness to members according to their rank in the population and not directly from their function value (Whitley 1989).

Instead of selecting all the members that pass to the new generation according to their fitness, as was done in Figure 3.3, a preset proportion can be selected at random, this proportion is referred to as a *generation gap* (De Jong and Sarma 1992). A generation gap usually increases the GAs exploratory power by letting offsprings interact with their parents.

Cross-Over Revisited

A different cross-over scheme than the one described in Figure 3.4 was used in the optimization engine, called a two-point crossover. Instead of one crossover point, two crossover points are chosen, and the string part of the two chromosomes between these two crossover points is exchanged. This is identical to considering the chromosome as a ring instead of a linear string. Two-point crossover has better characteristics than single-point crossover or other multiple-point crossover schemes when optimizing non-linear functions, since it facilitates the GA in finding relations between the parameters of the solution (Eshelman, Caruna and Schaffer 1989; Spears and De Jong 1991).

Dominance and Diploidy

The solutions in Figure 3.2 were coded as a single chromosome string. In higher life forms, chromosomes contain two sets of genes, rather than just one; known as diploid chromosomes as opposed to a haploid chromosome. Redundant allele values are resolved through a mechanism called *dominance*; at each locus one allele takes precedence over the other. By deciding that a 0 is dominant over a 1, and by introducing a new 1, marked as

$$\begin{array}{c} 11\bar{0}11 \\ \bar{1}0001 \end{array} \rightarrow 10011$$

Figure 3.5: A diploid chromosome is resolved using dominance.

$\bar{1}$, which is dominant over a 0, a diploid chromosome may be coded in a GA (Hollstien 1971), as is demonstrated in Figure 3.5.

Using diploid structured chromosomes enables the GA to maintain unused alleles in an inactive state, and reactivate them once they become useful again. Thus, diploidy offers an additional mechanism to preserve alleles that is more robust, and less disruptive to GA convergence, than mutation. It can be practical for optimizing functions with changing values, and hence require the GA to maintain its probing power and continuously try new solutions.

Most of the genetic literature in nature tends to concentrate on diploid chromosomes, while virtually all work on GAs concentrates on haploid chromosomes. This is primarily for simplicity since diploidy involves an overhead in a GA. Most studies of dominance considered it as a method to store structures in abeyance in periodic functions (Goldberg and Smith 1987), in spite its apparent suitability for time-evolving functions.

Niche Formation

Genetic algorithms can be configured to locate several local maxima. A member group of the population that is located on a distinct function peak is said to be forming a niche. To create niches, Holland (1975) suggested, by metaphor to nature, to modify the allocation of fitness such that all members that *share* the same peak in the objective function divide among themselves the function value of the peak. Thus, a member can have a high fitness if it rests on a peak with a low function value if that peak contains few other members. The difficulty is in determining whether two members share the same peak or are on two nearby peaks.

Sharing

A practical scheme that directly uses the sharing metaphor to induce niches is detailed in Goldberg and Richardson (1987). In this scheme, a *sharing function* is defined to determine the neighborhood and degree of sharing for each member in the population. The sharing function is such that it is equal unity if the two members are identical, and monotonously decreases to zero the more the mem-

bers are different from each other. The sharing function between members i and j, used by Goldberg and Richardson is:

$$s(i, j) = \begin{cases} 1 - \left(\frac{d(i, j)}{\sigma}\right)^{\alpha} & d(i, j) < \sigma \\ 0 & d(i, j) > \sigma \end{cases} \tag{3.3}$$

where $d(i, j)$ is the *distance* between the members which is defined as zero if the solutions are identical, and non-negative otherwise, α is a constant and σ is a preset distance threshold.

The fitness f_i of member i is derived from the unshared fitness $f(x_i)$ by:

$$f_i = \frac{f(x_i)}{\displaystyle\sum_{j=1}^{N} s(i, j)} \tag{3.4}$$

Thus, the more members there are in the vicinity of the i^{th} member, the higher the sum of its sharing function, and the lower is its fitness.

Two options for defining a distance are suggested by Goldberg *et al.* (1987), and are used throughout GA literature. *Phenotype distance* is defined as the Euclidean distance between the members' parameter vectors $d(i, j) = \|x_i - x_j\|$. *Genotype distance* is defined as the Hamming[†] distance between the two members' chromosome strings.

The workings of the sharing principle depends on the parameter σ which is the minimum distance between function peaks that enables their location. Setting σ requires knowledge of the function topology, and is further discussed in Appendix B.

Deterministic Crowding A different approach achieves niche formation through an intervention in the evolution cycle of the GA, named *deterministic crowding* (Mahfoud 1992). In this method a new generation is created by first selecting *at random* a pair of members from the population. Second, the selected members are mated and the offsprings are considered as candidates for reinsertion to the population in place of the parents. Each offspring is compared with the parent whose distance to it is smaller, and replaces it if its fitness is higher. Deterministic crowding was shown to be an

---

[†] The Hamming distance is defined as the number of different letters between two strings.

effective method for maintaining multiple peaks and ensuring population diversity, since it basically "kills" crowded member (see Appendix B for further discussion). Deterministic crowding rids of the necessity to preset a sharing distance σ, and hence is more independent of function topology. However, a distance still has to be defined to determine which offspring is candidate to replace which parent, Mahfoud (1992) used phenotype distance for this purpose.

3.3 Why Use GAs?

Several factors made GAs an attractive choice for the active display optimization problem. The objective function, which is the quality of the image as a function of the viewing parameters, is highly non-linear, it may contain several modes and may have discontinuities. GAs are fit to handle such objective functions. Additionally, GAs do not require derivatives which are difficult to calculate for the objective function.

The content of the airspace is continuously evolving in time, and the optimization engine is required to find optimal viewing parameters on-line, keeping up with the pace the airspace changes. By exploiting the GA's evolutionary nature and configuring it to track over time the objective function, rather than initiating a new search for the optimum each time events change, changes in the objective function can be detected quickly. Such a GA will strategically place its population members in an optimal compromise between seeking the current optimum and anticipating future optima. In this manner the GA will use past knowledge on the structure of the objective function when searching for a new optimum.

Moreover, it can be ascertained that each maximum in the objective function is related to a pair or a group of aircraft (the spatial situation of which is considered as deserving operator attention). Aircraft position changes continuously in the airspace, thus any objective function shape that relates to aircraft position changes continuously in time. The conclusion is that a new global maximum will most likely arise from a former local maximum that grew "steeper", while the former global maximum became "shallow". An optimization algorithm that keeps track of the local maxima can instantly detect such an event.

A reason against using GAs is their slow convergence rate compared with gradient based methods, typically it is exponential with the number of function evaluations, referred to as the number of trials. This handicap is circumvented

Figure 3.6: The perspective parameters for ground based ATC.

by configuring the GA to continuously monitor the airspace, thus having only to find the *changes* in the function. The slow convergence rate has an effect only on the initial "start-up time" of the algorithm.

3.4 GAs for the Active Display Optimization

3.4.1 Encoding the Parameter Set

In order to represent possible solutions in the GA, the parameters of a perspective projection need to be encoded as a chromosome string. The parameters chosen to define a perspective projection, see Figure 3.6, are: The coordinates of the point-of-regard (POR) P_x, P_y; the azimuth angle ψ and the inclination angle θ of the viewing direction; the viewing distance D, i.e. the distance between the center of the perspective projection and the POR; and the field-of-view (FOV) angle κ.

The parameters were encoded as a Gray[‡] binary string, each was encoded in a given range and precision, according to the expected parameter range and the required solution accuracy, these are summarized in Table 3.1. The parameters

---

[‡] Gray encoding are better suited for GA optimization than binary encoding, since in Gray encoding adjacent integers always differ by a single bit. Gray codes are generated by XOR-ing each digit in the binary code of a number with the digit one place left of it. Thus the Gray code of the integer 11 (binary 1011) is 1110, and of the integer 12 (binary 1100) is 1010.

Parameter	Range	Precision
Viewing azimuth ψ	0 - 360 Deg	0.25 Deg
Viewing inclination θ	0 - 90 Deg	0.25 Deg
P_x, P_y of POR	-150 - 150 NM	2 NM
Viewing distance D	0 - 150 NM	0.5 NM
Field-of-view κ	10 - 80 Deg	0.1 Deg

Table 3.1: The encoding of the solution parameters

were concatenated to form one string in the order: Ψ, Θ, P_x, P_y, D, κ yielding a 45 bit string. This order was chosen to place parameters which were assumed to have a related effect on the objective function closer together in the string, thus facilitating the GA in finding any such relation.

3.4.2 Constructing the GA

GA-1	
Population size	300
Selection	Deterministic with stochastic reminder. 0.3 Generation gap.
Mutation rate	1e-3
Crossover	Two-point crossover at a rate of 0.8.
GA-2	
Population size	50
Selection	Deterministic with stochastic reminder. No generation gap.
Mutation rate	1e-6
Crossover	Two-point crossover at a rate of 0.6.

Table 3.2: The parameters of GA-1 and GA-2.

As stated, the goal is to create a GA that does not converge but continuously explores for new evolving maxima. To study the effect of the various GA components, the GA is constructed in three steps. First, a simple GA, similar to the one presented in Section 3.2.1 is configured and evaluated. Second, diploidy is introduced and compared with the performance of the haploid GA. Third, an

appropriate niche formation technique for the objective function is developed and employed to form the final GA.

The Simple GA

A simple GA is configured for a high exploratory capability and a delayed convergence. This is achieved by a large population size, a high cross-over rate, and a moderate generation gap. This GA is referred to as GA-1. For comparison, a second GA was configured for fast convergence, referred to as GA-2. Table 3.2 summarizes the parameters of GA-1 and GA-2.

To test the GAs performance on a realistic, dynamic scenario, an airspace simulation was extracted from true recordings of air traffic data from the Denver ATC Center (Dorighi 1996), representing a moderately busy ATC center. The simulation included 20 aircraft which were all in an organized traffic flow, except for five "rogue" aircraft that violated a separation criteria sometime during the first four minutes of simulation. This scene is of similar characteristics to the scenes used in the experiment described in Chapter 6.

Two GA performance measures are shown in Figures 3.7,3.8; they are the result of a 4 minute simulation. The first performance measure is the objective function value of the best member in the population. This is the best solution the GA holds, and is the solution that is passed over to the decision algorithm to consider a possible parameter change. The second measure is the off-line performance which is defined as (Goldberg 1989):

$$P(T) = \frac{1}{T}\sum_{1}^{T}\max\{f(1)...f(T)\} \qquad (3.5)$$

where T is the number of trials, each trial being one objective function evaluation, and $f(\cdot)$ is the value of the objective function at each trial. Thus the off-line performance is a running average of the best population member across trials. An increase in the off-line performance indicates the GA is continuously improving the best member. The x-axis in Figures 3.7,3.8 is the number of trials and corresponds to the 4 minute simulation time. The different number of trials between Figure 3.7 and Figure 3.8 is due to a different overhead in managing the two different sized GAs, and other random factors.

To determine the GA's performance compared with the function's true maximum, the true maximum was calculated by "freezing" the airspace at different

Figure 3.7: The performance measures of GA-1.

Figure 3.8: The performance measures of GA-2.

Figure 3.9: The convergence map of GA-1, the percentage of lost and converged alleles.


Figure 3.10: The Convergence map of GA-2, the percentage of lost and converged alleles.

times, in intervals of 30 seconds, and letting a GA optimize this statically held airspace until it converged. These points, marked as triangles in Figures 3.7,3.8, show that both GA-1 and GA-2 do not reach the true maximum of the dynamic objective function. Initially GA-2 converges faster and to a higher score than GA-1, but as the objective function evolves and shifts in the solution space, GA-2 loses performance abruptly, while GA-1 sustains a constant performance level, as is evident by the level off-line performance curve.

An explanation can be obtained from the convergence maps of the two GAs, Figures 3.9,3.10. These show the percentage of lost and converged alleles for the two GAs. An allele is considered lost if it is not present in all the population, e.g. if in a certain bit position all the population has the value of zero, the "one" allele at that position is lost. Converged alleles appear in the population less than a certain percentage. As can be seen, GA-1 maintains a low level of lost alleles compared with GA-2 that almost immediately converges most of its population. One can conclude that the slow rate of convergence of GA-1 can be attributed to it allocating a substantial "effort" to maintaining population diversity for further exploration. GA-2, on the other hand, concentrates all its efforts in achieving a higher score causing it to lose its population diversity and hence its probing powers.

Introducing Diploid Chromosomes

Figure 3.9 indicates that the convergence of GA-1 was not stopped, but only slowed down. Longer simulation times show that eventually GA-1 converges its population, in particular at times of "low activity" when the objective function does not change much in time. By increasing the rate of mutation in GA-1, convergence may be further reduced, at the price of performance. Instead, diploidy can be used.

Table 3.3 summarizes the parameters of GA-3, a diploid GA which similar parameter values as GA-2, except for a larger population. The performance measures of GA-3 on the simulation run are shown in Figure 3.11. It is evident, from Figure 3.11, that this GA converges initially as fast as GA-2, but unlike GA-2, is able to continue tracking the objective function and reach a performance level which is much closer to the function's true maximum and at times even reaches it.

In a diploid GA, an allele is considered lost if both haploid halves of the dip-

Figure 3.11: The performance measures of GA-3.

A/C Velocity: 700 knots
Revolution Period: 18 minutes
Turn Rate: 20 Degrees/Minute
Turn Radius: 35 Miles

Center of turn

35 nm

10 nm
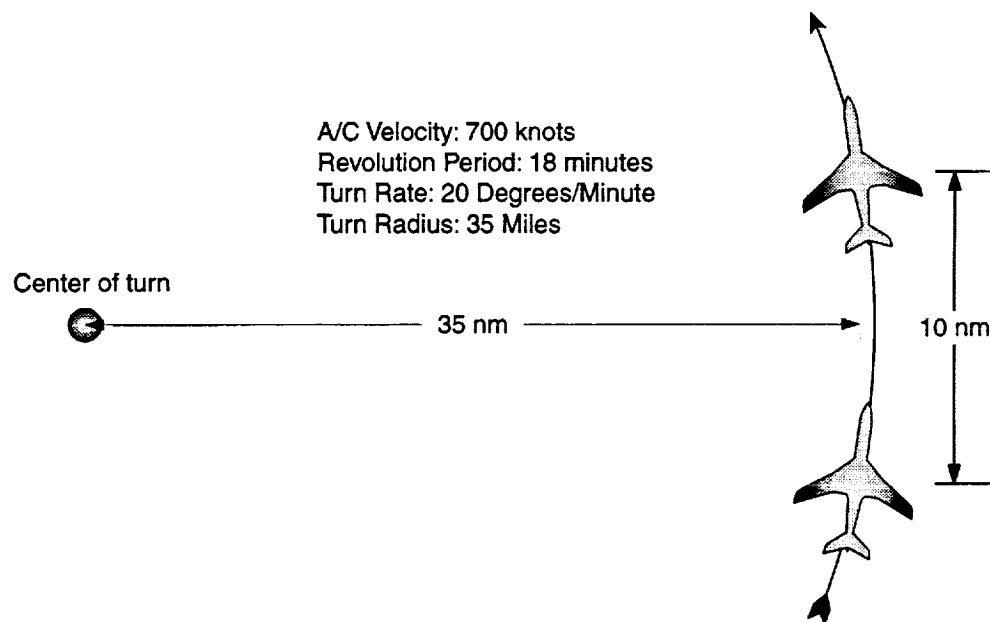
Figure 3.12: The two aircraft scenario.

loid chromosome do not hold this allele value, i.e. the allele is absent from both dominant and recessive genes. Examining the allele convergence of GA-3 revealed that zero alleles were lost or converged throughout the simulation.

Tracking a Peak in the Solution Space

To further demonstrate the diploid-GA's ability to track moving peaks in the solution space over time, a scene was chosen in which there is one major peak that remains at a constant level and shape, but moves in the solution space. The

scene consists of an aircraft pair turning about a common center, as is illustrated in Figure 3.12. The optimal viewing parameters are moving together with the aircraft, hence the objective function will have a constant-magnitude peak that is moving in the solution space.

The performance measures of GA-1 and GA-3 from a simulation of two complete turns of the aircraft pair are presented in Figures 3.13,3.14. As can be seen from Figure 3.14, the simple GA-1 quickly looses its ability to probe the solution space and rapidly loses performance. GA-3 on the other hand, manages to maintain a constant performance level throughout the aircraft motion. A slight increase in the performance level can be noted roughly at trial 70000, this is when the aircraft pair commences its second turn. This increase in performance can be attributed to the diploid GA memory mechanism. By storing alleles in a dormant state, it is capable of storing solution structures in abeyance, which the GA can reuse once they are useful again.

An additional indication of the behavior of GA-3 can be seen from Figures 3.15,3.16. These figures depict the solutions the GA holds as miniature cameras in space from a top-down view. The aircraft pair, depicted as white dots connected with lines to the ground grid, can be seen at two different simulation times. In these figures it is visible that the solution population tracked the aircraft pair through their motion.

The Niche Formation Technique

Diploidy ensures that GA-3 will not loose genetic material, i.e. will not lose alleles, and hence GA-3 can respond to gradual changes in the function's peak. Albeit, as can be seen from Figures 3.16,3.15 the solution population is still concentrated at one region of the solution space. Should a new peak emerge in a new location, e.g. a third aircraft enters the airspace, GA-3 will not be able to locate it. For this, niche formation will be used.

Since deterministic crowding (see page 40) showed superior niche formation capabilities in literature, it was selected as the basis for the niche formation GA to use in the active display optimization. However, when applied to the active display problem, this method failed to maintain the function peaks, as will be demonstrated later.

Deterministic crowding is dependent upon defining a distance between members. The GA literature that deals with niches always uses either pheno-

Figure 3.13: The performance of GA-3 in the turning A/C example.

Figure 3.14: The performance of GA-1 in the turning A/C example.

Figure 3.15: A solution distribution of GA-3

Figure 3.16: A solution distribution of GA-3

A: Equal-distance lines

B: Function contour

**Figure 3.17:** The "skewed" peaks (Part B) with respect to the distance metric (Part A) can lead to selection errors.

GA-3	
Population size	100
Selection	Deterministic with stochastic reminder. No generation gap.
Mutation rate	1e-6
Crossover	Two-point crossover at a rate of 0.6.
Chromosomes	Diploid structure.

Table 3.3: The parameters of GA-3.

type or genotype distance functions (Deb and Goldberg 1989; Mahfoud 1992; Miller and Shaw 1995; Spears 1994). These works tested their GAs on simple, synthetic function that had all their peaks both uniformly and symmetrically distributed in phenotype space (parameter space). However, in the active display problem the peaks are far from being uniformly nor symmetrically distributed, as some of the parameters represent angles and some distances.

Non-symmetrical peaks in parameter space render phenotype distance ineffective for niche formation, as is demonstrated in Figure 3.17. Remember that a deterministic crowding GA generates a new population by choosing two members for mating, and compares the fitness of each offspring with the parent closest to it in order to possibly replace it in the population. Assume M and P (mamma and papa) are two members of a GA each representing a two-dimensional solution. M and P are mated, to produce two offsprings, A and B (Abby

Figure 3.18: The shared visible volume and the camera distance.

and Benjamin?). Part A of Figure 3.17 shows equal (phenotype) distance curves around M and P, according to which, A's fitness will be compared with P's and B's fitness will be compared with M's. Part B of Figure 3.17 shows possible equal-function-value contour lines and demonstrates how a replacement error may occur in this scheme, causing one of the peaks (the peak on which B and P are situated) to lose its members. This example shows that for effective niche formation, the equi-distance lines should be parallel to equal-function-value lines, i.e. the peaks should appear as close as possible to spheres under the parameter mapping induced by the distance function. A similar analysis can be carried out for genotype distance.

A distance function was sought which bares more relation to the parameters' geometrical meaning. To demonstrate the complexity of defining such a distance, consider defining the distance between two members as being inversely proportional to the common visible volume between the two viewing projections represented by the two members, see Figure 3.18. While this distance seems to capture the essential geometrical meaning of the viewing parameters, it is not complete as it does not distinguish between viewpoints that view the same

volume from different orientations. It is also computation intensive therefore hindering on-line performance.

Both a computational simple and geometrically meaningful distance was sought, which led to the definition of the *camera distance* (see Figure 3.18) as the Euclidean distance between the camera positions the two solutions represent. The camera positions of the two solutions is:

$$x_{1,2} = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} r_x + D\cos\theta\cos\psi \\ r_y + D\cos\theta\sin\psi \\ r_z + D\sin\theta \end{bmatrix} \qquad (3.6)$$

where r_x, r_y, r_z are the coordinates of the POR, D is the viewing distance, and ψ, θ are the azimuth and inclination angles of the viewing direction, respectively, so that the camera distance is $d = \|x_1 - x_2\|$.

Eq. (3.6) is not a one-to-one mapping of the parameter space, i.e. it is enough that two members represent the same camera position for the distance between them to be zero. As a consequence, two niches which are identical in viewpoint position but differ only in viewing direction will not be created. However, the GA will hold effectively different viewing angles of the same visible volume. Also note that the field-of-view angle has no effect on the camera distance, as a consequence the GA is free to choose the best field-of-view regardless of niching considerations.

Functions that contain peaks of different size may also give rise to replacement errors. Such a hypothetical situation is demonstrated in Figure 3.19 for a one dimensional function. Like in the previous example, offspring A is considered for replacement of parent P and offspring B considered to replace parent M. After these replacements take place, peak *b* loses a member for the larger peak *a*. Note that the size which is relevant is the "base" of the peak and not its "height".

This problem was solved by introducing fitness sharing when comparing offsprings and parents. The effect of sharing on deterministic crowding is to favor offsprings/parents which distance themselves from the main population concentration. Thus, if a large peak attracts many solutions, their fitness will decrease, essentially causing a balance between the peak size and the number of population members it attracts. This new niche formation method is hereafter

Figure 3.19: The differences in the size of peaks can cause replacement errors.

referred to as deterministic crowding with sharing.

The shared fitness of each offspring which awaits insertion to the population is calculated by applying Eq. (3.4) using all members of the population in the summation, except for the parent which is a candidate for replacement. The parent's shared fitness is calculated using only the members in the population but not the offsprings.

Simulation with Niche Formation

To demonstrate and compare the niche formation technique, an example scenario with three aircraft was created such that its objective function will have pronounced peaks, illustrated in Figure 3.20. Initially aircraft a and b are flying head-on while aircraft c flies in parallel. As the scene progresses, the aircraft turn such that in the final state, aircraft b and c are flying head-on while aircraft a is flying parallel to them. The objective function of this scenario, initially has one peak, associated with viewing pair (a,b), the height of which decreases with time. A second peak, associated with viewing pair (b,c) slowly emerges and with time grows in height, finally becoming the new global maximum.

A deterministic crowding with sharing GA was configured, named GA-4, and is summarized in Table 3.4. To demonstrate the necessity of niche formation, the performance of the best member of GA-3 and that of GA-4 when optimizing the three aircraft scenario is illustrated in Figure 3.21. As can be seen, until a time of 150 seconds, when only one peak exists, GA-3 achieves a higher score. Once the second peak arises, the performance of the two GAs begins to balance, with GA-4 showing the benefits of its more diverse population towards the end of the period when the two peaks are present. At a time of 310 seconds,

Figure 3.20: The three aircraft scenario.

All turn rates: 18 deg/min
All turn radii are equal

80 nm

R 32 nm

100 nm

when only the second peak remains, GA-4 shows its superior behavior and its ability to quickly recuperate from the disappearance of one peak from the objective function by achieving a much better performance than GA-3.

A "visual" inspection of the population distribution at a time of 150 seconds is presented in Figures 3.22, 3.23 which represent the GAs' solution population from a "birds' eye" view as "black flocks" of tiny cameras in space. The aircraft are represented as white dots connected to the ground grid with drop lines. In Figure 3.22 the population of GA-3 is seen to clearly form a single cluster of solutions that are focused on the first aircraft pair. In contrast, in the population of GA-4, Figure 3.23, one can identify several clusters of solutions, at least one can be identified as focusing on the second aircraft pair. It is evident that looking only at the performance of the GA members is not sufficient, since the goal is not just to produce the best scoring members, but also to position the members such that new maxima can quickly be located. In other words, we strive to maintain population diversity.

To reflect this second goal, a new measure was defined, referred to as the

Figure 3.21: The performance of the best members of GA-3 and GA-4 in the three aircraft example.

running variance, such that the running variance ϑ_i of the i^{th} parameter is:

$$\vartheta_i = \frac{T \sum\limits_{k=1}^{T} \bar{x}_{k,i}^2 - \left(\sum\limits_{k=1}^{T} \bar{x}_{k,i} \right)^2}{T^2} \tag{3.7}$$

where $\bar{x}_{k,i}$ is the normalized value of the i^{th} parameter of the k^{th} member evaluated by the GA. The parameter is normalized such that its value range is between zero and unity. The running variance indicates the amount of variability which is maintained in each parameter through the GA's evolution.

A running variance of the three parameters representing angles, the azimuth angle, the inclination angle and the field-of-view angle of GA-3 and GA-4 are depicted in Figure 3.24. The angular parameters were chosen as a representative of all parameters. The running variance in Figure 3.24 is normalized by $1/12$ which is the variance of a uniformly distributed random variable in the range between zero and unity. From Figure 3.24 it is evident that GA-4 maintains a better population diversity in all parameters.

To compare the ability of GA-4 to maintain population diversity with that of other niche formation techniques, several other GAs have been defined, each differing from it in one aspect. Table 3.5 summarizes the parameters of these

Figure 3.22: The population of GA-3 at a time of 150 seconds.

Figure 3.23: The population of GA-4 at a time of 150 seconds.

GA-4	
Population size	50
Selection	Deterministic crowding with sharing
Mutation rate	1e-6
Metric	Camera distance
Chromosomes	Diploid structure

Table 3.4: The parameters of GA-4.

GAs, highlighting the difference of each with respect to GA-4.

The performance of the best member in each of GA-4 to GA-8 is portrayed in Figure 3.25. The normalized running variance of the azimuth, inclination and field-of-view parameters is depicted in Figures 3.27-3.28 respectively. From these figures it is evident that GA-4 is superior both in performance and in population diversity to any other niching method examined. GA-5 is the only GA which nearly reaches the performance and population diversity of GA-4. However, the deterministic crowding mechanism of GA-4 is superior to the selection by sharing mechanism of GA-5 which enables GA-4 to maintain useful population diversity and react faster to the changes in the objective function. Also note that GA-8 which uses a phenotype distance could have had some unfair advantage in the running variance graphs, as the running variance measures variability in phenotype space. From the poor performance of GA-7 and GA-8 one can deduce the effectiveness of using the geometrically meaningful camera distance.

Figure 3.24: The running variance of the angular parameters of GA-3 and GA-4 in the three aircraft example.

Figure 3.25: The performance of the best member of GA-4 to GA-8.

Figure 3.26: The running variance of the inclination parameter in GA-4 to GA-8.

Figure 3.27: The running variance of the azimuth parameter in GA-4 to GA-8.

Figure 3.28: The running variance of the field-of-view parameter in GA-4 to GA-8.

GA-5	sharing versus deterministic-crowding-with-sharing
Population size	50
Selection	Deterministic with stochastic reminder. no generation gap. selection based on shared fitness.
Mutation rate	1e-6
Crossover	Two-point crossover at a rate of 0.6.
metric	Camera distance
Chromosomes	Diploid structure
GA-6	**Simple deterministic crowding (without shared fitness)**
Population size	50
Selection	Deterministic crowding
Mutation rate	1e-6
Crossover	Two-point crossover.
Metric	Camera distance
Chromosomes	Diploid structure
GA-7	**Genotype distance versus camera distance**
Population size	50
Selection	Deterministic crowding with sharing
Mutation rate	1e-6
Crossover	Two-point crossover.
Metric	Genotype distance
Chromosomes	Diploid structure
GA-8	**Phenotype distance versus camera distance**
Population size	50
Selection	Deterministic crowding with sharing
Mutation rate	1e-6
Crossover	Two-point crossover.
Metric	Phenotype distance
Chromosomes	Diploid structure

Table 3.5: The parameters of GA-5 through GA-8, and their difference with respect to GA-4.

3.5 Summary

Achieving fast convergence for the evolving objective function and a tight maximum tracking is pertinent not as much because the function is expected to change rapidly, but due to the on-line nature of the active display system. In a mode in which the user is involved in choosing the timing of the change, either by directly requesting a viewpoint change or authorizing one, response time is critical to achieve a usable system. To this end, an optimization algorithm that evolves with the airspace was attempted, striving to create an agile algorithm that uses past knowledge of the state of the airspace to quickly locate the new maximum, rather than starting a fresh search at each time step. Diploid GAs have been first tried in a complex time evolving function and have exhibited their ability to maintain genetic material, enabling them to track time changing maxima. A new niche forming technique was developed, based on a geometrically meaningful distance which reflects the objective function's topology, rather than using phenotype or genotype distance as is common in literature. In addition, it was shown that the usage of deterministic crowding together with fitness sharing yields an effective niche forming GA which can deal with a changing number of different-sized peaks in the objective function.

4

Part Task Experiment: Static Images

4.1 Purpose

While the previous chapter dealt with finding the global optimum of a multi-modal optimization surface, this chapter deals with the question, whether and to which extent, this optimum indeed reflects the "best" display for the operator. This question can be reformulated by questioning whether the objective function that was formulated, accurately reflected the shortcomings in the human spatial reconstruction process and correctly analyzed the task to be performed. In order to answer this question, a preliminary part-task experiment was carried out in which active air traffic controllers were asked to interpret abstracted three-dimensional static ATC situations. The objectives were:

1) To evaluate the effect of the viewing parameters on the operator's ability to interpret perspective scenes.

2) To establish a relationship between the image optimization score and the operator's performance in interpreting the spatial scene.

3) To compare the operator's performance in interpreting the spatial scene for automatically and manually chosen viewing parameters.

4.2 Description of the Task

In order to rate the operator's ability to spatially interpret a scene, a task was designed requiring the operator to judge current separation distances between aircraft. Although this ability would be highly desirable in a flow control task, the choice of this performance score does not suggest that the perspective format should indeed be used for active flow control. In such a task, accurate judgements could very well be made with conventional plan view displays. The experiment consisted of two basic run types.

Critical pair selection runs measured the operator's ability to judge horizontal and vertical separation distances. A perspective scene was shown, abstracted from a hypothetical ATC situation, as in Figure 4.1. The aircraft in the scene are symbolically represented by horizontally oriented circles, located at the aircraft positions and connected to a rectangular ground grid by drop lines. The subjects are requested to identify the aircraft pairs that violate a minimal horizontal and vertical separation criterion. The separation distances are different for each run.

The criterion is communicated to the subject in a graphical manner at the start of each run, as shown in Figure 4.2, separate from the aircraft scene display. The criterion is displayed as a cylindrical volume, centered about an aircraft symbol and dropping a shadow on the ground. An aircraft pair is considered in violation of the criterion when each aircraft is within the volume of the other.

The volume is displayed with viewing parameters different from those the aircraft scene is viewed with. The transition between the two settings is done in a gradual manner resulting in a smooth motion pattern. This scheme was chosen to randomize the appearance of the volume and the motion cues involved in the transition contributed to a better understanding of the criterion. The volume display could be recalled by the subject at any instant of the experiment run by pressing a button.

Figure 4.1: The experimental scene with the manual viewpoint controls.

Figure 4.2: The critical selection volume display.

Viewing parameter setting runs served to examine manually chosen viewpoints. The operator was asked to choose the viewing parameters such that they would be able to interpret the perspective scene in the best possible way.

Based on previous experience with perspective displays (Grunwald 1996), the viewing parameter setting scheme was constrained to a preset pattern in which the viewing axis could be pivoted about a POR located in the ground plane. Thus, the viewing parameters that could be varied by the operator were: the azimuth and elevation angles of the viewing axis, the viewing range, i.e. the distance between the eye-point and the POR, the zoom angle and the location of the POR. By using *manipulators* integrated in the three-dimensional scene, see Figure 4.1, each one of these parameters could be set by independent operations.

4.3 Subject Background, Instruction and Training

Two female and eight male active Air Traffic Controllers participated in the experiment. Subjects age was between 29 and 46 years. All subjects had between 4-8 years of TRACON experience and three subjects also had control tower experience.

The subjects were told that the purpose of the experiment was to evaluate their ability to interpret spatial scenes. The two types of runs were demonstrated, during which the subjects were familiarized with the task, the separation criterion, the method of selecting critical aircraft and the viewing parameter setting method. Following this demonstration the subjects performed a series of 12 aircraft selection training runs. During the first six runs, the subjects received both immediate audio and visual feedback on their selections, in the last six runs, only audio feedback was given. In addition, they also performed six viewing parameter setting runs. The instruction, demonstration and training lasted between one and two hours.

The production included nine viewing parameter setting runs and 90 aircraft selection runs, and lasted between two and three hours. The available time to complete each run was limited and a dial in the lower right corner of the display, see Figure 4.1, showed the remaining time in minutes and seconds. The subjects were allowed to terminate the run when they felt that they identified all critical aircraft pairs.

4.4 Experiment Design

4.4.1 Independent Variables

The experiment's independent variables were:

1) The aircraft load level, i.e., the number of aircraft that was present in the scene.

2) The optimization level, i.e. the predicted quality of the viewing parameters, as computed by the objective function.

3) The parameter setting mode: manual or automatic.

4.4.2 Data Sets

Three load levels (20, 35 and 50 aircraft), and three image optimization score levels (high, medium and low), were considered. For each load level, three different scenes were tested, yielding a total of nine different scenes. Each scene had 5 critical aircraft pairs. For each scene, three different viewing parameter settings were computed for each optimization level, yielding a total of 81 different settings. Each setting comprised an experimental run. Thus, production runs included the 81 viewing parameter settings and the nine manual settings. The the total sequence of 90 runs was randomized, and the subjects were not told that the runs with their manually chosen settings were also present in the set.

The objective function consisted of a weighted sum of the viewing qualities of the ground and altitude separation[†] of all the aircraft pairs with a separation less than twice the selection criteria separation. The inverse of the actual separation distance was used as a weight, thus the closer an aircraft pair, the more important it was in the optimization. The familiarity cue model that was used in the perception model stipulated the observer knows the ground grid is square and of a known constant size, and that the drop bars connecting the aircraft symbols to the ground grid are perpendicular to the ground plane. The random viewing parameters at different optimization levels were generated by randomly choosing viewing parameters from the solution population history of a genetic

---

[†] The objective function consisted of the static elements of Eq. (A.1) described in Appendix A.

algorithm that was configured to explore a large portion of the solution space.

The raw experimental data for each selection run were the identities of the aircraft pairs selected by the subject, the time needed to complete the run, and the number of times and the total duration the critical volume was inspected. Each selected aircraft pair was classified as a hit or a false alarm depending whether it actually violated or did not violate the criterion. Additionally, the critical aircraft pairs that were not picked out by the subject were marked as misses.
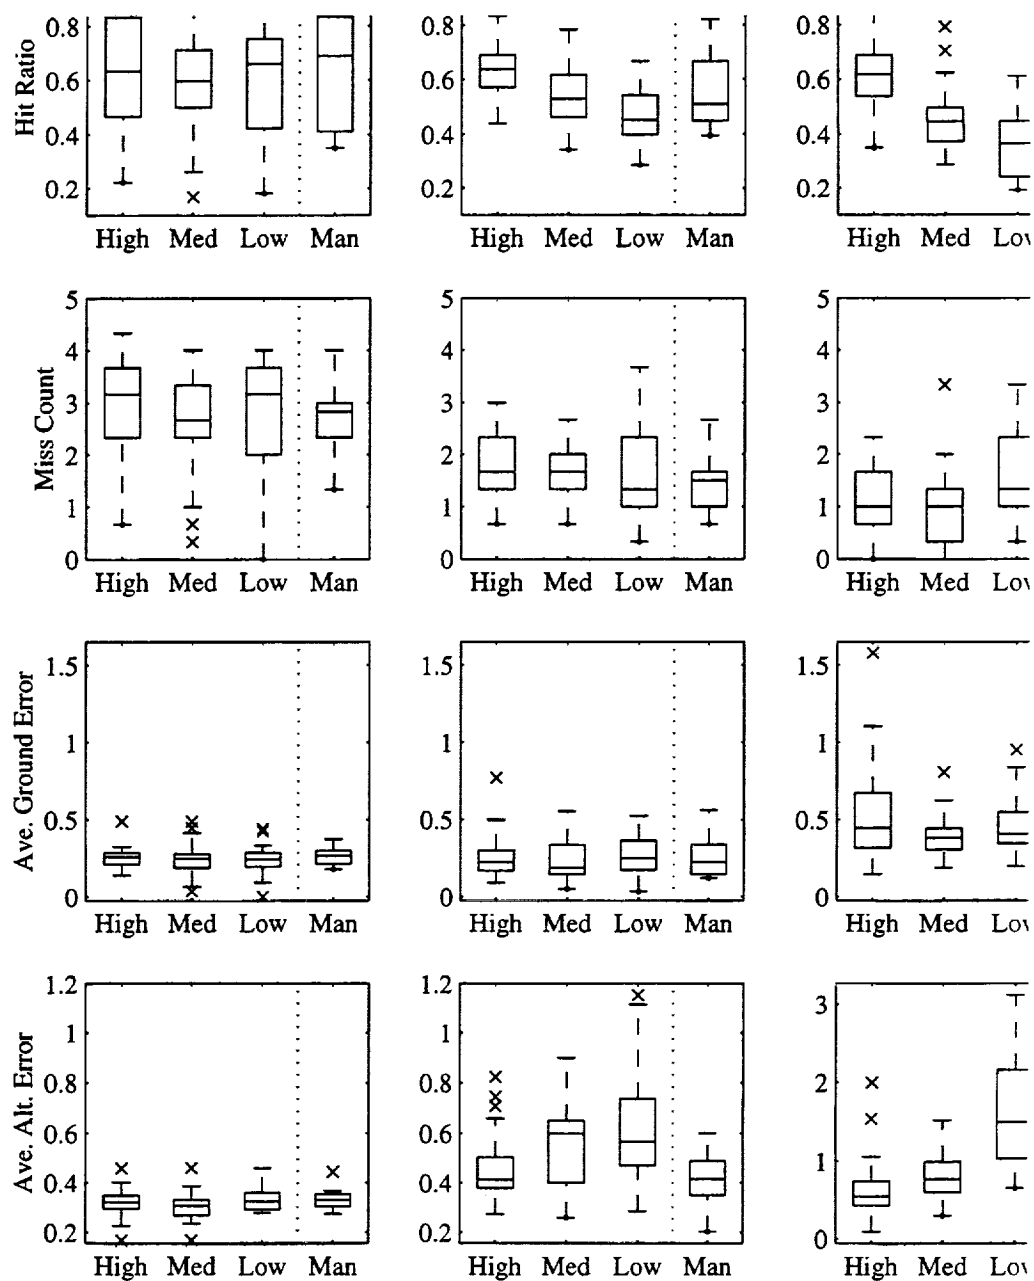
4.5 Experiment Results

The ratio of hits over false alarm tested using a χ^2 test on a contingency table. All the tests were done separately for each one of the three aircraft load levels. In the 35 and 50 load levels it was found, at a significance level of 0.001, that the ratio was different for the three optimization levels. In the 20 load level no significant difference was found. No significant difference in the ratio was found between viewpoints of the same optimization level and load level.

The data were further analyzed by performing a Friedman, 2-way, non–parametric test, and using a Page ordered alternatives test. The distribution of the data, from all subjects, for each experimental parameter setting (9 runs times 10 subjects) are depicted by the box plots of Figure 4.3. The box extends from the lower to the upper quartiles of the distribution with a line drawn at the median. The whiskers, extending from the boxes, represent the rest of the distribution. Outlier points are drawn by "x" marks.

The upper row of Figure 4.3 shows the hit ratio, i.e. the ratio of correct over total answers. As can be seen in the 20 load level, the optimization does not affect the hit ratio, while in the 35 and 50 load levels the hit ratio clearly improves with the optimization level. It also appears that this improvement is the strongest for the 50 load level.

The second row of Figure 4.3 shows the miss count distributions. An effect of the optimization level on the miss count was only detected in the 50 load level at a significance of 0.01, where the high and medium optimization levels were better then the low optimization.

The third and fourth rows of Figure 4.3 show the distribution of the average ground and altitude "error distances", respectively. The error distances rate the severity of false alarms and misses. For these cases, the normalized ground error

Figure 4.3: A summary of the experimental results by scene type and optimization

Experiment Results

distance is defined as the difference between the pair's actual ground separation and critical separation, normalized by the critical separation. The ground error is averaged over all false alarms due to ground error and all misses. The average altitude error distance is calculated in a similar fashion.

The average ground distance error was not affected by the optimization. However, at a significance of 0.01, the average altitude error reduces with the optimization level. This is true for the 35 and 50 load levels. Again, in the 20 load level no significant difference was observed.

The results clearly indicate that the optimization level affects subject performance and thus the ability to spatially interpret the scene. The sensitivity of the altitude error to the optimization level indicates the ability of the observer to reconstruct the vertical dimension in optimized viewpoints. In contrast, the insensitivity of the ground error to the optimization level indicates the observer uses the metric of the ground grid rather then spatially reconstructing.

Furthermore, in almost all scores, a trend is observed in which the more complex scenes are more affected by the optimization level. A possible explanation is, that for simple scenarios, i.e. the 20 load level, the detection problem is too simple, so that even from bad viewpoints the situation can be assessed reasonably well, whereas in complex scenes the choice of a correct viewpoint is more critical.

Figure 4.3 also shows that the performance in the manually chosen viewpoints was worse then in the highly optimized ones. This is true for the cases in which the optimization was effective.

4.6 Conclusions

The results clearly demonstrate that the effectiveness of perspective displays is highly affected by the choice of viewing parameters. The ability of the operator to spatially reconstruct scenes generally improves with the level of optimization, where the strongest performance improvement is for the high aircraft load levels. This indicates that complex scenes in particular, are benefiting from viewing parameter optimization. This might explain why 2-D displays are at times found to perform better then 3-D displays when simple scenarios are used, as was done by Jasek *et al.* (1995) and O'Brien *et al.* (1997). A proper comparison between 2-D and 3-D displays should be done for different load levels and with images for which the viewing parameters are optimized. The fact that the per-

formance for the manually chosen viewpoints was inferior to the one for the optimally chosen viewpoints, indicates the potential advantages of the method. The simple and abstract optimization function used in this experiment served to prove the method by which image quality is rated.

5
A Model for Spatial Perception from Moving Images

5.1 Introduction

The model, presented in Chapter 2, dealt with the spatial perception of three-dimensional layouts from static perspective images. Although this model proved itself very useful in understanding the human spatial reconstruction process, and was at the heart of the objective function used in the viewing parameter optimization scheme, it does not incorporate the time-varying nature of the ATC task. Therefore, it enables the prediction of judgement errors of current states only. From the definition of spatial awareness, given in Section 1.2.1, it follows that the controller is required to judge not only the current aircraft position, but also to judge their current velocities and flight paths. To estimate the controller's perception errors of dynamic scenes from moving images, a dynamic spatial perception model was formulated. Previous success with using

a Kalman filter to model ego-motion perception (Zacharias *et al.* 1985; Zacharias *et al.* 1995), inspired us to extend the static formulation of Chapter 2 to an optimal filtering framework, thus enabling the modeling of object motion perception from moving images.

5.2 General Description of the Model

5.2.1 The Geometry and Dynamics

The viewing geometry is similar to that used in Figure 2.1 of Chapter 2 and its parameters are the unknown vector χ which consists of the viewing parameters and the familiarity cue parameters, as stated in Eq. (2.4). When the viewed objects are in motion, the vector χ changes with time. The dynamic model assumes the observer is familiar with the kinematics of the motion, and attempts to estimate its parameters. By assuming the observer uses the dynamic model, this can be written as

$$\dot{\chi} = \mathcal{G}(a, t, \chi(t)) \tag{5.1}$$

where the function $\mathcal{G}$ is the *motion cue function* and a is the unknown *motion parameter vector*. The function $\mathcal{G}$, the vector a and Eq. (5.1) play an analog role to the familiarity cue functions $\mathcal{F}_j(\bullet)$, the parameter vector c, and Eq. (2.3) of the static model.

For example, assume that the parallelepiped in the example of page 24 is growing taller at a rate that the observer assumes to be constant. Then the parameter vector a might consist of the unknown rate of change and the function $\mathcal{G}$ will express the relation between the change in the parallelepiped height and the vector a.

The observer estimates the vector $X = (\chi, a)$ which contains both the vector χ and the motion familiarity cue vector a,

$$X = (\psi, \theta, \kappa, D, r, c, a)^T. \tag{5.2}$$

When viewing an air traffic situation with fixed viewing parameters, the derivatives of the viewing parameters, as reflected in the function $\mathcal{G}$ are zero.

5.2.2 The Reconstruction Process

According to the model, the observer assumes that the motion parameter vector q may have random variations, which are described by:

$$\dot{q}(t) = \underline{w}_a(t) \tag{5.3}$$

where $\underline{w}_a(t)$ is a zero-mean Gaussian noise process with a diagonal covariance matrix Q_a, expressing the parameter rate of change which is expected by the observer. For example, if the observer assumes the i^{th} element of q is likely to change by Δa_i in the time interval Δt, then

$$i^{th} \text{ diagonal element of } Q_a = \frac{\Delta a_i^2}{\Delta t}. \tag{5.4}$$

Augmenting the dynamic model of Eq. (5.1) with Eq. (5.3) yields:

$$\dot{X} = \begin{pmatrix} \dot{\chi} \\ \dot{q} \end{pmatrix} = \begin{pmatrix} \mathcal{G}(q, t, \chi(t)) \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ \underline{w}_a(t) \end{pmatrix}. \tag{5.5}$$

It is assumed that the observer estimates the vector X by using a measurement vector Z which consists of the two-dimensional coordinates of the viewed object in the perspective image. The vector Z relates to the vector X via the relation

$$Z = \mathcal{H}(X). \tag{5.6}$$

The function $\mathcal{H}(\bullet)$ can be derived by using the fundamental assumption of the static perception model. Namely, that the observer reconstructs the vector X by minimizing a cost function $J(Z)$, this can be written as:

$$X = \arg \min J(Z). \tag{5.7}$$

Hence, $\mathcal{H}(\bullet)$ is the inverse of the relation in Eq. (5.7).

The observer reconstructs the scene by an on-line estimation of the state X using the measurement vector Z. The model assumes that this estimation process can be modeled as a Kalman filter that estimates the states of the dynamic system in Eq. (5.5) while using the measurement equation of Eq. (5.6). By writing the equations for the estimation error covariance matrix of this Kalman fil-

ter, the judgement errors of the observer can be predicted.

Note, that Eq. (5.7) states that both parts of the vector X (the parameter vector χ and the motion parameter vector q) may be inferred from a single static image. This is true if the image contains geometrical representations of motion parameters, e.g. velocity vectors.

5.3 Mathematical Formulation

The first step to writing the filter covariance equation is to linearize Eqs. (5.5),(5.6) about the true viewing parameters and the true parameter vector X. Let the value of the parameters at the linearization point be designated by the subscript $(\cdot)_0$, then Eq. (5.5) can be linearized as:

$$\dot{x} = A(q_0, t)x + w(t) \tag{5.8}$$

where the matrix $A(q_0, t)$ is the linearization of the function $\mathcal{G}(q, t)$ about q_0:

$$A(q_0, t) = \begin{bmatrix} \dfrac{\partial \mathcal{G}}{\partial \underline{\chi}}\bigg|_{\chi_0} & \dfrac{\partial \mathcal{G}}{\partial q}\bigg|_{q_0} \\ 0 & 0 \end{bmatrix} \tag{5.9}$$

and $x = X - X_0$ is the difference between the true parameter values and the estimated parameter values.

To linearize Eq. (5.6), we note from Eq. (5.7) that the matching cost J is at a minimum, hence linearize the gradient of J, and equate to zero:

$$J_x(X, Z) = 0 = J_x(X_0, Z_0) + J_{xx}x + J_{xz}z \tag{5.10}$$

where Z_0 is the coordinate vector of the true image, and $z = Z - Z_0$ is the difference between the coordinate vector of the virtual image and that of the true image. Since for the true image the cost is at a minimum, $J_x(X_0, Z_0) = 0$, solving for z Eq. (5.10) becomes:

$$z = J_{xz}^+ J_{xx}x \tag{5.11}$$

where J_{xz}^+ is the Moore-Penrose Pseudo-Inverse of the matrix J_{xz}. By assuming a random measurement noise and substituting $H = J_{xz}^+ J_{xx}$, we get the measurement equation:

$$z = Hx + v(t) \tag{5.12}$$

where $v(t) \sim N(0, R)$. Note that $H = H(t)$ is a function of time, as it depends on the perspective projection and the position of the objects in the three-dimensional scene, as the cost J is affected by their values.

It is now possible to write the error covariance propagation equations for a continuous extended Kalman estimator which is estimating the states in Eq. (5.8), using the measurements of Eq. (5.12) as:

$$\dot{P} = A(t)P + PA(t)^T + GQG^T - PH(t)^T R^{-1} H(t)P \tag{5.13}$$

Eq. (5.13) describes the propagation of the covariance of the observer's judgment errors with time. In the ATC implementation, the $A(t)$ and $H(t)$ matrices can be considered as quasi-static, and the error covariance $P(t)$ of the filter can be assumed to be at a steady state. This is due to the typical scales of ATC in which the controlled volume is large compared with aircraft velocity, causing the aircraft to remain for long periods of time in the image, slowly moving in it.[†] Hence the algebraic form of the Riccati equation, obtained by setting $\dot{P} = 0$ in Eq. (5.13), is used.

5.4 Comparison of the Dynamic and Static Model

Since both models are based on the minimization of the matching cost, it is expected that in a static image, both models will yield the similar judgement error behavior. The scene that was used for the comparison is the scene with two disks with drop bars that was used in the analytical evaluation of the static model (Section 2.4 on page 30).

Figure 5.1 shows the predicted judgment errors of the ego parameters with the two models for different azimuth angles of the viewing direction. In producing Figure 5.1, the noise covariance matrices, R and Q_a, were calibrated such

---

[†] Typically, a controlled sector is 150 nautical miles in diameter, an aircraft velocity is 600 knots, taking the aircraft about 15 minutes to traverse the controlled sector.
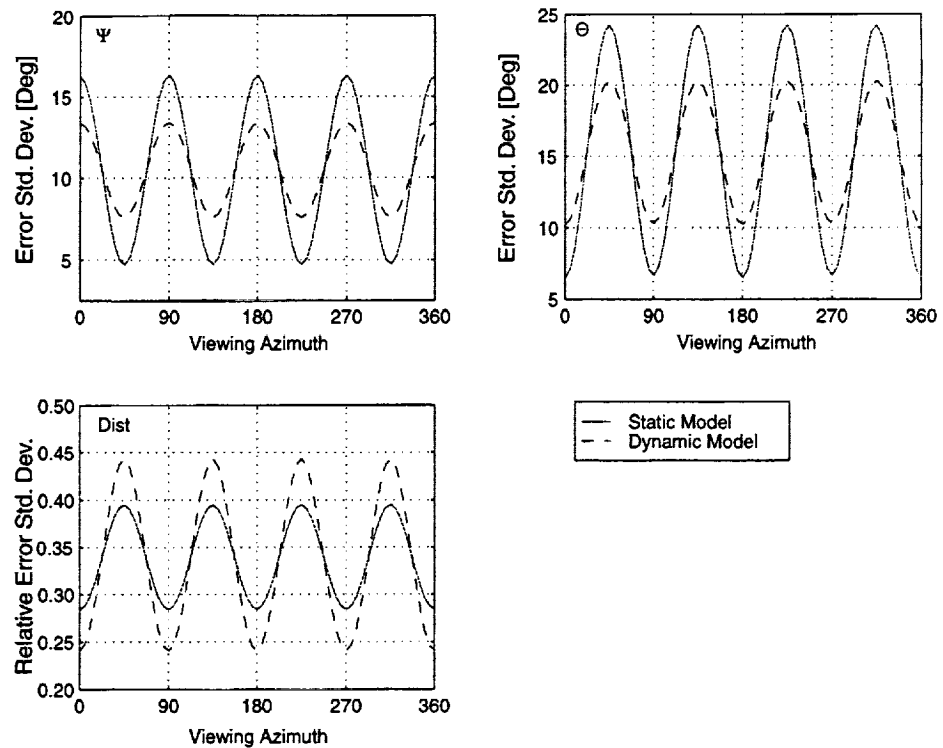
Figure 5.1: The judgment errors of the ego-parameters: viewing azimuth Ψ, viewing inclination θ and viewing distance, as a function of the viewing azimuth.

that the errors in the graphs of the static model predictions and the dynamic model predictions will be minimal. Figure 5.1 shows that the two models exhibit similar trends. The graphs are not identical since the two models assume different mechanisms in creating the judgment errors.

5.5 Summary

Spatial awareness requires the controller to perceive aircraft motion and future trajectories. Hence, by predicting the judgment errors of motion parameters, the dynamic perception model can be used to rate the degree in which a perspective image supports spatial awareness. This rating system was tested in the experiment presented in Chapter 6. The model is a reformulation of the static perception model in a filtering framework. Its novelty is in the ability to measure the judgement errors in perceiving the spatial motion of objects from perspective images.

6

Part-Task Experiment: Moving Images

6.1 Introduction

ATC scenes, that evolve in time, demand from the operator to judge future situations by projecting the current situation into the future. The dynamic spatial perception model of Chapter 5 describes the process by which the operator utilizes the dynamic changes in the display for estimating the motion of objects in the scene. Like with the static model, the dynamic model is incorporated into the objective function. In this case this objective function rates the dynamic image on how well the operator is able to judge current and future situations from this image.

The part-task experiment presented in this Chapter, is designed to evaluate whether the selected optimal viewing parameter setting indeed yields the "best" display from which the operator is able to judge *both current and future situations*. The experiment included two tasks. In the first one, abstracted from the ATC

flow control task, the subject was asked to select from an ATC scene evolving in time, the pairs of aircraft which might violate separation constraints, sometimes in the future. In the second task, the subject was asked a number of questions about a specific aircraft pair in the scene, related to the present and future relative position between the aircraft. Since the experiment was conducted fully automatically (without the experimenter verbally asking the questions), a unique response method has been developed. In this method a data tablet was used, sensitive to the pressure of a special pen. On this tablet, a "routing chart" was drawn, on which the questions were represented graphically. The method was inspired by the graphical layout of the Cooper-Harper scale, used in evaluating vehicle handling qualities. As the subject proceeded through the chart, the chosen path determined the next question to be asked. The unique advantage of the chart was, that it allowed questions about azimuth angles between aircraft to be answered in a natural way, by indicating this azimuth angle on a wind rose. Thus, the horizontally placed wind rose naturally represented angles in the horizontal plane.

6.2 Purpose

In the second part-task experiment abstracted ATC scenarios evolving in time were shown to subjects. They were asked to interpret dynamic and spatial relations and to project their future state. The objectives were:

1) To evaluate the effect of viewing parameters on the operator's ability to project future situations.

2) To evaluate the effect of viewing parameters on the operator's ability to interpret dynamic quantities.

3) To compare between the operator's spatial judgement performance in optimized perspective displays and in "conventional" plan-view displays.

4) To compare between the operator's spatial judgement performance in conventional and in enhanced graphical presentations.

6.3 The Graphical Display

Four display types were used which differed in their usage of perspective or plan-view projection and in the graphical symbology.

6.3.1 Conventional "Plain" Graphics

A graphical representation was modeled after the symbology common to current ATC displays, and is referred to as the "plain" graphics. The graphical representation includes only elements which are necessary for the experiment task and is not intended as an operational ATC display. A perspective display with plain graphics is presented in Figure 6.1. Each aircraft is represented as a circular disc, connected to the ground grid by a drop bar which consists of gray/white colored segments at 2000 feet intervals. The segments, which were not present in the first experiment, were added as a reference to aid the subjects in the detection of climb and descent rates. The ground is represented by a rectangular grid evenly spaced at 10 nautical mile intervals and spanning 150 nautical miles, representing a typical ATC control sector in En-Route Centers. The plan-view display with plain graphics is a top-down orthogonal projection of the symbology in the perspective display, augmented with data-tags displaying aircraft altitude in hundreds of feet; it is depicted in Figure 6.2.

The motion of the aircraft is shown for a duration of 15 seconds, referred to as the trace period, after which the aircraft are retraced to their initial position for the next trace. The aircraft display is blanked for 0.4 seconds during the retrace in order to avoid apparent motion reversal and to circumvent comparison of end-to-start positions. The relatively short duration of the trace period is sufficient to provide motion cues without changing the image considerably, thus keeping the optimization score approximately constant.
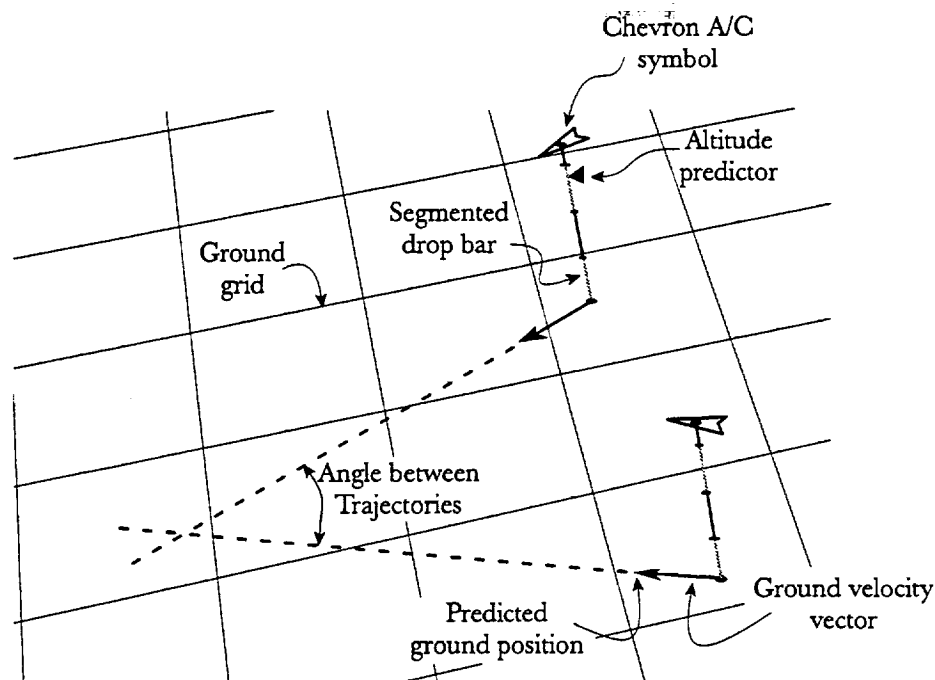
Figure 6.1: A perspective display with plain graphics.

Figure 6.2: A plan-view display with simple graphics.

Figure 6.3: The graphically enhanced aircraft representation in the perspective display and the angle between trajectories.

6.3.2 Enhanced Graphics

The enhanced graphics adds cues to aid the perception of aircraft ground and vertical velocities, as these data will be available in future ATC developments[†] (FAA 1996; RTCA 1995). A detailed view of the aircraft representation in the graphically enhanced perspective display is shown in Figure 6.3, the full display is depicted in Figure 6.4. The aircraft are represented as flat chevrons pointing at the direction of flight. At the bottom of each drop bar an arrow in the ground plane represents the ground velocity by pointing at the aircraft's predicted position 30 seconds ahead, assuming constant ground velocity, such that the length of the arrow and the size of its head are proportional to the ground speed. For climbing or descending aircraft, a triangle is drawn pointing at the predicted altitude 30 seconds ahead, assuming constant vertical speed. The triangle is colored red if the aircraft is descending and green if it is ascending, in which case the drop bar is extended to the future altitude. Figure 6.5 shows that the enhance-

---

[†] Current aircraft equipment does not provide rate of decent/climb information, but such information may be available in the future, when GPS is incorporated. Ground velocity is presently unavailable as well, however, air speed, as measured by the aircraft, is transmitted to the ATC station, which enables the controller to compare speeds of aircraft traveling at the same altitude and direction.

**Figure 6.4:** A perspective display with enhanced graphics.

Figure 6.5: A plan-view with enhanced graphics.

ments in the plan-view display includes the ground velocity arrow and a rate of climb/descent readout in feet/minutes. For non-zero rates, a green up-pointing arrow is shown for climb and a red down-pointing arrow for descent. The same trace and retrace periods as in the plain graphics were used.

6.4 Description of the Task

The subjects were asked to perform two different spatial tasks. The first task, a pair selection run, required the subject to project future situations; the second task, a questionnaire, required the subject to estimate spatial and dynamic quantities of selected aircraft pairs.

6.4.1 Pair Selection Runs

The purpose of these runs was to evaluate the subjects' ability to interpret moving spatial scenarios and project future horizontal and vertical separations. An animated ATC scene was shown, abstracted from a hypothetical ATC situation. The subjects were told that in this scene, four out of all aircraft pairs will violate a horizontal and vertical separation criteria sometime in the future. These pairs were referred to as "aircraft pairs in violation", and the subjects were asked to identify them. The available time to complete the task was limited and was displayed on a dial at the lower right corner of the screen, see Figures 6.1-6.5. The run was terminated automatically if the time limit was reached or if more than six aircraft pairs were selected. The subjects were allowed to manually terminate the run earlier if they felt they identified all the violating aircraft pairs.

The separation criteria were 10 nautical miles in ground distance and 1000 feet in altitude separation for all runs. The criteria were communicated to the subjects both verbally and using the separation volume display used in the first experiment, see Figure 4.2 on page 68. The subjects were told the grid size and the drop bar segment length. Each aircraft pair selected by the subject was classified as a "hit" if it was in violation or as a "false alarm" if it wasn't. In addition, violating aircraft pairs that were not selected by the subjects were classified as a "miss".

6.4.2 Questionnaire on Aircraft Pair State

The purpose of these runs was to evaluate the subject's performance in estimating spatial and dynamic features of specified aircraft pairs and in projecting their

future state. An air traffic situation is shown to the subjects, identical to the pair selection runs. A queried pair was highlighted, one aircraft in red and the other in blue. The subjects had to answer a series of questions designed to evaluate their spatial judgment performance of the aircraft situation. A novel method for presenting the questions and responding to them was used. The questions were drawn in a map that was placed on a data tablet. The subjects answered the questions by traversing the map, and clicking with a special pen along the route. A full-scale map is presented in Figure 6.17 (page 104), a sample traversal path is shown in Figure 6.6, the question answered by each pen click along this path are explained in Table 6.1. Depending on the route in the map which the subject chose, the following judgements were made: the time the aircraft will reach equal altitude, the angle between trajectories (see Figure 6.3), current and future (two minutes ahead) relative position, and the confidence level of these judgements. The advantages of using a map to perform the judgements was twofold: (1) judgements of horizontal situations were answered on a horizontal device, avoiding a possible source of errors (Cunningham and Pavel 1991; Mittelstaedt 1991); (2) the map provided a visual representation of the questions enabling rapid traversal once the subjects were familiarized with its layout.

The goal was to complete the questions in one minute. A dial showing the time left was displayed at the lower right corner of the screen and a bell was rang at the end of the time. The subject was permitted to complete the map traversal even if the goal time was reached.

The set of aircraft pairs that was queried included the violating pairs and those the subject selected in the selection run. A random aircraft pair was added to the queried set in case all four violating pairs were correctly selected to prevent the subject from recognizing the queried pairs.

6.4.3 Run Sequence

The entire set of runs was randomized for each subject. The pair selection run always preceded the questionnaire run for a specific airspace and viewpoint. Thus the subject would alternate between answering a pair selection run and an aircraft questionnaire run. The aircraft pairs in the questionnaire run were randomly ordered.

Figure 6.6: The map with the questions on the relative aircraft pair situation and a sample traversal path, indicating the pen clicks.

Click	The Question Answered
1	Start, aircraft pair is highlighted
2	Ground separation between aircraft is closing
3	Altitude separation between aircraft is closing
4	Aircraft altitude will be equal in 2 minutes
5	Heading of blue aircraft with respect to red's heading is 105 degrees clockwise.
6	Currently, blue aircraft is at a distance of 30 miles from red, at an angle of 45 degrees counter-clockwise with respect to red's heading.
7	In two minutes, blue aircraft will be at a distance of 15 miles from red, at an angle of 90 degrees clockwise with respect to red's heading
8	Subject is very confident of his/her answers.
9	End, aircraft pair is de-highlighted.

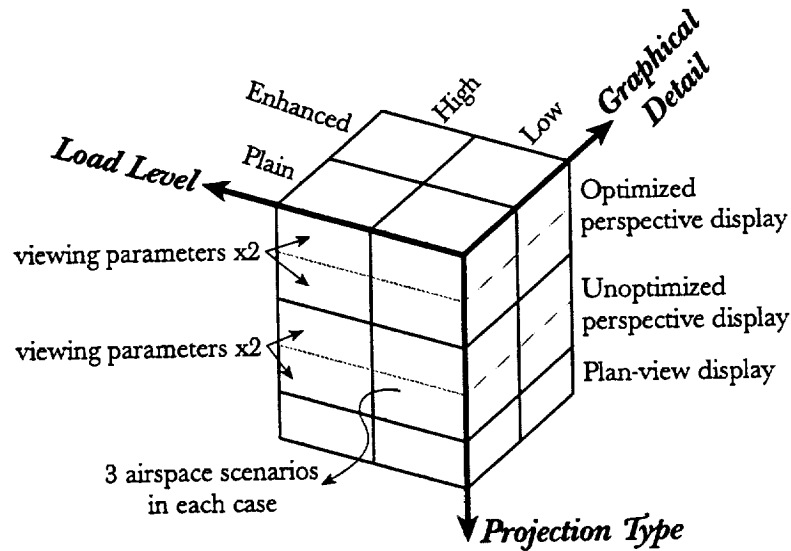**Table 6.1:** The question answered by each click of the sample path.

Figure 6.7: The experiment data-sets.

6.5 Experiment Design

The experiment's independent variables were:

1) *The aircraft load level*, i.e. the number of aircraft in the airspace. Two load levels were tested, 15 aircraft representing current aircraft loads in busy ATC centers, and 25 aircraft.

2) *Projection type*, unoptimized perspective display (low image score), optimized perspective display (high image score) and plan-view display.

3) *Level of graphical detail*, i.e. plain graphics and enhanced graphics.

Two viewing parameter settings were generated at each optimization level, and three airspace scenarios were generated for each combination of independent variables, as is depicted in Figure 6.7. The objective function used in generating the optimized viewpoints utilizes the dynamic perception model and is given in Eqs. (A.1),(A.2) on page 111. The familiarity cues that were used in the dynamic perception model stipulated the observer knows the size of the ground grid, that the drop bars are perpendicular to the ground plane, and assumes the aircraft are moving at a constant velocity. In the graphically enhanced scenes, the observer also assumes that the arrow is proportional to the aircraft velocity. The different optimization levels were generated by randomly choosing viewpoints with different scores from the solution population history of a genetic algorithm that was configured to explore a large portion of the solution space.

# 6.6 Subject Background, Instruction and Training

Two subject groups participated in the experiment; five active air traffic controllers, of which one was female, and five graduate students of which two were female.

The subjects were told that the purpose of the experiment was to evaluate their ability to interpret spatial scenes. The two types of runs were demonstrated in all the display types, during which the subjects were familiarized with the task, the separation criteria, the method of selecting aircraft pairs and the usage of the tablet. Following this demonstration the subjects performed a series of eight aircraft selection training runs and eight aircraft questionnaire runs. During the first two runs of each type, the subjects received both immediate audio and visual feedbacks on their answers. The last six runs of each type were conducted with only audio feedback in the selection run, and with the instructional graphics being presented after repetitive mistakes in the questionnaire runs. The instruction, demonstration and training lasted about two hours. The production (data collection) included a questionnaire run and a selection run in each of the 60 scenes (see Figure 6.7), and lasted between five to six hours.

6.7 Results

The results are summarized in Figures 6.8-6.16, which present bars grouped by display type: "2D" being the plan view display; "High" being the optimized perspective display, i.e. high image score; "Low" being the unoptimized perspective display, i.e. low image score. The error bars (only upper half is shown) represent one standard error. Effect was tested using an ANOVA, and a post-hoc Games and Howell (1976) comparison was used to check for differences at a significance level of 5% ($\alpha = 0.05$). Horizontal lines are drawn between bars for which the effect of display type, inside each group, was found significant.

Figures 6.8,6.9 present the mean number of hits and false alarms in the selection runs. The hit results were further analyzed in Figures 6.10,6.11 which show the probability of a violating pair to be selected $P_r(\text{Hit})$, and a decomposition to two cases; $P_r(\text{Hit} | \text{VS} \neq 0)$ is the probability of a violating pair to be selected if one, or both, of the aircraft has a non-zero vertical speed, referred to as a vertical violation; $P_r(\text{Hit} | \text{VS} = 0)$ is the probability of a violating pair to be selected if both aircraft have a zero vertical speed, i.e. the aircraft motion is

Figure 6.8: The mean number of hits as a function of the display type and the aircraft load.

Figure 6.9: The mean number of false-alarms as a function of the display type and the air-craft load.

Figure 6.10: The probability of a violating aircraft pair to be hit at a high aircraft load level, decomposed by the presence of vertical speed.

Figure 6.11: The probability of a violating aircraft pair to be hit at a low aircraft load level, decomposed by the presence of vertical speed.

Figure 6.12: The estimation and prediction errors of the ground distance between aircraft

Figure 6.13: The estimation and prediction errors of the relative aircraft position.

Figure 6.16: The subject confidence level in estimating the aircraft pair situation.

entirely in a horizontal plane, referred to as a horizontal violation.

Judgement errors from the questionnaire runs, and the confidence level, are depicted in Figures 6.12-6.16. The errors in estimation and projection of the ground distance between aircraft, relative to their true distance, is presented in Figure 6.12. The position error, which is the distance between the indicated blue aircraft position and its true position, relative to the true distance between aircraft is shown in Figure 6.13. The error in estimating the time aircraft will reach equal altitude, relative to the true time is presented in Figure 6.15. The confidence level is grouped by the type of the subjects in Figure 6.16.

6.7.1 The Effect of Viewing Parameter Optimization

Viewing parameter optimization has a significant effect on the spatial awareness in complex images, where complexity is brought forth both by graphical detail and aircraft load. This is reflected in the increase in the number of hits and the decrease in the number of false-alarms in Figures 6.8,6.9 for enhanced graphics at high load, and the lack of a significant difference in the other cases. This reiterates the findings of the first experiment, and establishes them for motion. From Figure 6.10 one sees that optimization affects both when vertical speed is present, and when it isn't, which may be attributed to the distribution of the optimization effort between the different geometrical features of aircraft pairs.

The judgement and prediction errors of the subjects decreases with opti-

Figure 6.14: The error in estimation of the angle between aircraft trajectories.

Figure 6.15: The error in the estimation of the time-to-equal-altitude of the aircraft pair.

mized viewing parameters, as is visible in Figures 6.12-6.14. The confidence level increases with the optimization (Figure 6.16) indicating the subjects preferred optimized displays.

6.7.2 Plan-View Vs. Perspective

The optimized perspective displays provided better spatial awareness than plan-view displays in complex scenes (enhanced graphics and high load), as is evident in the higher number of hits in Figure 6.8. The advantage of graphically enhanced perspective displays in the hit performance is apparent regardless of the presence of vertical speed, see Figure 6.10, which might be attributed to their tendency to naturally allocate the user's attention to "prime" areas of the screen. As image complexity is reduced, this advantage apparently diminishes, until in the plain graphics at low loads, plan-view displays have the upper-hand over the perspective displays; possibly due to the accurate read-outs of the data-tags which loose effectiveness in congested displays.

In the horizontal plane, Figures 6.12-6.14 shows that judgement and prediction performance of the plan-view display were comparable to those of optimized perspective ones, disproving a common criticism that ground distance judgments cannot be done as accurately in perspective displays as in plan-view ones (Gregory 1977; Wickens, Haskell and Harte 1989). Also, Figure 6.16 suggests that the subjects felt as confident to perform judgements in optimized perspective displays as in the plan-view ones. Moreover, the perspective display's natural integration of both the vertical dimension and the graphical enhancements, can explain the better judgment performance of the time-to-equal-altitude, see Figure 6.15.

6.7.3 The Effect of Graphical Enhancements

The effect of graphical enhancements was significant in improving the hit and false-alarm performance of both perspective displays in Figures 6.8,6.9. In particular their effect is visible on vertical violations, see Figures 6.10,6.11, for horizontal violations, a significant effect was observed only for optimized displays. The higher effectiveness of the graphical enhancements in perspective displays is an indication of their natural interpretation.

The graphical enhancements were added in order to improve the ability to comprehend aircraft velocities. Indeed, they had no effect on present distance

and position estimation, see Figures 6.12,6.13. Alas, neither did they have an effect on the ability to project future distances and positions, nor on trajectory angle estimation (Figure 6.14). An improvement was only observed in the unoptimized (plain) perspective displays, which was comparable to the improvement achieved from optimizing the viewing parameters without enhancing the graphics. The high variability of the time-to-equal-altitude judgements in the plain graphics (Figure 6.15) as opposed to that of the enhanced one, indicates that the time estimation was only possible with graphical enhancements.

A general improvement with graphical enhancements in the confidence level of both subject groups was observed, see Figure 6.16. The mixed results suggest that for an effective spatial perception enhancement, a comprehensive optimization scheme should be applied, selecting both the viewing parameters and the level and type of graphical enhancements.

6.7.4 The Effect of Aircraft Load

While the plan-view and unoptimized perspective displays suffered a high performance loss when increasing aircraft load, see Figures 6.8,6.9, optimized perspective displays suffered the least. This was particularly noticed in the probability of hit for vertical violations in the plain-graphics plan-view displays (Figures 6.10,6.11), indicating the difficulty of detecting vertical speeds from altitude data-tags in congested displays. Since in the questionnaire runs the queried pair was highlighted, aircraft load had no effect on these runs.

6.7.5 The Effect of Subject Background

A significant difference between the two subject groups was observed only in the confidence level, Figure 6.16, the ATC generally being more confident. The confidence of both subject groups in optimized perspective displays was higher than it was in unoptimized ones, and equal to that of plan-view displays. However, a higher confidence in plain-graphics plan-view displays was reported by the ATC subjects, which can be explained by their experience with similar displays. The lack of a significant difference between the subject groups in other measures can be explained by the nature of the experiment task which required spatial awareness, but no ATC knowledge.

6.8 Summary

The experiment proved the feasibility of designing perspective displays for the enhanced perception of geometrical features. Optimized perspective displays yielded a better spatial awareness than unoptimized ones, and an overall better performance than plan-view displays in complex scenes. A novel method, using a map drawn on a tablet, was used to register the answers of the subjects in judging the spatial state of aircraft pairs. These judgements showed that optimized perspective displays enable equivalent performance in ground plane estimations and superior performance in the vertical plane. This improvement was also reflected in a user preference towards optimized perspective displays.

Figure 6.17: The full scale tablet map, as it was used in the experiment.

7 Discussion

A novel method for rating the quality of both static and dynamic perspective displays has been developed. This rating depended on the viewing parameters, the scene contents, and the task to be performed. In order to compute the rating, a perception model was used, which predicted the errors that a human observer would make in reconstructing the three-dimensional layout from the two-dimensional perspective display. The rating was used in a GA-based optimization engine, to select the "best" viewing parameter setting for the given display, scene contents, and task. The optimized display was shown to effectively support spatial awareness.

The effectiveness of optimizing the viewing parameters of the display has been clearly demonstrated in a series of experiments, in which the subjects performed a spatial awareness task. The performance in this task, when using optimized perspective displays, was superior to that of plan-view displays or that of perspective displays with manually chosen viewing parameters. Hence, the

importance of properly selecting the viewing parameters was clearly shown.

An active display concept was conceived that automatically selects the optimal viewing parameters to enhance the spatial perception of geometrical features. A novel GA based optimization approach was presented to exploit the time continuity of the optimized objective function by locating local maxima early in their "evolution". This approach was possible due to a new niche formation technique which combined two existing techniques together with a new definition for a geometrically meaningful distance function.

Graphical enhancements, such as velocity vectors or altitude predictors, increase their effectiveness in optimized perspective displays. Optimizing both the format of the graphical enhancements, as well as the viewing parameters, may be mutually supportive in improving the level of spatial awareness obtained from the display.

From the components of the active display system, our research was mainly concerned with the scene analyzer and the optimization engine with its spatial perception model. The other two components, i.e. the decision algorithm and the path planner, are subjects for further development. GAs may be further exploited for the active display problem by adapting them to Pareto optimization,[†] which may enable agility in adjusting to changes in user priorities. A complete active display system should also consider and incorporate the auxiliary data that will be available in future ATC scenarios. Having a fully-fledged active display system will enable to test the concept in full-task experiments and in their natural ATC environment. A promising area for future development is the cockpit displays of traffic information. It constitutes an ideal application for optimized perspective displays, as a good measure of spatial awareness is obtainable from an "overview" glance of the image, requiring numerical data only for adding accuracy.

Optimized perspective displays are not limited to air traffic control applications. They may be used in other applications that require the presentation and

---

[†] Pareto optimization deals with vector functions $f_1(x)...f_n(x)$. Solution x is said to be dominant over solution y if $f_i(x) \geq f_i(y)$ for all $i = 1...n$, and there is a j such that $1 < j < n$ and $f_j(x) > f_j(y)$. The dominant front is the set of all solutions in the solution space which are not dominated by any other solution, thus it is the set of all points that maximize the scalar function $\sum w_i f_i(x)$ for all possible values of w_i. GAs are suitable for Pareto optimization since the dominance relation which is defined only between two points, can easily be incorporated to a selection scheme.

the processing of complex, time-varying spatial information. Computer aided design (CAD), teleoperator navigation and database visualization are but a few examples of such applications. The active display system can enable a *task centered* interface, by having the user indicate the geometrical features of interest, rather than manually manipulating the viewing parameters for best viewing these features. An optimization, aimed at deliberately directing judgment errors in desired directions, can create visual effects, e.g. an architect may desire a building to appear taller than it is, a photographer may desire a room to appear more spacious than it is. Thus, the spatial perception model can accurately control visual effects that are currently being practiced by artists using rules of thumb. It is well-known that the "classical" artists have introduced deliberate geometrical distortions in the projection geometry, in order to communicate the spatial features in the best possible way. Apart from the projection geometry, the spatial perception model can easily be extended to include texture, lighting, shaded light, occlusion and stereoscopic vision. By including the placement of light sources and object textures in the objective function, by appropriate formulation of the perception model, the optimization of more complex images, than the ones used in this study, will be possible. The basic principle, used in the perception model, namely, that the variance of the judgement errors is proportional to the curvature of the matching cost function, can be easily extended to different scene attributes, like texture and shading.

This study proved perspective displays with viewing parameter optimization to be an effective means of delivering spatial information. Their design, and implementation in the experiments, was aimed at specific spatial tasks; the first being the identification of certain spatial criteria (separation violation) whether present or future, the second task being judging spatial features. These tasks were designed to measure the spatial awareness and comprehension of the subjects, and do not necessarily reflect the actual task air traffic controllers perform. Controllers are required to process spatial data and perform spatial tasks, such as the rerouting of air traffic in the airspace. The crux of the task is the requirement to work with spatial information. By configuring the objective function to accurately reflect the operator's task, optimized perspective displays provide the the optimal way in which the required spatial information can be transferred to the operator.

A Rating Image Quality

An objective function was constructed to rate the image quality for use in the optimization engine. This function is based on the predicted errors in the perception of geometrical features the scene analyzer picked for enhancement. The scene analyzer and the construction of the objective function are presented.

A.1 The Scene Analyzer

The scene analyzer identifies events in the airspace the geometrical features of which are to be enhanced. An event in the airspace is defined as an aircraft situation in the airspace which is of interest to the controller. Several studies were conducted to establish what information controllers use when dealing with air traffic (Amaldi *et al.* 1996; CourteixKherouf 1998). This was achieved by interrupting controllers and requesting them to reconstruct from memory the airspace they were just controlling. The following spatial data were found to be noted by controllers:

- Aircraft that are expected to violate a separation criteria if let to continue on their current course.

- Aircraft that are headed for the same destination.

The scene analyzer assigns weights to the geometrical features passed to the optimizer. The weights reflect the relative importance of each geometrical feature and are used by the optimizer in constructing the objective function. The following events were identified:

1) Aircraft that are currently in violation of a separation criteria. These were assigned the maximum weight $w = w_{max}$. The separation criteria used in selecting these pairs was larger than the critical criteria controllers must guarantee between aircraft or the criteria the subjects were shown in the experiments. Typically, a separation twice the critical separation was used to ensure that aircraft that are not in violation, but are at a separation close to the critical one, are also enhanced by the optimizer.

2) Aircraft that are predicted to violate a separation criteria in the future if let to continue in their current course. The same separation criteria as in the previous aircraft pair type was used. The weight was assigned to these pairs relative to the time-to-violation (TTV) of the pair, such that if the TTV is zero, the weight assigned is w_{max} and when the TTV is equal 20 minutes, the weight assigned is w_1. Values of TTV beyond 20 minutes, which is considered the "horizon", are ignored.

3) Aircraft headed to the same destination. The weight on these aircraft was assigned to be the lowest, and was set relative to the time-to-destination (TTD) of the first aircraft in the destination and the time separation at the destination (TSD) between the aircraft, such that it was always between w_2 and w_3. If TTD+TSD/2 is greater than 30 minutes, or TSD/2 is greater than TTD, this aircraft pair is ignored.

This arrangement is suitable for a TRACON control, for En-Route control, the time constants will have to be adjusted to the different time scales of En-Route control. The weights w_{max}, w_1, w_2, w_3 were set by experimentation, and by observing the relations: $w_{max} > w_1$, $w_1 \cong w_2$ and $w_2 > w_3$.

A.2 The Objective Function

The optimization engine maximizes the objective function, which reflects the quality of viewing of the relevant geometrical features in the image. The score q_i of the i^{th} aircraft pair is

$$q_i = \max\left(\frac{d}{\sigma_d}, M_d\right) + \max\left(\frac{a}{\sigma_a}, M_a\right) + \max\left(\frac{v}{\sigma_v}, M_v\right) + \max\left(\frac{c}{\sigma_c}, M_c\right) \qquad (A.1)$$

where d, a, v and c are the actual ground separation, altitude separation, relative ground velocity and relative climb/descent rate between aircraft, respectively; $\sigma_d, \sigma_a, \sigma_v$ and σ_c are the perception errors of d, a, v, c, respectively, and are calculated by the perception model. M_d, M_a, M_v, M_c are constants.

Eq. (A.1) states the viewing quality score is relative to a saturated sum of the reciprocals of the relative errors in each geometrical feature of the aircraft pair. The constants $M_{(*)}$ indicate the saturation level and set a desired error level in each feature. The aim of the saturation function is to distribute the optimization effort to decreasing the errors in all features. A saturation was suitable for this task due to the singular nature of the errors, which for different viewing parameters can vary from infinitesimally small values to infinitely large ones.

The image score Q is set according to the doctrine that the image quality is the average viewing quality of all aircraft pairs, as long as the worst viewed aircraft pair has a reasonable viewing quality, yielding:

$$Q = \min\left(\begin{array}{c} \frac{1}{N}\sum_1^N w_i q_i \\ \lambda \min(q_i) \end{array}\right). \qquad (A.2)$$

According to Eq. (A.2), the image quality is taken as the smaller of the weighted average of the aircraft pair viewing qualities or λ times the quality of the worst viewed aircraft pair. The weights w_i are set by the scene analyzer and are normalized such that their sum is unity. This scheme was selected in order to avoid situations in which the optimizer generates viewpoints which are focused on a single aircraft pair, raising its viewing score in expense of the other pairs in the scene.

B Analyses of Niching Schemes

B.1 The Sharing Parameter

The workings of the sharing principle depends on the parameter σ of the sharing function in Eq. (3.3), which is the maximum distance between solutions necessary to form as many niches as there are peaks in the function. Deb *et al.* (1989) suggest a method for choosing the parameter σ under the assumption that there are q peaks in the function, each occupying an equal space in the p-dimensional solution hyperspace, yielding

$$\sigma = \frac{r}{\sqrt[p]{q}} \tag{B.1}$$

with r being the radius of the p-dimensional solution hyperspace, and depends on the maximum and minimum value of each parameter.

An analysis can provide insight to the effectiveness of the sharing scheme

and the meaning of the sharing parameter σ. Assume that there are two peaks in the function, with unshared fitness values of f_1 and f_2 respectively. Assume further that the population consists of n members and is divided into two proportions n_1 and n_2 on each peak, respectively. To maintain a stable population on each peak, the shared fitness of the solutions on each of the peaks must be equal, yielding:

$$\frac{f_1}{n_1 + n_2 s(1,2)} = \frac{f_2}{n_2 + n_1 s(1,2)} \qquad (B.2)$$

or

$$\frac{f_1}{f_2} = \frac{n_1 + n_2 s(1,2)}{n_2 + n_1 s(1,2)} = \gamma \qquad (B.3)$$

where $\gamma \geq 1$ is the ratio between the peak fitness, $s(1,2)$ is the sharing function value between members on peak 1 and peak 2. Eq. (B.3) can be rearranged to yield:

$$\frac{n_2}{n} = \frac{1 - \gamma s(1,2)}{(1 - s(1,2))(1 + \gamma)}. \qquad (B.4)$$

To have a non-zero population on the second peak (with lower fitness value), the numerator of Eq. (B.4) must be greater than zero, which together with Eq. (3.1) on page 35 provides a condition on the distance between the two peaks:

$$d(1,2) \geq \left(1 - \frac{1}{\gamma}\right)\sigma. \qquad (B.5)$$

Therefore, there is a lower bound on the distance between two peaks $d(1,2)$, which is dependent on the sharing distance parameter σ and the ratio between peaks. Note that if the peaks are equal, $\gamma = 1$ any distance can exist between the two peaks, in Genotype sharing this means even a one bit difference is sufficient to create niches. The current analysis was performed without making any assumption on the structure of $d(1,2)$. A similar analysis that uses the assumption that $d(1,2)$ is a genotype distance was carried out by Deb *et al.* (1989), who wrongfully used this analysis as an argumentation against genotype sharing.

B.2 On Deterministic Crowding

Deterministic crowding is a variation on a former algorithm, named crowding (De Jong 1975). The idea behind crowding is to maintain population diversity by diluting clustered members. In crowding, two members of the population are selected according to their fitness, mated, and the offspring are inserted back to the population in the following manner: A *crowding factor* (CF) number of members are selected at random, and each offspring replaces the member of the CF sample which has the smallest distance to it. The CF is a preset number that can be selected in the range between two to the size of the population, and is usually taken as a small number to reduce the computation effort.

Crowding per se is not considered as an effective means for locating and maintaining multiple peaks in functions. Consider a function with multiple peaks $p_1...p_k$, and consider there is a member from peak p_1 awaiting insertion to the population. Crowding chooses CF candidates from the population and replaces the closest one. Assuming no errors in the comparison, an element of peak p_1 will be replaced if it is among the CF candidates. As long as all peaks other than p_1 together contain CF or more elements, these other peaks are vulnerable to loss of members. When peak p_1 has $n - CF + 1$ members, one of its members will always be in the CF sample, and the other peaks (at most $CF - 1$ such peaks) will not lose their members to p_1. Sampling and replacement errors[†] during the initial evolution of the GA usually cause the crowding GA to maintain but two peaks.

Deterministic crowding gains its effectiveness as a niching method due to the observation that usually the closest member in the population to an offspring is one of its parents (Mahfoud 1992). By deterministically selecting the parent with the smaller distance, deterministic crowding greatly reduces the number of sampling and replacement errors. When the objective function has skewed peaks, as in the active display optimization, usage of phenotype distance reintroduced replacement errors when choosing which offspring should be considered to replace which parent.

[†] A replacement error is defined as the event in which a member from one peak replaces a member from another peak. Sampling error is the event in which none of the candidates for replacement belong to the same peak as the member awaiting insertion, which necessarily leads to a replacement error.

C References

Adams, M. J., Tenny, Y. J. and Pew R. W. **1995**.
Situation Awareness and the Cognitive Management of Complex Systems. *Human Factors 37*, 1, 85–104.

Amaldi, P., Boudes, N. and Cellier, J. M. **1996**.
Information Processing Strategies in Air Traffic Control. In *Proceedings of 20th Congress of ICAS the International Council of the Aeronautical Sciences*. Italy, 2036–2044.

Ames Jr., A. **1925**.
The Illusion of Depth from Single Pictures. *Journal of the Optical Society of America 10*, 137–148.

Azuma, R., Daily, M. and Krozel, J. **1996**.
Advanced Human-Computer Interfaces for Air Traffic Management and Simulation. In *Proceedings of the AIAA Flight Simulation Technologies Conference*. San Diego, CA, 656–666.

Barfield, W. and Rosenberg, C. **1995**.
Judgments of Azimuth and Elevation as a Function of Monoscopic and Binocular Depth Cues Using a Perspective Display. *Human Factors 37*, 1, 173–181.

Burnett, M. S. and Barfield, W. **1991**.
Perspective Versus Plan View Air traffic Control Displays. Survey and Empirical Results. In *Proceedings of the Human Factors Society 35th Annual Meeting*. San Francisco, CA, 87–91.

Chiruvolu, R., Hwang, V. S. and Sheridan T. B. **1992**.
Virtual Display Aids for Teleoperation. In *Proceedings of SPIE - The International Society for Optical Engineering*. Boston, MA, 299–310.

Courteix-Kherouf, S. **1998**.
   Complementary Use of Datalink and Voice Frequency Communication Between Pilots and Air Traffic Controllers in a Simulated Situation. In *Proceedings of HCI-AERO'98, International Conference on Human-Computer Interaction in Aeronautics*. Montreal, CA, 111–116.

Cunningham, H. A. and Pavel, M. **1991**.
   Target Axis Effects Under Transformed Visual-Motor Mappings. In *Pictorial Communication in Virtual and Real Environments*, Ellis S. R., Kaiser M. and Grunwald A. J. *Eds*, Taylor & Francis, 283–294.

Cutting, J. E. **1991**.
   On the Efficacy of Cinema, or what the Visual System Did Not Evolve To Do. In *Pictorial Communication in Virtual and Real Environments*, Ellis S. R., Kaiser M. and Grunwald A. J. *Eds*, Taylor & Francis, 486–495.

Das, H., Sheridan T. B., Slotine and JeanJeaques, E. **1989**.
   Kinematic Control and Visual Display of Redundant Teleoperators. In *Proceedings of the IEEE International Conference of Systems, Man and Cybernetics*. Vol. 2. 1072–1077.

De Jong, K. A. **1975**.
   An Analysis of the Behavior of a Class of Genetic Adaptive Systems. Dissertation Abstracts International, University of Michigan, 5140.

De Jong, K. and Sarma J. **1992**.
   Generation Gaps Revisited. In *Proceedings of the International Conference on Genetic Algorithms and Their Application*.

Deb, K. and Goldberg, D. E. **1989**.
   An Investigation of Niche and Species Formation in Genetic Function Optimization. In *Proceedings of the International Conference on Genetic Algorithms and Their Applications*. 42–50.

Dorighi, N. **1996**.
   Private communication. NASA, Ames research center.

Ebert, R. E. and MacMillan, A. G. **1985**.
   Misperception of Small Cars. In *Trends in Ergonomics/Human Factor II*, Ebert R.E. and Ebert C. G. *Eds*, Elsevier Science Publishers, North Holland, Netherlands, 33–39.

Ellis, S. R., Tharp, G. K., Grunwald, A. J. and Smith, S. **1991**.
   Exocentric Judgments in Real Environments and Stereoscopic Displays. In *Proceedings of the Human Factors Society 35th Annual Meeting*. Vol. 2. San Francisco, CA, 1442–1446.

Ellis, S. R. and McGreevy, M. W. **1987**.
   Perspective Traffic Display Format and Airline Pilot Traffic Avoidance. *Human Factors 29*, 4 (August), 371–382.

Endsley, M. R. **1995**.
   Towards a Theory of Situation Awareness in Dynamic Systems. *Human Factors 37*, 1, 32–64.

Eshelman, L. J., Caruna, R. and Schaffer, J. D. **1989**.
   Biases in the Crossover Landscape. In *Proceedings of the Third International Conference on Genetic Algorithms*. 10–19.

Etter, D. M., Hicks, M. J. and Cho, K. H. **1982**.
   Recursive Adaptive Filter Design Using an Adaptive Genetic Algorithm. In *IEEE International Conference on Acoustics, Speech and Signal Processing*. Vol. 2. 635–638.

FAA **1996**.
   Aviation System Capital Investment Plan. Federal Aviation Administration, Washington D.C.

Farber, J. and Rosinski R. R. **1978**.
   Geometric Transformations of Pictured Space. *Perception 7*, 269–282.

Games, P. A. and Howell J. F. **1976**.
   Pairwise Multiple Comparison Procedures with Unequal n's and, 113–125.

Goldberg, D. E. **1989**.
   Genetic Algorithms in Search, Optimization and Machine Learning. Addison Wesley.

116

Goldberg, D. E. and Richardson, J. **1987.**
    Genetic Algorithms with Sharing for Multimodal Function Optimization. In *Genetic Algorithms and Their Applications: Proceedings of the Second International Conference on Genetic Algorithms.* 41–49.

Goldberg, D. E. and Samtani, M. P. **1986.**
    Engineering Optimization via Genetic Algorithm. In *Proceedings of the Ninth Conference on Electronic Computation.* 471–482.

Goldberg, D. E. and Smith R. E. **1987.**
    Nonstationary Function Optimization Using Genetic Algorithms with Dominance and Diploidy. In *Genetic Algorithms and Their Applications: Proceedings of the Second International Conference on Genetic Algorithms.* 59–68.

Gregory, R. L. **1977.**
    Eye and Brain. Weidenfeld & Nicolson, London.

Gregory, K. T. and Ellis, S. R. **1990.**
    The Effects of Training on Errors of Perceived Direction in Perspective Displays. NASA, National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA, -90081.

Grunwald, A. J. **1984.**
    Tunnel Display for Four-Dimensional Fixed-Wing Aircraft Approaches. *Journal of Guidance, Control and Dynamics 7,* 3 (May-June), 369–377.

Grunwald, A. J. **1996.**
    Advanced Interactive Display Formats for Terminal Area Traffic Control. Technion - Israel Institute of Technology.

Grunwald, A. J. and Ellis, S. R. **1988.**
    Interaction Orbital Proximity Operations Planning System. NASA Ames Research Center, Moffett Field, CA, 2839.

Grunwald, A. J., Ellis, S. R. and Smith, S. **1988.**
    Mathematical Model for Spatial Orientation from Pictorial Perspective Displays. *IEEE Transactions on Systems, Man and Cybernetics 18,* 3, 425–437.

Hacisalihzade, S. S., Stark, L. W. and Allen, J. S. **1992.**
    Visual Perception and Sequences of Eye Movement Fixations: A Stochastic Modeling Approach. *IEEE Transactions on Systems, Man and Cybernetics 22,* 3 (May-June), 474–481.

Holland, J. H. **1975.**
    Adaptation in Natural and Artificial Systems. The University of Michigan Press.

Hollstien, R. B. **1971.**
    Artificial Genetic Adaptation in Computer Control Systems. In *Dissertation Abstracts International,* University of Michigan, 1510B, 3.

Jasek, M., Pioch, N. and Zeltzer, D. **1995.**
    Enhanced Visual Displays for Air Traffic Control Collision Prediction. In *Proceedings of ICAT'95 - International Conference on Artificial Reality and Tele-Existence.* Makuhari, Chiba, Japan, 635–640.

Kim, W. S., Ellis, S. R., Tyler, M. E., Hanaford, B. and Stark, L. W. **1987.**
    Quantitative Evaluation of Perspective and Stereoscopic Displays in Three-Axis Manual Tracking Tasks. *IEEE Transactions on Systems, Man and Cybernetics 17,* 1 (January-February), 61–72.

La Gournerie, J. **1859.**
    Traite de Perspective lineaire contenant les traces pour les tableaux plans et courbes, les bas-releifs, et les decorations theatrales, avec une theorie des effects de perspective. Dalmont et Dunod, Paris.

Lumsden, E. A. **1983.**
    Perception of Radial Distance as a Function of Magnification and Truncation of Depicted Spatial Layout. *Perception and Psychophysics 33,* 177–182.

Mahfoud, S. W. **1992.**
    Crowding and Preselection Revisited. In *Parallel Problem Solving From Nature, 2,* R. Manner and B. Manderick *Eds,* Elsevier Science Publishers, Amsterdam, 27–36.

McGreevy, M. W. and Ellis, S. R. **1986.**
The Effect of Perspective Geometry on Judged Direction in Spatial Information Instruments. *Human Factors 28,* 4 (August), 439–456.

Miller, B. L. and Shaw, M. J. **1995.**
Genetic Algorithms with Dynamic Niche Sharing for Multimodal Function Optimization. University of Illinois at Urbana-Champaign, Urbana IL, 95010.

Minga, A. K. **1986.**
Genetic Algorithms in Aerospace Design. In *AIAA Student Conference.* Huntsville, AL.

Mittelstaedt, H. **1991.**
Interactions of Form and Orientation. In *Pictorial Communication in Virtual and Real Environments,* Ellis S. R., Kaiser M. and Grunwald A. J. *Eds,* Taylor & Francis, 377–389.

NRC **1997.**
Flight to the Future: Human Factors in Air Traffic Control. Wickens, C. D. and Mavor, A. S. and McGee J. *Eds,* Panel on Human Factors in Air Traffic Control Automation, National Research Council; National Academy Press, Washington, D.C.

O'Brien, J. V. and Wickens, C. D. **1997.**
Free Flight Cockpit Displays of Traffic and Weather Effects of Dimensionality and Data Base Integration. In *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting.* Albuquerque, NM, 18–22.

Papoulis, A. **1984.**
Probability, Random Variables and Stochastic Processes. Chapter 4.

Pelegrin, M. **1998.**
Air Transport System Structure Around 2010. In *Proceedings of HCI-AERO'98 International Conference on Human-Computer Interaction in Aeronautics.* Montreal, CA, 7–14.

Perry, T. S. **1997.**
In Search of the Future of Air Traffic Control. *IEEE Spectrum 8,* (August), 19–35.

Purdy, W. C. **1960.**
The Hypothesis of Psychophysical Correspondence in Space Perception. General Electric Advanced Electronics Center, Ithaca NY, 6056.

RTCA **1995.**
RTCA Task Force 3 Final Report on Free Flight Implementation. RTCA, Inc., Washington, D.C.

Rogowitz, B. E. and Allebach, P. J. **1990.**
Human Vision and Electronic Imaging: Models, Methods, and Applications. Rogowitz, B. E. and Allebach, P. J. *Eds,* Belingham, WA, USA.

Sakaino, H. and Sonehara, N. **1996.**
Adaptive Kalman Filter Model for Stable Perception of Visual Space Unaffected by Eye Movement. In *Proceedings of the 1996 IEEE Signal Processing Society Workshop.* NJ, USA, 492–501.

Sarter, N. and 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.

Sedgwick, H. A. **1991.**
The Effects of Viewpoint on the Virtual Space of Pictures. In *Pictorial Communication in Virtual and Real Environments,* Ellis S. R., Dorighi N. and Grunwald A. J. *Eds,* Taylor and Francis, Chapter 30.

Shaviv, G. E. and Grunwald, A. J. **1997.**
Advanced Perspective Display Formats for Air Traffic Control. In *Proceedings of the 37th Israeli Annual Conference on Aerospace Sciences.* Israel, 257–265.

Sheridan, T. B. **1987.**
Supervisory Control. In *Handbook of Human Factors,* G. Salvendy *Eds,* Wiley, New York.

Spears, W. M. **1994.**
Simple Subpopulation Schemes. In *Proceedings of the 1994 Evolutionary Programming Conference.*

Spears, W. M. and De Jong, K. **1991**.
: An Analysis of Multi Point Crossover. In *Foundations of Genetic Algorithms*. 301–315.

Stein, E. S. **1992**.
: Air Traffic Control Visual Scanning. Federal Aviation Administration, Atlantic City, NJ, -9216.

Van Gent, R. N., Hoekstra, J. M. and Ruigrok, R. C. J. **1998**.
: Free Flight with Airborne Separation Assurance. In *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*. Montreal, CA, 63–70.

Whitley, D. **1989**.
: The Genitor Algorithm and Selection Pressure: Why Rank-Based allocation of Reproductive Trials is Best. In *Proceedings of the International Conference on Genetic Algorithms and Their Application*. 116–121.

Wickens, C. D. **1992**.
: Engineering Psychology and Human Performance. Harper Colliners Publishers, Chapter 4, 116–166.

Wickens, C. D., Haskell, I. and Harte, K. **1989**.
: Ergonomic Perspective of Flight Path Displays. *IEEE Control Systems Magazine.*

Wickens, C. D., Merwin, D. H. and Lin, E. L. **1994**.
: Implications of Graphic Enhancements of Scientific Data: Dimensional Integrability, Stereopsis, motion and mesh. *Human Factors 36*, 1, 44–61.

Zacharias, G. L., Caglayan, A. K. and Sinacori, J. B. **1985**.
: A Model for Visual Flow-Field Cueing and Self-Motion Estimation. *IEEE Transactions on Systems, Man and Cybernetics 15*, 3 (May-June), 385–389.

Zacharias, G. L., Miao, A. X. and Warren, R. **1995**.
: Multistage Integration Model for Human Egomotion Perception. *Journal of Guidance, Control and Dynamics 18*, 5 (September-October), 937–944.