Vision-based navigation using multi-rate feedback from optic flow and scene reconstruction

A. R. Arvai
amandar7@ufl.edu
J. J. Kehoe
jjk@ufl.edu
R. Lind
ricklind@ufl.edu
Aerospace Engineering
University of Florida
Florida, USA

ABSTRACT

Vision-based control is being aggressively pursued for autonomous systems. Such control is particularly valuable for path planning to achieve mission objectives like target tracking and obstacle avoidance. This paper presents a multi-rate strategy that utilises a fast-rate optic flow approach and a slow-rate scene reconstruction approach. The vehicle uses scene reconstruction to generate an accurate map for path planning; however, optic flow is used to avoid obstacles while the scene reconstruction is computed. A switch element is used in the feedback path to determine whether information relating to the reconstructed map or the optical flow should be used for navigation. The resulting controller is able to generate flight trajectories and perform obstacle avoidance within a computational cost which is reasonable given performance demands and computational resources on a wide range of aircraft. A simulation demonstrates the performance of an aircraft that uses the multi-rate controller to avoid an obstacle which is only observed after a turn. Essentially, the fast-rate optic flow indicates the presence of the obstacle during the time that slow-rate scene reconstruction is being performed. The resulting flight path is able to follow mission objectives and avoid a collision.

1.0 INTRODUCTION

Autonomy is an enabling technology for many missions envisioned for future generations of unpiloted air vehicles (UAV). Some missions will use micro and small UAV platforms that will be difficult for a human operator to observe and command. Some missions will require rapid decision making based on more targets than a human can process. Finally, missions may utilise swarms of UAV systems for which the workload would be too immense for a human to control.

Vision-based feedback can be a critical component to autonomy. Inertial sensors, such as gyros and accelerometers, provide information about the state of the vehicle; however, they do not provide any information about the environment. Vision can be used to extract information, such as location of obstacles, about the flight space. As such, vision-based feedback provides a mechanism for autonomous navigation to maneuver through unknown or uncertain environments. Even urban environments can be considered given that appropriately sized vehicles, having wingspan less than 2ft and turn radius less than 20ft, can respond quickly to information obtained from cameras which often have field of view between 70 degrees and 140 degrees.

The concept of visual-servo, or vision-based, control is appropriate for tasks such as guidance and navigation. Researchers in robotics have been particularly active in this area along with more recent applications in aerospace and manufacturing. Most techniques share some commonality; namely, a sequence of image processing and vision processing are performed to extract information which is then analysed to make an optimal control decision. The basic unit of information from an image is a feature point which indicates some pixel of particular interest due to, for example, colour or intensity gradient near that pixel. Among the techniques that utilise feature

Paper No. 3402. Manuscript received 15 January 2008, revised version received 10 March 2011, accepted 17 March 2011.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

points, the approaches related to optic flow and scene reconstruction are of direct interest to this paper.

Optic flow refers to the velocity of the feature points across images taken at different times[1,2]. Motion towards an obstacle will result in large flow that increases as the distance to that obstacle decreases. In this way, optic flow is used for control by indicating regions that may contain close obstacles. The approach has been used for various applications ranging from guide robots for the visually impaired[3] to micro air vehicles[4,5,6,7]. Some successful missions that demonstrated optic flow include obstacle avoidance for ground vehicles[3,8,9], road navigation[10], terrain following for helicopters[4], and the autonomous landing of aerial vehicles[8]. A form of optic flow using probability distributions of correspondences between feature points has also been used for obstacle avoidance[11].

The computation of optic flow has received considerable attention in the context of enabling autonomy. A fast algorithm was derived using intensity gradients as features for pattern matching in combination with a brightness constraint[3]. The computation time is further decreased by compressing the peripheral parts of optical flow to reduce the amount of input data[9]. Sensors with extremely fast rates of computing optic flow have been generated using inspiration from the biology of flying insects[5].

Alternatively, the techniques using scene reconstruction have only limited experience with control. The approach for scene reconstruction attempts to estimate the relative locations, in an inertial coordinate system, of features points in a set of images. The underlying concept of scene reconstruction is actually straightforward to understand by considering projective geometry of aircraft positions and feature points; however, the mathematics required to accurately compute optimal solutions to the problem are quite complex[12]. The techniques of structure from motion are perhaps the most matured approach to scene reconstruction[13]. A recent approach uses adaptive learning to estimate a detailed terrain from a set of feature points[14]. In each case, a strategy for path planning enables navigation by choosing a trajectory to avoid the estimated locations of obstacles. Scene reconstruction has been used for a variety of ground vehicles[15,16] and aerial simulations[14,17,18].

Each type of vision-based control, optic flow and scene reconstruction, has strengths and weaknesses. Optic flow presents only inferences about obstacle locations but can be computed extremely fast; conversely, scene reconstruction presents highly-detailed maps of obstacle location but requires an intense computational effort. Obviously autonomous aircraft will require both fast rate as the vehicle flies through cluttered environments and low computational cost due to limited processing speeds of on-board hardware. The issue of processing increases dramatically for small aircraft that are designed to operate within urban environments but have only payload capability of ounces that must be distributed between sensors and batteries and telemetry along with the computational engine.

This paper introduces a multi-rate strategy for navigation that combines optic flow and scene reconstruction. A slow loop operates on the scene reconstruction and, while the computations are being performed, a fast loop uses optic flow to avoid obstacles. In this way, the aircraft uses optic flow to ensure collisions will not occur during the scene reconstruction. An optimality condition for navigation is presented for optic flow that considers only the component from translation motion to avoid obstacles. A switching strategy is then implemented to enable multi-rate feedback and blend the path planning from scene reconstruction and optic flow.

A series of simulations are presented to demonstrate the methodology. An aircraft maneuvers through a cluttered environment using only optic flow, only scene reconstruction, and the multi-rate strategy. The approaches using only one type of vision-based control are not able to satisfy the mission objectives whereas the multi-rate controller is able to navigate through the obstacles and still reach the target waypoint.

Also, the contribution of this paper lies mostly with optic flow and its use in multi-rate control; consequently, the discussion focuses on optic flow and multi-rate. A basic outline of scene reconstruction is presented, along with references that provide extensive details, because the associated computations are quite complicated and well beyond the scope of this paper[12,19].

2.0 FEATURE POINTS

This paper will use feature points as the foundation for any vision-based feedback. These points generally correlate to items in the environment of special significance. Some examples of items that often constitute feature points are corners, edges and light sources. Such feature points can provide information about the overall object in the sense that a set of corners can outline a building. Feature points do not necessarily provide enough information to completely describe an environment but, in practice, they usually provide sufficient information for obstacle avoidance.

A camera effectively maps the 3D environment onto a 2D image plane. This image plane is defined as the plane normal to the camera's central axis located a focal length, f, away from the origin of the camera basis. The image plane has coordinates bounded by the field of view such that the vertical component satisfies $\mu \in [\underline{\mu}, \bar{\mu}]$ and the horizontal component satisfies $\nu \in [\underline{\nu}, \bar{\nu}]$.

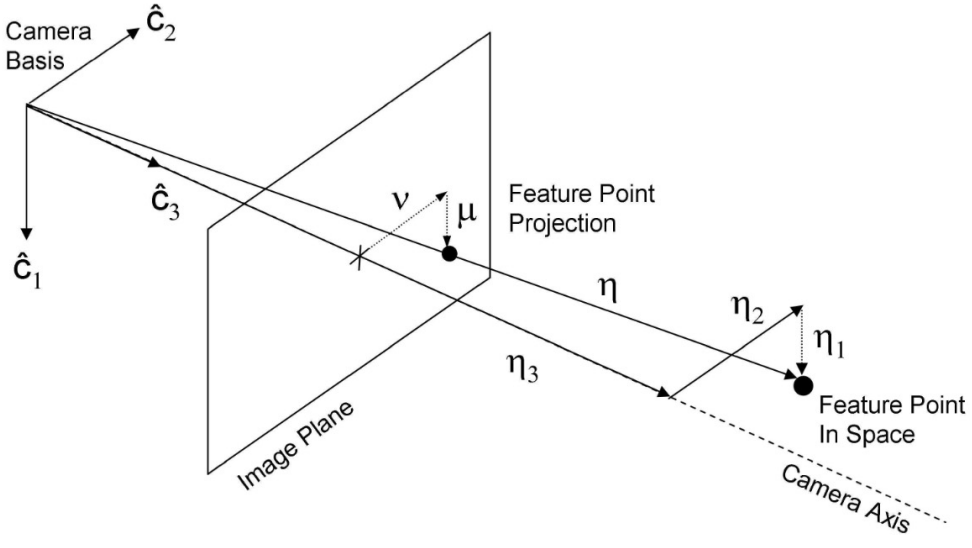


Figure 1. Mapping from environment to Image Plane.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

The mapping provided by the camera is described in Fig. 1. The vector, η, between the camera and the environmental item is decomposed into components along and perpendicular to the camera axis or line of sight. The projection onto the image plane is then offset from the center of the image by μ and ν as a function of these components of η and the focal length of *f*.

The co-ordinates of the feature point in the image plane are given by Equation (1) and Equation (2). These equations utilise a standard pin-hole camera model. Also, the camera origin is assumed to coincide with the camera lens.

$$\mu = f\frac{\eta_1}{\eta_3} \quad \ldots (1)$$

$$\nu = f\frac{\eta_2}{\eta_3} \quad \ldots (2)$$

The extraction of feature points is critical to the success of the vision-based control described in this paper. Such extraction begins with identifying the features in the image using, among others, techniques based on edge detection, colour distribution, intensity variation or basic differentiation of image properties[20]. The extraction continues as additional images are measured by tracking a feature point throughout these images using, among others, techniques for segmented correlation and Kalman filtering[15,16,21]. Extracting feature points is not trivial and degrades with increasing noise and increasing motion of the aircraft although the techniques are routinely used for robotic systems.

3.0 VISION-BASED CONTROL STRATEGIES

3.1 Scene reconstruction

3.1.1 Concept

Scene reconstruction attempts to actually estimate the inertial co-ordinates of obstacles within the environment observed by the camera. Essentially, this reconstruction seeks to create a virtual map of the environment. Such a map would obviously be highly beneficial for any autonomous system because a flight path could be computed with confidence. The techniques of scene reconstruction can roughly be built from elements of epipolar geometry, stereopsis, structure from motion, and terrain mapping.

Epipolar geometry is a fundamental computation used for computer vision[22,23,24]. The basic concept notes that a feature point which is observed on multiple images must satisfy a geometric relationship. In other words, the location of an obstacle can be determined by triangulation using the lines of sight from two camera positions. Of course, this concept requires the positions and attitudes of the cameras to be known.

The reconstruction of a complete scene, still assuming the positions and attitudes of cameras are known, can be achieved by stereopsis[25,26]. The technique uses correspondences between multiple feature points across multiple frames to reconstruct an entire scene. Stereopsis is basically stereo vision that fuses features observed from images to estimate the depth of those features. The mathematics are obviously built upon the epipolar constraint but the ability to robustly account for a large number of features adds significant complexity[27].

Relaxing the assumptions of known camera characteristics dramatically increases the difficulty of scene reconstruction. Structure from motion is introduced to simultaneously estimate both obstacle properties and camera properties[13]. Nonlinear relationships between the camera and obstacles are derived whose solutions correspond to the scene reconstruction. The relationships can be approximated as affine for small motions[28] and solved using a nonlinear least-squares factorisation[29]. The full problem based on projective geometry can be addressed using, for example, conics and iterative solutions[30].

Finally, recent advances consider scene reconstruction explicitly for vision-based control. These advances use an adaptive and multi-level method based on wavelets to provide various resolutions of block geometry[31]. The method actually uses learning theory in association with the adaptive wavelets such that information from multiple flights is used to methodically generate a virtual map that relates all feature points. Such a method can operate in near real-time speed and, consequently, provide information for obstacle avoidance during flight.

3.1.2 Strategy

Autopilot design is relatively straightforward using scene reconstruction. The computer vision provides a virtual map of the environment so the autopilot merely needs to find an optimal path through that map. Essentially, the problem becomes a classical version of trajectory generation given a map of the environment. The real difficulty arises with the need for on-line processing and the fast updates which are particularly important for maneuvering through cluttered environments.

Among the variety of approaches for path planning, a rapid algorithm has been directly developed for vision-based control with scene reconstruction[32]. The algorithm for this approach solves sequences of optimisation problems based on local characteristics within the environment. The local solutions are then connected throughout the flight to generate a path through the environment. The approach does not yet guarantee stability; however, simulated demonstrations indicate good performance at a reasonable computational cost.

A variety of sensors can be used for scene reconstruction. A video camera is the traditional sensor and the entire process can actually operate using only vision information[33]; however, the quality of scene reconstruction is obviously improved with additional information. The approach is directly augmented to incorporate sensors such as inertial gyros and laser range finders.

3.2 Optic flow

3.2.1 Concept

Optic flow represents the velocity of feature points as projected onto the image plane[8]. Essentially, the relative motion between an aircraft and an obstacle results in movement of the associated feature points if that motion is not purely along the line of sight of the camera. The optic flow can thus provide inferences about the location of obstacles during flight.

The velocities associated with optic flow result from time derivatives of the feature-point positions in Equation (1) and Equation (2). The resulting expressions, given in Equation (3) and Equation (4), describe the motion in the two-dimensional image plane;

$$\dot{\mu} = f\left(\frac{\eta_3\dot{\eta}_1 - \eta_1\dot{\eta}_3}{\eta_3^2}\right) \quad \ldots (3)$$

$$\dot{\nu} = f\left(\frac{\eta_3\dot{\eta}_2 - \eta_2\dot{\eta}_3}{\eta_3^2}\right) \quad \ldots (4)$$

The concept of optic flow can be demonstrated for the artificial scenario in Fig. 2. An aircraft is flown between two obstacles and images are recorded during the flight. In this case, the flight path is a northerly trajectory at constant altitude. The obstacles lie equal distances to the east and west on either side of this trajectory; however, the obstacle to the east is closer to the initial location of the

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

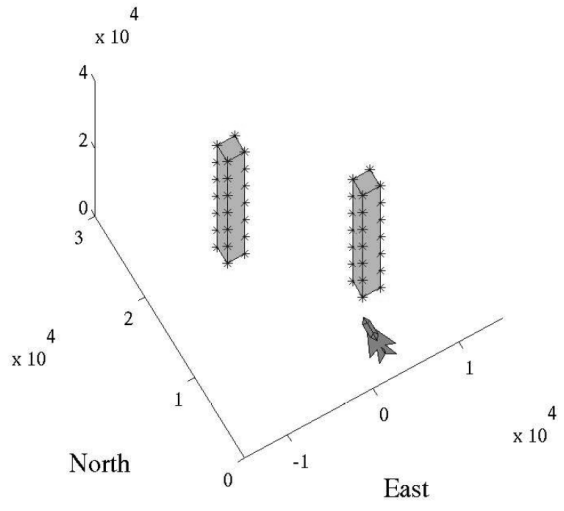


Figure 2. Scenario of aircraft and obstacles.

aircraft. Also, a set of feature points on each obstacle are denoted by the star symbols.

Vision-based feedback is generated by simulating straight and level flight along the flight path. A pair of images, shown in Fig. 3, represent the feature points taken 0·04 sec apart. Several characteristics are notable for the feature points in each image. In particular, the feature points on the right side of each image are positioned farther from the vertical centerline, are spaced farther apart along the vertical direction, and are reduced in number when compared to the feature points on the left side. These characteristics uniformly result from the relative location of the camera to the obstacles. The obstacle to the east of the vehicle, which corresponds to the right side of each image, is closer to the camera so only a reduced portion of that obstacle lies within the field of view. The effect of distance, as described in Equation (1) and Equation (2), results in these visual characteristics.

The optic flow of these feature points is given in Fig. 4. The vectors are actually scaled by a factor of 25 to increase their visibility and ease of presentation. These vectors clearly demonstrate the mathematical properties expected from Equation (3) and Equation (4) for such a scenario. The feature points tagged to the obstacle to the east, or to the right in the image plane, have larger optic flow due to the reduced distance. Similarly, the variation in causes points above or below the horizontal centerline to have a vertical velocity in the image plane.

This example is somewhat simplistic but clearly demonstrates a relationship between optic flow and obstacle detection. The magnitude and direction of optic flow can be used to infer information about relative position between obstacles and a camera. This argument is particularly strong when the obstacles are fixed in position but the general concept also applies to scenarios of moving obstacles.

Also, the measurement of optic flow is obviously affected by feedback noise. Uncertainty in optic flow results from errors in camera calibration, discrete pixelisation, feature point localisation, and integration with other sensors. These uncertainties and noise will cause incorrect estimates of optical flow; however, realistic levels of noise in vision systems have been shown to cause relatively minor variations in vision-based feedback[34,35].

3.2.2 Translational optic flow

The general expressions for optic flow, as given in Equation (3) and Equation (4), describe the velocity for any time-varying position of a feature point. A specific form of these expressions can be derived for the situation of an aircraft flying relative to fixed features. In this case, the expressions in Equation (5) and Equation (6) relate the optic flow in terms of aircraft states and camera parameters. Standard nomenclature is used for aircraft states such as p for roll rate, q for pitch rate, r for yaw rate, u for forward velocity, v for side velocity, and w for vertical velocity in the expressions.

$$\dot{\mu} = f\left[\left(q + p\nu + q\mu^2 + r\mu\nu\right) + \left(\frac{u\mu - w}{\eta_3}\right)\right] \quad \ldots (5)$$

$$\dot{\nu} = f\left[\left(r - p\mu + r\nu^2 + q\mu\nu\right) + \left(\frac{u\nu + v}{\eta_3}\right)\right] \quad \ldots (6)$$

An evaluation of optic flow in terms of aircraft states is particularly valuable for obstacle avoidance. This optic flow, as evidenced by Equation (5) and Equation (6), can essentially be decomposed into translational and rotational parts. The primary risk of collision between an aircraft and a fixed obstacle is associated with translational motion; therefore, the rotational component of optic flow is not critical for this particular task. Aerial vehicles often use roll maneuvers to turn the aircraft but it remains the translation during the turn, rather than the roll itself, that is the fundamental concern for collisions.

The translational component of optic flow, which is of primary interest for obstacle avoidance, can be computed using measurements of the aircraft velocities. The desired values, $\dot{\tilde{\mu}}$, and $\dot{\tilde{\nu}}$, result from subtraction of the appropriate combinations of aircraft states and feature point co-ordinates from the measured optic flow as given in Equation (7) and Equation (8).

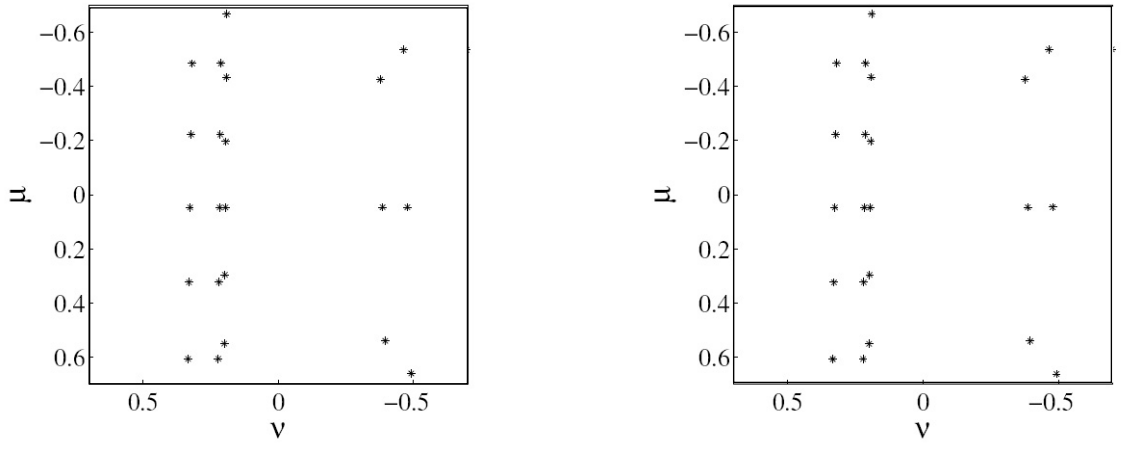


Figure 3. Feature point positions in image plane at timesteps 1 (Left) and 2 (Right).

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

$$\dot{\tilde{\mu}} = \dot{\mu} - f\left(q + p\nu + q\mu^2 + r\mu\nu\right) \quad \ldots (7)$$

$$\dot{\tilde{\nu}} = \dot{\nu} - f\left(r - p\mu + r\nu^2 + q\mu\nu\right) \quad \ldots (8)$$

A comparison between optic flow and translational optic flow is demonstrated using the scenario in Fig. 2. In this new case, the aircraft flies due north but continuously rolls during the flight. The feature points in Fig. 5 correspond to images taken 0·04sec apart while the aircraft has a roll rate of 2·608 rad/sec, a pitch rate of 0·0126 rad/sec, and a yaw rate of –0·1176 rad/sec.

The optic flow in Fig. 6 describes the motion of the feature points in Fig. 5. The optic flow is clearly influenced by the rotational velocities of the aircraft such that the feature points are primarily rotating in the image plane. Conversely, the translational component of optic flow shows considerably smaller magnitudes and different directions than the total optic flow. The translational component in Fig. 6 actually agrees quite well with the optic flow in Fig. 4 that resulted from purely translational motion of the aircraft along this same flight path. Note the translational component has been scaled by a factor of 25 for ease of presentations to ensure the vectors are visible.

3.2.3 Strategy

A control strategy is formulated to navigate the aircraft away from regions of high-magnitude optic flow or, alternatively, regions of nearby obstacles. Essentially, navigation commands are generated to minimise a cost function. This cost function is derived to account for characteristics, such as magnitude and location, of feature points and their associated optic flow. The regions in the image plane where this cost function is low indicate the desired directions toward which the aircraft should fly.

The magnitude of optic flow is obviously a metric of importance because of its inferred relationship to obstacle distance. As such, the cost function utilises this velocity value to penalise obstacles which may present imminent threat. The optic flow is computed by tracking feature points, of which N are assumed, through a sequence of images. The n^{th} point has co-ordinates of μ_n and ν_n and translational optic flow of and with velocity given in Equation (9).

$$V_n = \sqrt{\dot{\tilde{\mu}}_n^2 + \dot{\tilde{\nu}}_n^2} \quad \ldots (9)$$

Also, the cost function penalises flight paths that pass close to potential obstacles. The two-dimensional image plane does not

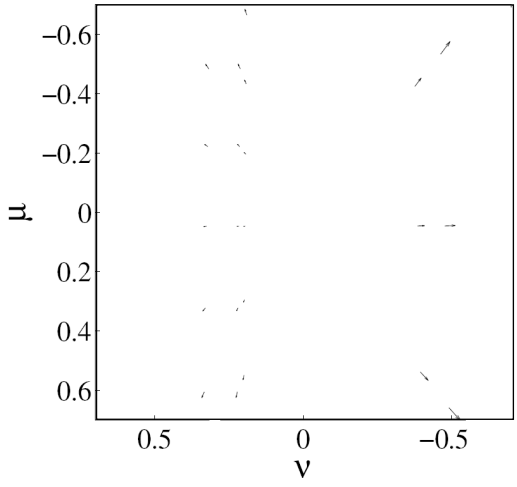


Figure 4. Optic flow from Fig. 3.

provide direct information for this penalty; however, the feature points can be related to such information. A point in the image plane correlates to a line of sight, as shown in Fig. 1, so a flight path is safe if no obstacles lie along the line of sight associated with the flight path. As such, the cost function will consider distance, $d_n(\mu,\nu)$, from every point on the image plane to the n^{th} feature point as in Equation (10).

$$d_n(\mu,\nu) = \sqrt{(\mu - \mu_n)^2 + (\nu - \nu_n)^2} \quad \ldots (10)$$

A cost function can be computed by relating the parameters for velocity and distance as in Equation (11). This cost function essentially normalises the quartic of the translational optic flow by the distance to the associated feature point. The quartic provides a disproportionately high cost for feature points with high velocities that lie close to the flight path.

$$J(\mu,\nu) = \sum_{n=1}^{N} \frac{V_n^4}{d_n(\mu,\nu)} \quad \ldots (11)$$

A flight path is chosen as the line of sight along which the cost function in Equation (11) is minimal. This line of sight correlates to the image co-ordinates, μ_{opt} and ν_{opt}, with lowest cost as given in Equation (12).

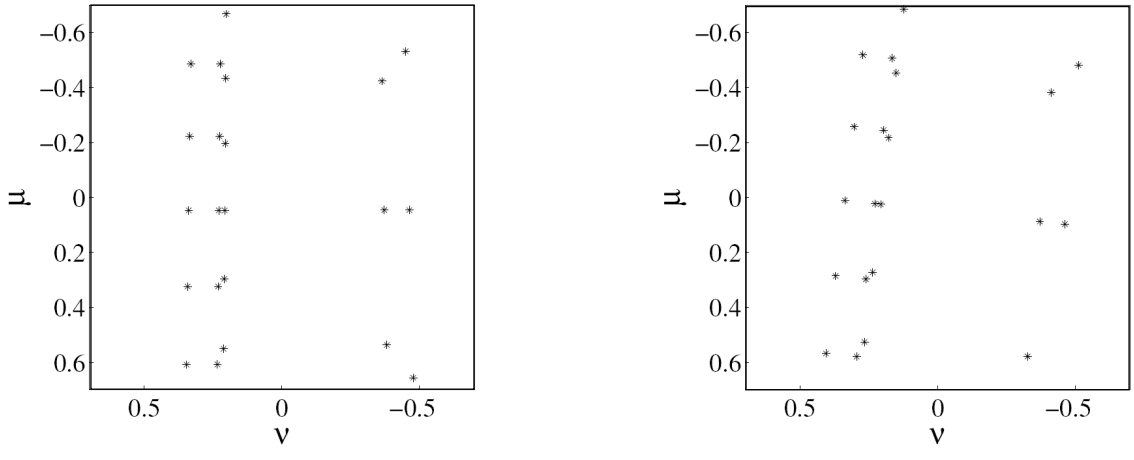


Figure 5. Feature point locations at timesteps 1 (left) and 2 (right).

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

$$\left[\mu_{opt}, \nu_{opt}\right] = \arg \min_{\substack{\mu\in[\underline{\mu},\bar{\mu}] \\ \nu\in[\underline{\nu},\bar{\nu}]}} J(\mu,\nu) \qquad \ldots (12)$$

Navigation commands are generated to steer the aircraft toward the optimal image co-ordinates or, alternatively, the optimal line of sight. The angles for heading and pitch attitude are computed, assuming angle of attack is given by α and angle of sideslip is negligible, as in Equation (13) and Equation (14.) These optimal commands, Δ_ψ and Δ_θ, can be directly extracted from Fig. 1 by including the angle of attack.

$$\Delta_\psi = -\mathrm{Tan}^{-1}\left(\frac{\nu_{opt}}{f}\right) \qquad \ldots (13)$$

$$\Delta_\theta = \alpha - \mathrm{Tan}^{-1}\left(\frac{\mu_{opt}}{f}\right) \qquad \ldots (14)$$

The control strategy is demonstrated using the scenario presented in Fig. 2. Simulated flights generated the optic flow in Fig. 4 or the translational optic flow in Fig. 6 based on feature points. The cost function, $J(\mu,\nu)$, associated with these optic flows is shown in Fig. 7. This cost function indicates the regions to the west or straight ahead are the safest. As such, the navigation commands of $\Delta_\psi = -0{\cdot}61$ rad and $\Delta_\theta = 0$ are computed using Equation (13) and Equation (14). A standard inner-loop command tracker is implemented to perform the trajectory following[36]. The commands indicate the strategy is able to steer the aircraft away from obstacles to minimise the threat of collision.

Variations to the cost function are easily considered by introducing relative weights to the velocity or distance in Equation (11). Such weights will raise the cost to emphasise a penalty on optic flow of feature points or the distance from those feature points to the current heading. Several such weights were considered with the resulting formulation in Equation (11) chosen to work reasonably well across a large range of aircraft velocities and simulated environments. Also, this current cost assumes that the aircraft is moving faster than any features in the image so any high levels of optic flow are easily related to close obstacles.

4.0 MULTI-RATE CONTROL

Certainly both vision-based control strategies, scene reconstruction and optic flow, present enabling technologies for autonomous aircraft but each approach also has limitations. Scene reconstruction can provide accurate maps of obstacles such that the resulting flight paths are confidently safe; however, the computational cost for scene reconstruction greatly limits its utility for on-board applications. Optic flow can actually be computed using hardware with extremely fast rates; however, the information provides only a rough inference about obstacles so the resulting flight paths are generated with some level of caution. As such, a clear trade-off is observed between the level of information for decision making and computation time required to provide that information.

A multi-rate controller is proposed that takes advantages of the strengths of each type of vision-based feedback while limiting the effects of their respect weaknesses. The fundamental scheme will utilise scene reconstruction in a slow loop while the optic flow runs in a fast loop. Each loop computes navigation commands, Δ_ψ and Δ_θ, so a switching element determines which set of commands are actually passed to the maneuvering controller as in Fig. 8.

The switching element is critical to effectively using both types of vision-based control. This switch is implemented based on a trigger that reflects the total magnitude of the optic flow field in the image as given in Equation (15). In this case, the trigger uses the quartic of the optic flow for each feature point to disproportionately emphasise the largest vectors which are assumed to relate to the closest obstacles.

$$F = \sum_{n=1}^{N} V_n^4 \qquad \ldots (15)$$

The law governing the switch compares F to some threshold, σ, to determine which type of vision-based control should command the aircraft. Essentially, a large value of optic flow means either the vehicle is close to an obstacle and/or the vehicle is viewing a densely-cluttered environment. Either situation presents a need for fast-rate control so the optic flow should dominate. The actual value of σ is chosen based on feature density, aircraft velocity, and maneuver aggressiveness which will vary for each situation. Essentially, the threshold is chosen to reflect capabilities of the vehicle through iterative simulations. The resulting switch is given in Equation (16).

$$\begin{bmatrix}\Delta_\psi \\ \Delta_\theta\end{bmatrix} = \begin{Bmatrix}\text{output of optic flow} & : \text{if } F > \sigma \\ \text{output of scenere construction} & : \text{else}\end{Bmatrix} \qquad \ldots (16)$$

The multi-rate controller should enable the aircraft to operate better than a single-rate controller using either optic flow or scene reconstruction. Such control will be particularly beneficial when entering a cluttered environment in which some obstacles are hidden from view at various times during the flight. Consider an

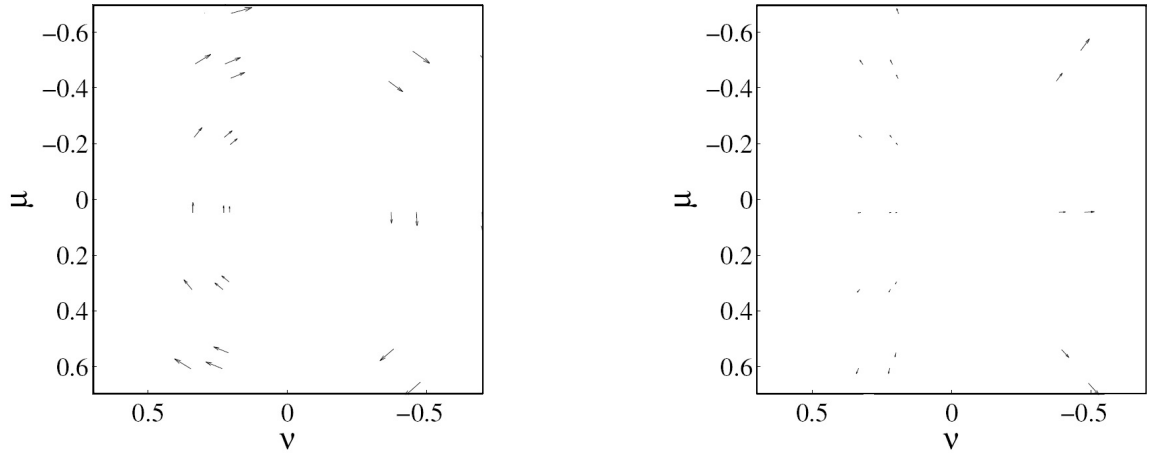


Figure 6. Optic flow (Left) and scaled-by-25 optic flow with rotational components removed (right).

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

image taken at $t = 0$ during which the aircraft is turning. The controller using scene reconstruction will not respond to features in that image until processing is finished. The vehicle will remain turning during this time even though new features, including potentially unseen obstacles, will come into the view of field as the vehicle changes position and heading. The optic flow, which will evaluate the image considerably faster, will able to quickly respond to these newly-viewed obstacles. In this way, the multi-rate controller allows the high-resolution path planning associated with scene reconstruction but maintains obstacle avoidance during the computation time by utilising optic flow.

Some aspect of the switching logic and associated trigger may need to be tuned based on specifics to each problem. The trigger essentially relates sensitivity to when optic flow indicates a potential obstacle is too close. Such sensitivity should be computed by relating the turn radius of the vehicle to the optic flow anticipated when an obstacle approaches that turn radius. The use of a discrete switch, opposed to a blend, also relates a choice of how quickly the transition of structure-from-motion to optic flow must occur given the velocity towards obstacles.

Also, the stability with this formulation can not be proven for the closed-loop system in Fig. 8. That system contains nonlinear flight dynamics, nonlinear image processing, and multi-rate control commands which are, collectively, outside the regime in which theoretical statements about stability can not be made. A series of several dozen scenarios involving different aircraft and different obstacles have been simulated and uniformly the controller was able to generate a reasonable flight path.

5.0 EXAMPLE

5.1 Mission

A set of simulations are presented to demonstrate the vision-based controllers. These simulations utilise a common scenario to compare the ability of each controller to steer an aircraft past obstacles. In this case, the scenario is a set of obstacles, of which one is large and seven are small, populating an area through which an aircraft must fly. The controllers are evaluated on the flight path each generates by considering time to reach a target waypoint and distance to obstacles during the flight.

The traversing of an urban environment by a micro air vehicle is a mission of particular interest to the University Florida. The scenario of these simulations is intended to mimic that mission; however, models of the flight dynamics of a micro air vehicle are not matured. Instead, the simulations utilise a high-fidelity

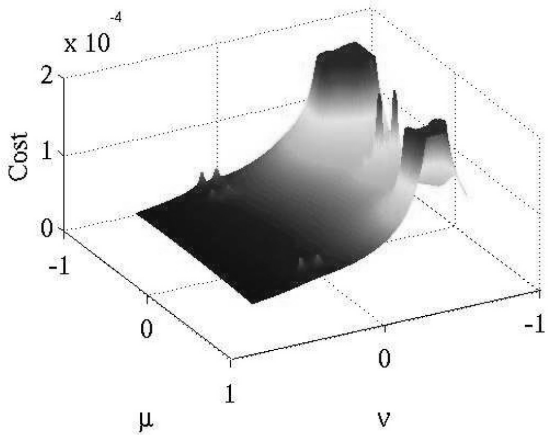


Figure 7. Cost Function for Fig. 3 and Fig. 6.

simulation of an F-16[37]. The F-16 is not an ideal model for this simulation but the nonlinear dynamics provide an estimate of representative flight dynamics. The simulations are able to show the controller performance on a realistic model so, like any controller design, the gains will need to be tuned for a different model but the basic approach can at least be validated.

Also, the model includes nonlinear open-loop dynamics and a maneuvering autopilot. This autopilot includes a stability augmentation system to smooth the flight path and a control augmentation system to track attitude commands. An additional outer-loop controller allows the vehicle to follow waypoints using the strategy from scene reconstruction[36].

The mission chosen for these simulations highlights the advantages of multi-rate control. A set of obstacles is situated behind a long wall such that they are initially hidden from view. The mission objective seeks a path to the northwest. After the vehicle flies north around the wall, the hidden obstacles become visible. The objective is thus to maneuver between these obstacles while navigating toward the indicated destination located to the west. The aircraft is also commanded to maintain a constant altitude.

The simulation represents responses from slow-rate control using only scene reconstruction, high-rate control using only optic flow, and multi-rate control using both scene reconstruction and optic flow. The scene reconstruction updates every 120 seconds while the optic flow updates every 0·01 seconds to reflect processor speed on current payloads of micro air vehicles[38].

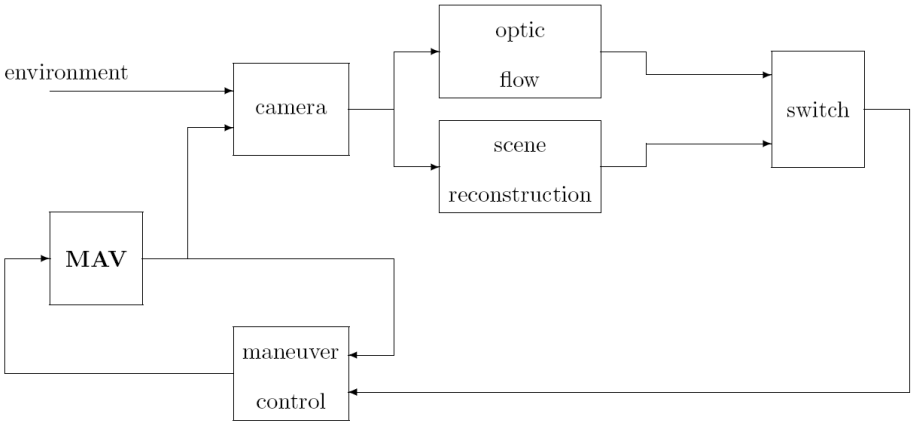


Figure 8. Closed-loop system with multi-rate control.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

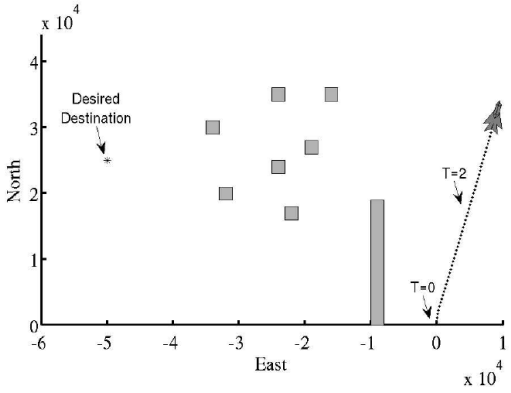


Figure 9. Flight path using optic flow.

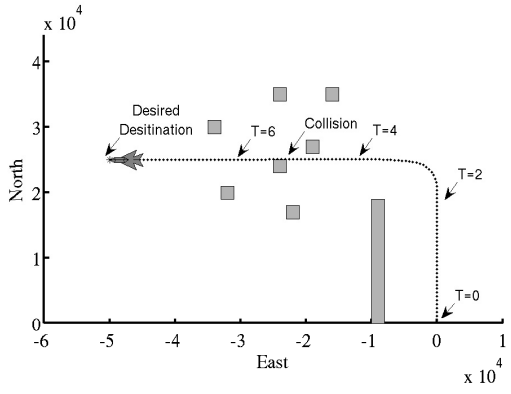


Figure 10. Flight path using scene reconstruction.

5.2 Control based on optic flow

A flight path is simulated in response to control using only optic flow. This controller uses the cost function in Equation (11) to analyse the optic flow of the image plane at each timestep. It then selects a line of sight which is assumed to be the safest based on that cost function. The appropriate commands for heading and pitch are subsequently generated and passed to the maneuvering controller.

The resulting flight path is shown in Fig. 9. The vehicle initially sees the optic flow to the left and consequently turns right. It therefore avoids the obstacles but does not fly toward the northwest; instead, it steers toward the northeast. This controller therefore does not accomplish the mission objectives.

The path in Fig. 9 is not entirely unexpected given the simplistic nature of the implemented controller. Several approaches could be used that are much more advanced than simply steering toward least flow; however, this simplistic controller is used because it can operate at extremely high rate. As such, the results are not indicative of a limitation in optic flow as much as they are indicative of a limitation in high-rate optic flow. This example demonstrates that optic flow can certainly provide obstacle avoidance but its reactive nature does not directly allow for waypoint navigation.

5.3 Control based on scene reconstruction

The representative response to a controller using only scene reconstruction is also simulated. The simulation did not actually implement the entire scene reconstruction; instead, results from ongoing research at the University of Florida and the University of South Carolina are used to mimic the anticipated performance. The flight path is actually indicative of a controller that had perfect feature point extraction, perfect feature point tracking, perfect state estimation, perfect terrain mapping, and perfect path planning. Obviously such an assumption is unrealistic but it does provide a benchmark to demonstrate a potential measure of performance.

The fight path from this assumption of perfect analysis is presented in Fig. 10. This result is only understood by noting that the updates are computed at the green points noted on the path. The vehicle records an image at the initial time $t = 0$, which shows only the wall. Assuming that the feature point analysis takes 2 minutes, or runs at 0·008Hz, a resulting trajectory for this image is not created until point $t = 2$min. More obstacles have come into view at approximately $t = 3$min; however, a trajectory resulting from the new image will not be computed until $t = 4$min. By this time, the aircraft has already collided with an obstacle.

The flight path in Fig. 10 is meant to indicate a possible problem with scene reconstruction for aircraft. The approach has been used with considerable success for some systems, such as ground vehicles and even helicopters, that can stop and/or hover; however, an aircraft is continually moving forward so the computation delay can cause significant problems. This example demonstrates that scene reconstruction can indeed be used for path planning but collisions can result from unseen obstacles when the vehicle turns out of the field of view used to generate the path.

5.4 Multi-rate control

The multi-rate control scheme is also introduced to the simulation. This simulation uses the low-rate scene reconstruction to compute a path but watches for impending obstacles using the high-rate optic flow while following that path. Should a threat be detected, optic flow navigates the aircraft to safety. In this case, the optic flow operates at 100·0Hz while the scene reconstruction operates at 0·008Hz.

The initial commands from the two schemes, as shown in Figs 9 and 10, try to steer the vehicles in different directions. So, the threshold, $\sigma = 0{\cdot}0002$, in Equation (16) is used such that the optic flow controller does not change the vehicle path until the magnitude of flow passes a critical value. The value of σ is chosen by noting the relationship between vehicle dynamics and optic flow of obstacles. Essentially, the value was iteratively selected by choosing a threshold that allows the vehicle to approach obstacles but still turn away as the vehicle nears a distance to those obstacles which is comparable to the turn radius. In this way, the vehicle follows the scene reconstruction but avoids obstacles.

Figure 11 presents the results of the integrated multi-rate controller. The path initially follows the commands from the scene reconstruction because the optic flow does not indicate any imminent threats. The command switches to optic flow at $t = 3{\cdot}8$min when the optic flow notes the aircraft is approaching a small obstacle. The vehicle continues using optic flow to avoid obstacles until $t = 4{\cdot}5$min at which time command was transferred back to scene reconstruction. The controller was again switched to optic flow at $t = 5{\cdot}7$min to avoid a final obstacle and continued to follow commands based on optic flow until $t = 6{\cdot}4$min. The controller based on scene reconstruction commands the vehicle during the remainder of the simulation until the final waypoint is reached.

The controller twice switched into obstacle avoidance using optic flow during the simulation. The trigger for such switching, as described in Equation (15), occurs when the amount of optic flow shown in Fig. 12 exceeds $\sigma = 0{\cdot}0002$. Clearly this weighted metric of optic flow increases dramatically as the vehicle approaches the obstacles.

The actual commands for navigation generated by optic flow utilise a cost function for the image. The values of cost function, given in Equation (11), are shown in Fig. 13 to correspond to $t = 4$min and in Fig. 14 to correspond to $t = 6$min. Each plot indicates a disproportionately high cost for close obstacles due to the

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

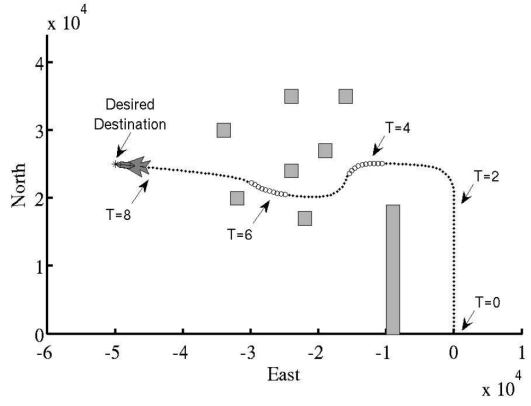


Figure 11. Flight path using multi-rate control.

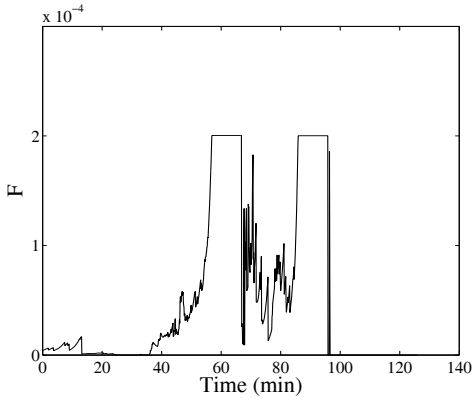


Figure 12. Weighted magnitude of optic flow.

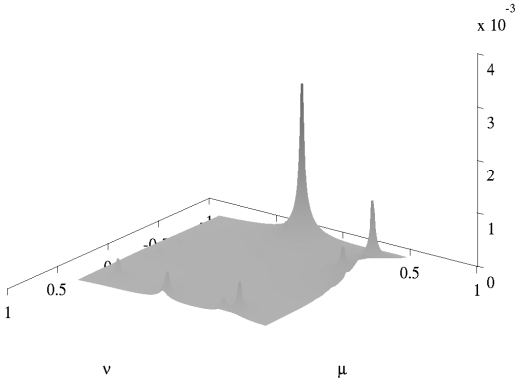


Figure 13. Cost function.

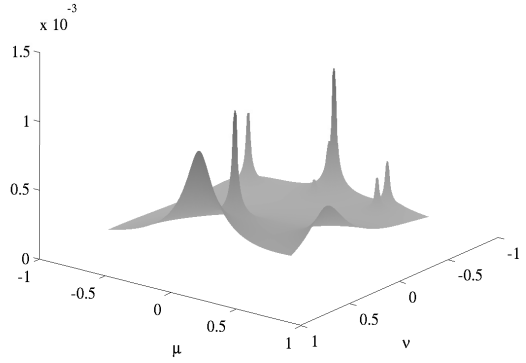


Figure 14. Cost function.

quartic term in the cost function. In each case, the feature points of obstacles near the vehicle are associated with extremely high cost. The flight path of the vehicles, shown in Fig. 11, is thus chosen by steering towards the minimum point in the image plane associated with the lowest cost.

5.5 Noise sensitivity

The sensitivity to noise is evaluated through a series of simulations. In each case, the locations in the image plane of each feature point are varied due to the addition of some magnitude-bounded random perturbation. The optic flow is then computed from these noisy locations of features points and used for guidance commands.

The flight path of the aircraft is given in Fig. 15 for noise ranging from 0·5 pixels to 5·0 pixels. The flight path is nearly identical using noise-free measurements and 0·5-pixel noise. A slight detour to the south is taken when using measurements with noise levels between 1·0 pixel and 3·0 pixel. The addition of noise of 5·0 pixels causes the vehicle to detour to the north.

The approach is thus shown to have reasonable sensitivity to noise. Noise levels of 1·0 pixel are often assumed for image processing and the flight path is essentially unchanged for noise values less than 1·0 pixel. A small variation is encountered with uniform paths resulting for noise up to 3·0 pixel. Even a high level of 5·0 pixels of noise is able to be used and still a flight path is generated that avoids obstacles while flying through the city.

6.0 CONCLUSION

This paper demonstrates that a multi-rate scheme is valuable for vision-based control. Both optic flow and scene reconstruction are powerful approaches; however, each has relative suitability for on-board autonomous vehicles. The scene reconstruction can clearly provide path plans but the computation time is excessive for fast vehicles operating within dense obstacles. The optic flow can be implemented on hardware that is fast but provides limited information. So, the high-rate optic flow merged with the low-rate scene reconstruction provides a beneficial structure which adopts the strengths of each approach while avoiding the weaknesses. A simulation shows how an autonomous vehicle can utilise both types of vision-based feedback to avoid obstacles while still pursuing a mission objective. This multi-rate approach represents a powerful technology for enabling vision-based navigation with particular emphasis on operating in densely-populated regions such as micro air vehicles flying in urban environments.

7.0 ACKNOWLEDGMENTS

This work was supported jointly by the Air Force Research Laboratory and the Air Force Office of Scientific Research under F49620-03-1-0381 with Johnny Evers, Neal Glassman, Sharon Heise, and Robert Sierakowski as project monitors.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

REFERENCES

1. LUCAS, B. and KANADE, T. An iterative image registration technique with an application to stereo vision, Proceedings of the DARPA Image Understanding Workshop, 1981, pp 121-130.
2. HORN, B.K. and SCHUNK, B.G. Determining optical flow, *Artificial Intelligence*, 1981, **17**, pp 185-203.
3. SONG, K.-T. and HUANG, J.-H. Fast optical flow estimation and its application to real-time obstacle avoidance, IEEE International Conference on Robotics and Automation, 2001, **3**, pp 2891-2896.
4. RUFFIER, F., VIOLLET, S., AMIC, S. and FRANCESCHINI, N. Bio-inspired optical flow circuits for the visual guidance of micro-air vehicles, International Symposium on Circuits and Systems, May 2003, **3**, pp 846-849.
5. BARROWS, G.L. Future visual microsensors for mini/micro-UAV Applications, 7th IEEE International Workshop on Cellular Neural Networks and their Applications, July 2002, pp 498-506.
6. BARROWS, G.L., CHAHL, J.S. and SRINIVASAN, M.V. Biomimetic visual sensing and flight control, 2002 Bristol UAV Conference, April 2002.
7. BARROWS, G.L. and NEELY, C. Mixed-mode VLSI optic flow sensors for in-flight control of a micro air vehicle, *SPIE*, July 2000.
8. YOUNG, G.-S., HONG, T.-H., HERMAN, M. and YANG, J.C.S. Obstacle detection for a vehicle using optical flow, Intelligent Vehicles Symposium, June 1992, pp 185-190.
9. BARATOFF, G., TOEPFER, C., WENDE, M. and NEUMANN, H. Real-time navigation and obstacle avoidance from optical flow on a Space-variant map, IEEE International Symposium on Intelligent Control, September 1998, pp 289-294.
10. CASTRO, A.P.A., SILVA, J.D. and SIMONI, P.O. Image based autonomous navigation with fuzzy logic control, International Joint Conference on Neural Networks, July 2001, **3**, pp 2200-2205.
11. MERRELL, P.C., LEE, D.J. and BEARD, R.W. Obstacle Avoidance for Unmanned Air Vehicles Using Optical Flow Probability Distributions, SPIE Optics East, *Robotics Technologies and Architectures, Mobile Robot XVII,* October 2004, Volume 5609-04.
12. MA, Y., SOATTO, S., KOSECKA, J. and SASTRY, S.S. *An Invitation to 3D Vision*, Springer-Verlag Publishers, New York, NY, USA, 2004.
13. ULLMAN, S. The Interpretation of Visual Motion, MIT Press, Cambridge, MA, USA, 1979.
14. KURDILA, A., NECHYBA, M., PRAZENICA, R., DAHMEN, W., BINEV, P., DEVORE, R. and SHARPLEY, R. Vision-based control of micro air vehicle: progress and problems in estimation, IEEE Conference on Decision and Control, December 2004, pp 1635-1642.
15. WANG, H. and BRADY, M. A structure-from-motion algorithm for robot vehicle guidance, Intelligent Vehicles Symposium, July 1992, pp 30-35.
16. STEPHENS, M.J., BLISSETT, R.J., CHARNLEY, D., SPARKS, E.P. and PIKE, J.M. Outdoor vehicle navigation using passive 3D vision, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, June 1989, pp 556-562.
17. KANADE, T., AMIDI, O. and KE, Q. Real-time and 3D vision for autonomous small and micro air vehicles, IEEE Conference on Decision and Control, December 2004, pp 1655-1662.
18. SINOPOLI, B., MICHELI, M., DONATO, G. and KOO, T.J. Vision based navigation for an unmanned aerial vehicle, IEEE International Conference on Robotics and Automation, 2001, **2**, pp 1757-1764.
19. FORSYTH, D.A. and PONCE, J. *Computer Vision: A Modern Approach*, Prentice-Hall Publishers, Upper Saddle River, NJ, USA, 2003.
20. LORIGO, L.M., BROOKS, R.A. and GRIMSOU, W.E.L. Visually-guided obstacle avoidance in unstructured environments, IEEE International Conference on Intelligent Robots and Systems, September 1997, **1**, pp 373-379.
21. YAO, Y.-S. and CHELLAPPA, R. Dynamic feature point tracking in an image sequence, 12th IAPR International Conference on Pattern Recognition, October 1994, **1**, pp 654-657.
22. LONGUET-HIGGINS, H. A computer algorithm for reconstructing a scene from two projections, *Nature*, 1981, **293**, pp 133-135.
23. HUANG, T. and FAUGERAS, O. Some properties of the e-matrix in two-view motion estimation, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1989, **11**, (12), pp 1310-1312.
24. ZHANG, Z., DERICHE, R., FAUGERAS, O. and LUONG, Q.-T. A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry, *Artificial Intelligence J*, 1995, **78**, pp 87-119.
25. GRIMSON, W. *From Images to Surfaces: A Computational Study of the Human Early Visual System*, MIT Press, Cambridge, MA, USA, 1981.
26. FAUGERAS, O. *Three-Dimensional Computer Vision*, MIT Press, Cambridge, MA, USA, 1993.
27. OKUTAMI, M. and KANADE, T. A multiple-baseline stereo system, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1993, **15**, (4), pp 353-363.
28. KOENDERINK, J. and VAN DOORN, A. Affine structure from motion, *J Optical Society of America: Part A*, 1990, **8**, pp 377-385.
29. TOMASI, C. and KANADE, T. Shape and motion from image strems under orthography: a factorization method, *Int J Computer Vision*, 1992, **9**, (2), pp 137-154.
30. HARTLEY, R. and ZISSERMAN, A. *Multiple View Geometry in Computer Vision*, Cambridge University Press, Cambridge, UK, 2000.
31. BARRON, A.R., CHOEN, A., DAHMEN, W. and DEVORE, R. Approximation and learning by greedy algorithms, *Annals of Statistics*, 2008, **36**, (1), pp 64-94.
32. ZABARANKIN, M., KURDILA, A., PROKOPYEV, O., GOEL, A., BOGINSKI, V., CAUSEY, R., URYASEV, S. and PARDALOS, P. Vision-based trajectory planning for autonomous micro air vehicles in urban environments, *IEEE Robotics and Automation Magazine*, in review.
33. WEBB, T.P., PRAZENICA, R.J., KURDILA, A.J. and LIND, R. Vision-based state estimation for uninhabited aerial vehicle, AIAA-2005-5869, August 2005.
34. CAUSEY, R.S, LIND, R. and PRAZENICA, R.J. Effects of camera uncertainty on vision-based feedback, AIAA-2005-6097.
35. ROBERTS, P.J., WALKER, R.A. and O'SHEA, P. Fixed wing UAV navigation and control through Integrated GNSS and vision, AIAA-2005-5867, August 2005.
36. KEHOE, J.J., CAUSEY, R.S., ARVAI, A. and LIND, R. Optical flow to estimate angular rates and aerodynamic angles of nonlinear aircraft, submitted to American Control Conference, June 2006.
37. STEVENS, B.L. and LEWIS, F.L. Aircraft control and simulation, Wiley, 2003.
38. KEHOE, J.J., CAUSEY, R.S., ABDULRAHIM, M. and LIND, R. Waypoint navigation for a micro air vehicle using vision-based attitude estimation, *Aeronaut J*, December 2006, **110**, (1114), pp 821-829.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

Pilot task demand load during RNAV approaches with a Cessna Citation

M. M. Heiligers
m.m.heiligers@tudelft.nl
K. T. P. van Bennekom, T. J. van Tuinen, Th. van Holten and M. Mulder
Delft University of Technology, Faculty of Aerospace Engineering
Delft
The Netherlands

ABSTRACT

This research aims to develop a method which predicts the task demand load as experienced by pilots while flying an area navigation (RNAV) approach. First, this will yield insight in which aspects of an approach actually influence pilot task demand load. And second, during the design of approaches this method can be used to rapidly evaluate a potential approach and to 'optimise' an approach with respect to pilot task demand load. During previous research, focusing on approaches flown with a B747, a list of factors that influence pilot task demand load has been obtained, as well as a method to keep pilot task demand load at an acceptable level. The method consists of seven guidelines to be adhered to during approach design. This paper shows that the list of factors and the method do not only apply to a B747 aircraft but are generally applicable to other aircraft as well. This is underpinned by results from both flight simulator tests and real flight tests with TU Delft's Cessna Citation laboratory aircraft. Additionally, it is shown that there are no discrepancies between the list of factors influencing pilot task demand load resulting from the flight simulator tests and the list of factors resulting from the real flight tests.

NOMENCLATURE

AAL	above airport level
ALAR Task Force	Approach and Landing Accident Reduction Task Force
APP	approach
APP01 to APP10	test approach 1 to 10
ATC	Air Traffic Control
B747	Boeing 747
CAVOK	ceiling and visibility OK
CDA	continuous descent approach
CDU	command display unit
CL	checklist
Eratio	energy ratio
FAF	final approach fix
FD	flight director
Flaps APP	flaps approach (flaps 15)
Flaps LAND	flaps land (flaps 40)
FMS	flight management system
GD	gear down
GS	glideslope
HDG	heading
IAF	initial approach fix
IAS	indicated airspeed
IF	intermediate fix
ILS	instrument landing system
LNAV	lateral navigation
LOC	localiser
MCC	multi-crew co-ordination
ML	mental load
NASA TLX	NASA Task Load indeX
NDB	non directional beacon
NLR	National Aerospace Laboratory
PANS-OPS	Procedures for air navigation services – aircraft operations
PF	pilot flying
PFD	primary flight display
PH-LAB	registration of the Cessna Citation test aircraft
PM	pilot monitoring
PMM	point mass model
QNH	pressure at mean sea level
RNAV	area navigation
RSME	rating scale for mental effort
RW	runway
SIMONA	TU Delft's flight simulator
SOP(s)	standard operation procedure(s)
SRS	SIMONA research simulator
TDL	task demand load
THR	threshold
TU Delft	Delft University of Technology
VNAV	vertical navigation
VOR	VHF omni-directional beacon
VREF	reference speed

Paper No. 3635. Manuscript received 29 September 2010, revised version received 7 March 2011, accepted 17 March 2011.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

1.0 INTRODUCTION

This research aims to develop a method which predicts the task demand load (TDL) as experienced by the pilot while flying an approach. TDL is defined as the mental workload imposed by the system to be controlled or supervised[1]. As opposed by mental load, the workload experienced by a particular operator. First, this will yield insight in which aspects of an approach actually influence pilot TDL. And second, during the design of approaches this method can be used to rapidly evaluate a potential approach and to 'optimise' an approach with respect to pilot TDL.

The rationale within this research is that approaches should be designed such that they can be flown according to Standard Operating Procedures (SOPs) and that a stabilised approach at 1,000ft can be achieved. This is based on the conclusions of the Flight Safety Foundation Approach-and-landing Accident Reduction Task Force[2]. These conclusions read, amongst others, that 'Establishing and adhering to adequate SOPs and flight-crew decision-making processes improve approach-and-landing safety' and that 'Unstabilised and rushed approaches contribute to approach-and-landing accidents'. Therefore, within this research, pilot TDL is predicted for approaches while flying according to SOPs and while aiming to achieve a stabilised approach.

The approaches considered in this research are Area Navigation (RNAV) approaches or, more specifically, RNAV transitions since the final part of the approach is guided by the Instrument Landing System (ILS). The approaches are flown using the Flight Management System (FMS), Autothrottle and Autopilot with Vertical Navigation (VNAV) and Lateral Navigation (LNAV) modes. On Localiser intercept heading the autothrottle and autopilot are switched off, and the remainder of the approach is flown using the Flight Director (FD).

Given the level of automation described above, given a certain aircraft with its corresponding SOPs, and given a certain approach, we aim to map pilot TDL and the factors that contribute to pilot TDL. The factors contributing to pilot TDL considered in this research are the properties of the approach trajectory and its speed and altitude constraints (for instance, the Localiser intercept speed or distance available on Localiser Intercept Heading), the meteorological conditions (wind direction and wind speed) and the flight mechanical properties of the aircraft (for instance, how easy it is to dissipate energy)[3]. To investigate pilot TDL we thus focus on factors that can be described as 'the environment' of the pilot, instead of focusing on the constraints of the pilot himself (like memory capacity, time delay, etc.)[3]. In this respect our work is influenced by the principles of cognitive work analysis[4]. This approach deliberately deviates from the idea behind models such as the Procedure-Oriented Crew model (PROCRU)[5,6] or the Man-Machine Integrated Design and Analysis System (MIDAS)[7-9] that use human operator models which do focus on the constraints of the human operator. It is anticipated that by focusing on the environment of the pilot instead of on the limitations of the pilot himself much simpler models can be achieved to predict pilot TDL than by using human operator models.

During previous research[10-14], factors have been identified that influence pilot TDL for pilots flying an RNAV approach with a B747. These factors were identified based on flight simulator tests. In this paper it will be demonstrated that the same factors also influence pilot TDL when flying an approach with a Cessna Citation. This indicates that the set of factors that has been identified is a generally applicable set of factors, and not only valid for the B747. Additionally, it is investigated whether the same set of factors influences pilot TDL during flight simulator tests and during real flight. To this end, in this paper, both the results of a flight simulator experiment for a Cessna Citation and the results of real flight tests with a Cessna Citation aircraft are compared. It will be demonstrated that the same set of factors results from both experiments. Finally, the simulation tool that was developed for the B747 in order to analyse an approach with respect to the factors that were proven to influence pilot TDL is adjusted in order to include the Cessna Citation. It will be shown that the simulation tool also works and provides reliable results for the Cessna Citation.

This paper is structured as follows. First, the basic principles of this research are introduced as well as the scope of the research. After that, the results of previous research[11-13] which focused on the B747 are briefly explained. Subsequently, the human in the loop experiments are presented, these experiments are conducted for the Cessna Citation aircraft both in a flight simulator and during real flight tests. To conclude, the simulation tool that is adjusted to include the Cessna Citation is explained, and its predictions are illustrated by a case study.

2.0 BASIC PRINCIPLES OF THIS RESEARCH

At the heart of the project lies the development of a method that will provide guidelines that can be used during approach design in order to keep pilot TDL during the approach at an acceptable

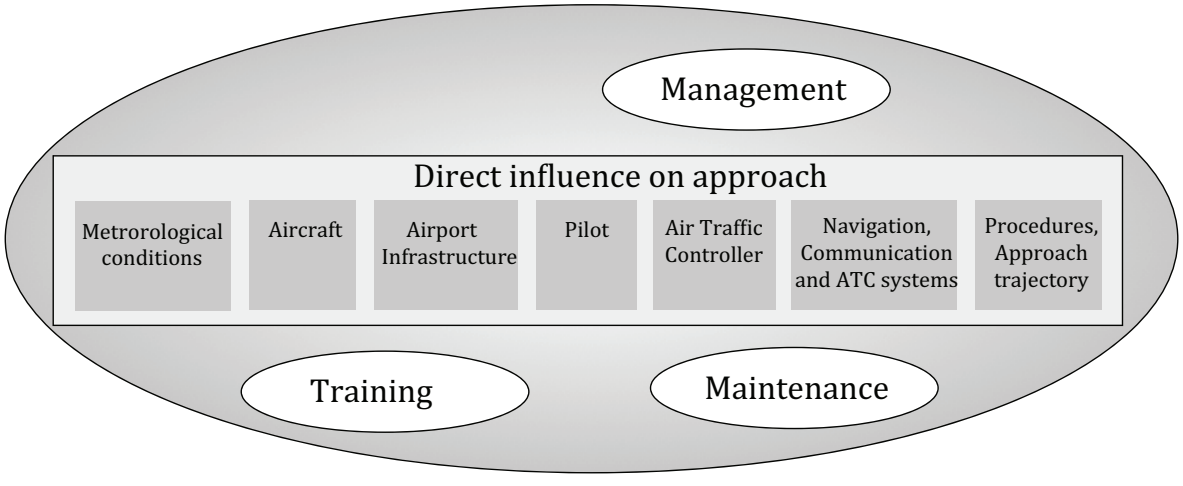


Figure 1. Direct and indirect factors that influence the safety of airport approaches.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

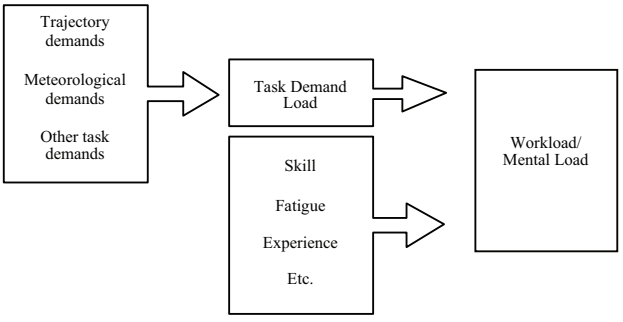


Figure 2. Difference between task demand load and workload, adapted from Ref. 14.

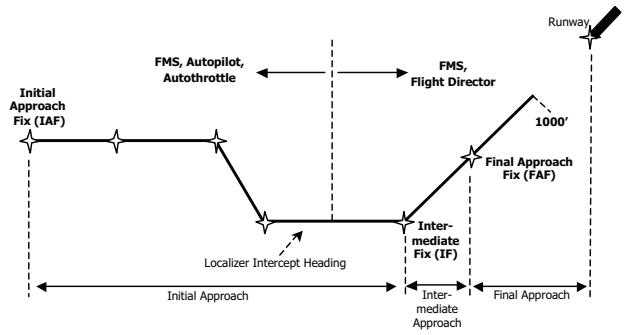


Figure 3. Part of flight considered (top view) and automation used.

level. In order to analyse whether a newly designed approach actually meets all the guidelines, a computer simulation program is developed. This simulation program incorporates the aspects that affect pilot TDL during approach, including standard operating procedures, altitude-profiles, velocity-profiles, etcetera. It should also be possible to enter different types of aircraft and to change the meteorological conditions (turbulence intensity, amount of wind). These properties are the descriptors of the environment that form the 'input' of the computer program as they constitute the specific characteristics of the approach to be evaluated. The 'output' of the simulation program is an indication whether the guidelines to keep pilot TDL at an acceptable level are met. This section will explain the basic principles of the method and computer simulation, the assumptions and the choices that have been made as to what is incorporated within this research, and also what is considered to be beyond the scope of this research.

2.1 Factors of the air transport system included

Many different factors and the interactions between those factors have an influence on the execution of an approach, see Fig. 1. This research concentrates on the 'pilot' box in Fig. 1. It will, e.g., not consider the Air Traffic Controller's TDL. To determine pilot TDL, this research will only take into account the factors that have a direct influence on an approach (see Fig. 1), most importantly the characteristics of the trajectory, the type of aircraft and the meteorological conditions.

2.2 Task demand load

This research aims to develop a method to predict pilot TDL. Task Demand Load is defined as the mental workload imposed by the system to be controlled or supervised[1], see also Fig. 2. The TDL is not to be mistaken for the mental workload experienced by the human operator, which is referred to as Mental Load (ML). Many of the well-known methods to measure workload, like the NASA Task Load indeX, measure ML, not TDL.

Within this research several experiments are performed during which pilots are asked to comment on approaches regarding the amount of effort these approaches require, or their effect on the difficulty as experienced by the pilot. When pilots give their opinion on these matters, they obviously base their opinion on the mental workload they experienced. This results in the situation that in order to obtain information about the task demand load, pilots are asked about the mental workload they experienced during the experiments, unfortunately there is no other way. However, by choosing pilots with different levels of experience etc., by testing the approaches in random order and by converting the pilots' ratings to *z*-scores it is assumed that through the comments of the pilots a good indication of the task demand load can be obtained.

2.3 Approaches considered and automation used

Obviously, the TDL depends directly on the type of approach that is considered. This research focuses on Area Navigation (RNAV) Approaches. Although it is appreciated that non-precision approaches such as Non-Directional Beacon (NDB) approaches are, in general, more difficult for a pilot to fly than RNAV approaches[16], a deliberate choice is made to focus on RNAV approaches only, since these are expected to become more and more frequently used in the future. The last part of the RNAV approach is assumed to be flown using the Instrument Landing System (ILS).

The part of the flight that is considered starts at the Initial Approach Fix (IAF) and comprises the entire approach (Initial Approach, Intermediate Approach and Final Approach) until 1,000ft above airport level, see Fig. 3. Based on interviews with pilots it was decided to use two different levels of automation during the approach: until Localiser Intercept Heading the approach is flown using the FMS, Autopilots and Autothrottle. At Localiser Intercept Heading (but before Localiser capture) the pilot switches to Flight Director (FD) mode and disconnects the Autothrottle, the remainder of the approach is thus flown using the FMS and FD, which implies manual control by the pilot.

2.4 Non-nominal conditions and emergencies

Non-nominal conditions and emergencies such as engine failure are not considered in this research. The goal is to determine pilot TDL for published RNAV approaches under nominal conditions. When any emergencies such as engine failure occur, the crew will most likely not be required to follow the RNAV approach anyway, but will be vectored to the runway in the most convenient way.

Additionally, the assumption for less severe non-nominal situations is that when flying under nominal conditions, the RNAV approach should provide enough 'margin' with respect to pilot TDL, such that the pilot has enough spare capacity and time to deal with non-nominal conditions. This implies that the TDL that is predicted by this research for nominal conditions should be well below the absolute maximum TDL a pilot can cope with in order to guarantee this margin.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

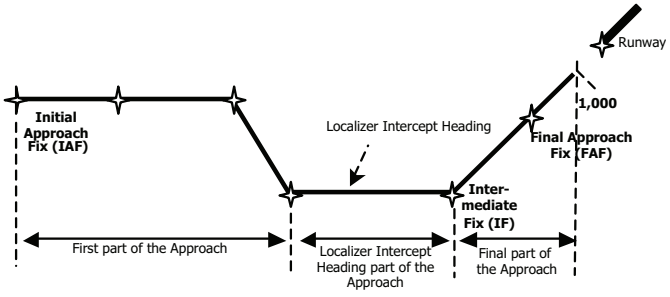


Figure 4. Division of the approach into three parts.

2.5 Boundary conditions: Stabilised approach and standard operating procedures

The TDL experienced by the pilot also depends on the boundary conditions that are set, e.g. the accuracy with which the approach needs to be flown. The boundary conditions chosen for this research are that the approach should be performed according to standard operating procedures and that pilots should aim to achieve a stabilised approach at 1,000ft above airport elevation. This decision is based on the conclusions of the ALAR Task Force[2].

To determine whether a stabilised approach is achieved at 1,000ft, the following nine criteria[2] are used:

1. The aircraft is on the correct flight path;
2. Only small changes in heading/pitch are required to maintain the correct flight path
3. The aircraft speed is not more than VREF + 20kt Indicated Airspeed (IAS) and not less than VREF;
4. The aircraft is in the correct landing configuration;
5. Sink rate is no greater than 1,000ft per minute; if an approach requires a sink rate greater than 1,000ft per minute, a special briefing should be conducted;
6. Power setting is appropriate for the aircraft configuration and is not below the minimum power for approach as defined by the aircraft operating manual;
7. All briefings and checklists have been conducted;
8. Specific types of approaches are stabilised if they also fulfill the following: instrument landing system (ILS) approaches must be flown within one dot of the glide slope and localiser; a Category II or Category III ILS approach must be flown within the expanded localiser band; during a circling approach, wings should be level on final when the aircraft reaches 300ft above airport elevation; and
9. Unique approach procedures or abnormal conditions requiring a deviation from the above elements of a stabilised approach require a special briefing.

2.6 Level of detail of computer simulation models

As briefly explained in the introduction, it is the goal to incorporate very detailed models of the environment of the pilot in the computer simulation, and to add to this a rather simple model for the pilot. Therefore, the aircraft with its kinematic and dynamic constraints, the 3D properties of the trajectory, the velocity profile, turbulence, wind, etcetera, in other words: the factors that have a direct influence on an approach as given in Fig. 1, are modeled as detailed and accurate as possible. Whereas the pilot model is kept as simple as possible. This simple pilot model consists of a manual control model (which in effect only contains a pure gain plus time delay) and a model for performing actions such as selecting flaps and gear according to the SOPs.

3.0 RESULTS OF PREVIOUS RESEARCH FOR B747

Using these basic principles and assumptions, a method (consisting of guidelines to keep pilot TDL at an acceptable level) and two computer simulations have been developed for the B747. The results are briefly explained in this section.

Based on two sets of B747 flight simulator experiments with nine B747 pilots participating in each experiment[12,13], a list of factors has been identified that influence pilot TDL during approach. This list of factors is considered to be complete, which means that there are not any other factors that fall within the previously defined scope of this research that influence pilot TDL. The factors that influence pilot TDL are grouped per approach part (see Fig. 4) and can be summarised as follows:

For the first part of the approach:

- The major contributor to pilot TDL is the fact whether or not the altitude and velocity constraints can be met at the waypoints. This is only true when the effect of not meeting the constraints continues into the Localiser part or final part of the approach. If the consequences of not meeting the constraints remain within the first part of the approach this does not influence pilot TDL.
- The number of waypoints, number of heading changes and the altitude profile (horizontal approach, CDA, stepped approach) do not influence pilot TDL. This is due to the fact that this part of the approach is flown in LNAV and VNAV modes with autopilot and autothrottle.

For the Localiser intercept part of the approach:

- The time available to perform all actions (which is directly related to the distance available on Localiser intercept heading) is the most important factor for pilot TDL. Actions that need to be performed for the B747 on Localiser intercept heading are: select flaps 10, select heading select, arm the approach, and (due to the choices made for this research, see Fig. 3) switch off the autopilot and autothrottle.
- Next to this, pilot TDL is also influenced by the Localiser intercept speed, the Localiser intercept angle, and whether the constraints at the waypoints can be met.

For the final part of the approach:

- The most important factors influencing pilot TDL seem to be whether or not a stabilised approach can be achieved at 1,000ft, the distance between IF and FAF and the airspeed on final. Whether an approach is stabilised can, for a B747, be determined from: (1) whether the constraints at the waypoints can be met during the final part of the approach, (2) the value of the vertical speed (should be below the sink rate warning) which in itself is a function of airspeed on final and glideslope angle, and (3) the FAF altitude and distance between IF and FAF (resulting in the line-up distance) since these two factors together determine whether there is enough time available to perform all actions required for a stabilised approach. All these factors thus influence pilot TDL during the final part of the approach.

The method to predict pilot TDL during approach for a B747 is based on the above factors. The method basically consists of seven guidelines for the design of approaches. When these guidelines are met, pilot TDL during the approach will be acceptable. Starting point for the guidelines is that pilots should fly the approach according to SOPs and that they should aim to achieve a stabilised approach at 1,000ft.

Concluding, the guidelines for the contributors to pilot TDL for the B747 then are that:

1. aircraft should be able to meet the altitude and airspeed constraints throughout the approach, especially during the final part of the approach, and during the first part of the approach if this has consequences for the subsequent parts of the approach;

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

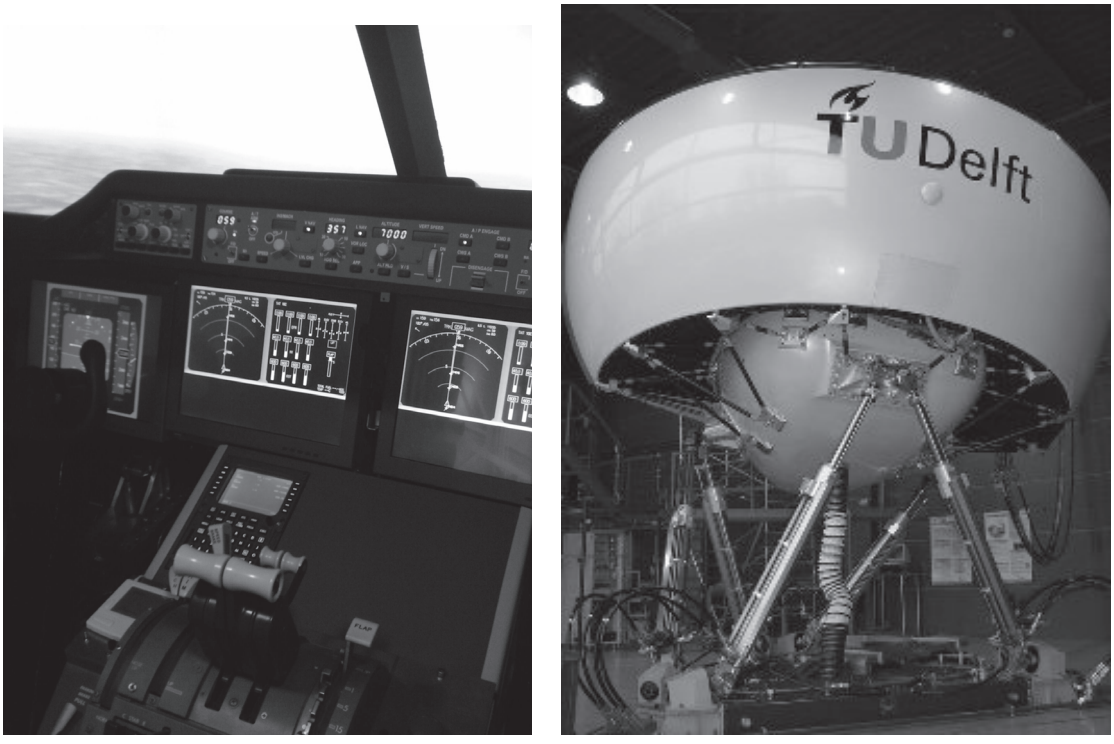


Figure 5. The SIMONA Research Simulator.

2. there should be sufficient time to perform all actions on Localiser intercept heading;
3. it should be possible to achieve a stabilised approach. Whether a stabilised approach can be achieved for a B747 depends on 1. whether the aircraft can dissipate enough energy during the final part of the approach, 2. the value of the vertical speed (should be below the sink rate warning) which in itself is a function of airspeed on final and glideslope angle, and 3. the FAF altitude and distance between IF and FAF (resulting in the line-up distance);
4. the distance between IF and FAF should be sufficient;
5. the vertical speed should be below the sink rate warning;
6. the Localiser intercept speed should not be too high, and that
7. the Localiser intercept angle should not be too large.

A quantification of these guidelines is given in Ref. 13.
It is hypothesised that the same factors will also determine pilot TDL for other aircraft types. In this respect it should be noted that the factors 'FAF altitude' and 'time available on Localiser intercept heading' are factors that originate from the fact that for the B747 the SOPs require pilots to perform a number of actions on Localiser intercept heading and between the FAF and 1,000ft, and that they should have sufficient time to do so. If, for another aircraft type, these actions are required to be performed in another part of the approach, then care should be taken that sufficient time is available in that particular part of the approach. In that case, the factors 'time available on Localiser intercept heading' and 'FAF altitude' might not influence pilot TDL for that particular aircraft type. All other factors in the list above are assumed to be valid for all aircraft types.

In order to obtain a prediction whether these guidelines are met for an approach, a comprehensive Monte Carlo computer simulation was developed for the B747, together with a relatively simple point mass model computer simulation[11-13]. By running either of these computer simulations for a specific approach and analysing their output combined with the requirements in Procedures for Air Navigation Services Aircraft Operations (PANS-OPS)[17], an indication can be obtained of whether or not the guidelines as described above are met[11-13].

Questions arising from this previous research are whether the guidelines outlined above are indeed valid for other aircraft types, and whether the fact that they were obtained from flight simulator tests might cause discrepancies with guidelines for real flight. To find an answer to these questions, this paper describes flight simulator experiments and real flight tests with a Cessna Citation aircraft. Given the goal of this research (to develop a method to keep pilot TDL at an acceptable level during RNAV approaches) the choice for a Cessna Citation aircraft might not be obvious, since it is not among the aircraft types that will most frequently fly these kinds of approaches. However, it is the only aircraft type for which non-linear aerodynamic models were available for the simulations, and the only aircraft that was available for flight tests. Therefore this paper concentrates on the Cessna Citation. It should be noted in this respect that the particular aircraft that was used does not have a VNAV mode or an autothrottle. When comparing the results found for the Citation to the results found for the B747 this should be kept in mind.

4.0 HUMAN IN THE LOOP EXPERIMENT

Two separate human-in-the-loop experiments were performed for the Cessna Citation aircraft which involved pilots flying different approaches under varying conditions. The first experiment was conducted in a six-degree-of-freedom flight simulator, while the second was conducted in a Cessna Citation II aircraft.

4.1 Experiment goal

The first goal of the experiment was to test whether the same factors that were found to influence pilot TDL for the B747[14] are general factors, and also influence pilot TDL for other aircraft, in this case a Cessna Citation. The second goal was to investigate whether the factors that influence pilot TDL during flight simulator tests also influence pilot TDL during real flight.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

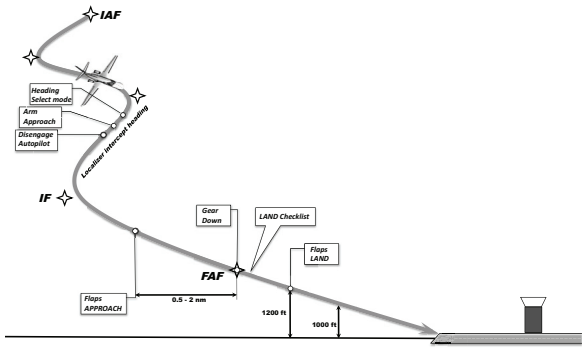


Figure 6. Standard operating procedures for Cessna Citation.

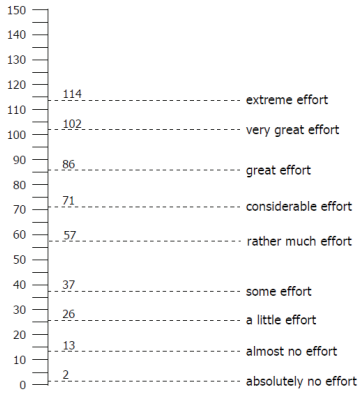


Figure 7. Rating scale mental effort, during the experiment the Dutch translation was used.

4.2 Method of experiment 1, flight simulator tests

4.2.1 Apparatus and Citation model

The experiment was performed in the six-degree-of-freedom TU Delft SIMONA Research simulator (SRS), see Fig. 5. The Cessna Citation aerodynamic models as well as the yaw damper are based on the Cessna Citation I[18]. Autopilot modes available during the experiment were: LNAV, Heading Select, Altitude Hold, Vertical Speed and Indicated Airspeed (IAS). In Flight Director operation additional modes available were Glide slope mode and Localiser mode.

The autopilot and flight director models are based on the autopilots developed for the model of the B747[13], which were derived from Ref. 19. The LNAV mode is based on the VOR modes described in Ref. 19.

All approaches that were flown, were pre-programmed in the FMS/CDU. The appropriate approach was loaded in the FMS before the start of the approach, and during the experiment pilots could switch between the 'Progress' and 'Legs' pages, but could not use the CDU interactively, or modify the approach.

There were some discrepancies between the SRS and a Citation that are of importance to the experiment. First, the cockpit lay-out in the SRS differed from reality (see Fig. 5). Second, the aircraft was not trimmed when the pilot switched from autopilot to flight director. Third, the altitude capture was slightly abrupt as compared to reality. This resulted in a minor 'bubble feeling' when an altitude was captured. Fourth, with respect to the flight director, at LOC intercept the flight director over exaggerated the bank angle required for a correct intercept. In practice this meant that during LOC intercept the pilots rolled the aircraft to a correct bank angle and waited for the FD bars to return to the centre position. All pilots were briefed about these discrepancies before the tests commenced. Also, all these aspects were the same for all pilots and constant during all runs.

4.2.2 Subjects and instructions

Six Citation pilots participated in the experiment, total flight hours ranging from 500 to 13,200 hours (M = 7,100 hours, s = 5,047 hours). Their flight hours on the Cessna Citation ranged from 200 to 1,500 hours (M = 750 hours, s = 521 hours). The pilots were paired up to form a crew of pilot flying (PF) and pilot monitoring (PM). The task of each pilot was to fly 10 different approaches as PF and the same 10 approaches as PM, starting at the Initial Approach Fix (IAF) and ending at around 800' above airport level (AAL).

Pilots were instructed to fly the approaches according to SOPs. The SOPs are used as stated for the Aircraft Operations Manual of the Cessna Citation II[20], see Fig. 6. It is assumed that before every run the approach checklist has been completed. On Localiser intercept heading pilots should switch to heading select mode, arm the approach and switch from AP to FD (the latter requirement is not prescribed by SOPs but is based on the choices made for this research). The pilots are then required to select Flaps APP between 2nm and 0·5nm from the FAF. On the FAF the Gear is selected, followed immediately by the landing checklist. At 1,200' above the airfield level Flaps LAND are selected.

Two weeks before the experiment the pilots received a briefing by mail. On the day of the experiment they were briefed as well. The pilots were asked to adhere very strictly to SOPs, even if they could foresee that by adhering to SOPs they would not meet certain constraints at waypoints or would end up unstabilised at 1,000ft. Additionally they were asked to perform their tasks according to the principles of Multiple Crew Coordination (MCC) and to fly passenger comfort. They were briefed about the discrepancies between the SRS and the Citation (as explained in the previous paragraph), and were informed that there would be no emergencies (e.g. engine failure) during the flight. They were told that they could fly the approach as published on the approach and landing charts, implying that ATC would not interfere. The pilots were not allowed to use speedbrakes, since approaches should be designed such that they can be flown without the use of speedbrakes.

4.2.3 Procedure

Before starting the experiment the pilots could familiarise themselves with the SRS and their task during three to five (depending on the pilot) practice approaches. After that the experiment started. Before every approach the pilots could take as much time as they thought necessary to study the approach and landing charts, to brief the approach and to prepare the SRS for the next approach. The simulation was started when the pilots indicated they were ready.

After every approach the pilots (PF) were asked to fill in a run questionnaire. Each run questionnaire consisted of three parts: the first part required a rating of the total approach on the Rating Scale Mental Effort (RSME)[21], see Fig. 7 and required additional RSME ratings for the three individual parts of the approach, resulting in three RSME sub-ratings per approach. The RSME is constructed according to the 'magnitude estimation' method[22] and can therefore be regarded as interval data. The Dutch version of the scale (which was also used for this research) was used and validated in[21,23]. Though originally intended to measure only one aspect of a task, it is used here to get an indication of the total task because of its simplicity and ease of use when compared to, for example, a NASA TLX rating procedure[24].

The second part of the run questionnaire contained two questions asking the pilot's opinion on whether the pilot would have adhered

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

to SOPs during real flight, and whether the pilot would have used speedbrakes during real flight. The third part of the run questionnaire consisted of specific closed format questions per approach part regarding the factors hypothesised to influence pilot TDL in that specific approach part. For an example, see Fig. 8. To analyse the pilots' answers, the response options for all questions have been coded from 1 – 5, and are regarded interval data. Although there is much controversy about whether these response options can be considered ordinal or interval data[25,26], it is, in this case, deemed appropriate to treat the data as interval scale because the response options were arranged horizontally and were equally spaced apart, and the verbal labels connoted more-or-less evenly-spaced gradations, most of them symmetrical about a neutral middle.

At the end of the day, after all approaches were flown, the pilots (PF) filled in an end of day questionnaire. The first part of the end of day questionnaire regarded the realism of the flight simulator and the realism of the experiment as a whole. The second part contained general questions about factors that might possibly influence pilot TDL during approach.

4.3 Method of experiment 2, Cessna Citation flight tests

4.3.1 Apparatus

The second experiment was performed in the Cessna Citation II laboratory aircraft (PH-LAB) see Fig. 9. The PH-LAB is jointly owned by Delft University of Technology and the National Aerospace Laboratory (NLR). As with the SRS experiments, the same autopilot modes were available to the pilots and the FMS/CDU was used in the same manner.

The flight tests were performed at Malta International Airport. All tests were performed in either CAVOK (officially this means no clouds under 5,000ft, in practice it was a clear blue sky) or in FEW012 to FEW033, which means that ⅛ to ¼ of the sky is covered with clouds with a cloud base at 1,200ft to 3,300ft, respectively. The few small clouds that were present during the tests did not have an effect on the visibility. Maximum windspeeds encountered during the tests were 12kt, with a maximum tailwind of 7kt. During the test week the QNH only varied between 1,014 and 1,017hPa, which means that there was not much variability in air density, and therefore only a very small effect on the flight mechanics due to this change in air density. All in all it can be stated that the meteorological conditions during the real flight tests were very similar to the conditions during the flight simulator experiment.

The number of waypoints was	Very large	Large	Neutral	Small	Very small
As a result I found this part of the approach	Very difficult	Difficult	No influence	Easy	Very easy

Figure 8. Example of question from the run questionnaire for the flight simulator tests.

4.3.2 Subjects and instructions

For this experiment the same six pilots participated and the same 10 approaches were flown. All pilots were given the same instructions as were given during the SRS experiments.

4.3.3 Procedure

The procedure used during the Citation flight tests was similar to the SRS experiments. Before starting the experiment the pilots could familiarise themselves with the Citation and their task during one or two (depending on the pilot) practice approaches. After which the experiment started. Before every approach the pilots could study the approach and landing charts, to brief the approach. The experiment started when the aircraft crossed the IAF and ended around 800ft above airport level (AAL). At this altitude a go-around was initiated and the aircraft was maneuvered to the IAF of the next approach.

Obviously the aircraft could not be paused after each approach, so there was considerably less time available to fill in a run questionnaire. Instead, during the go-around, the PM was given control over the aircraft while the PF completed a single page questionnaire. This consisted of an RSME scale for the entire approach and three RSME scales for the different parts. Using the headsets the researchers could then ask some specific questions about the approach for additional information.

When the PF had completed all ten approaches an end of day questionnaire was filled out. This questionnaire resembles the end of day questionnaire used during the SRS experiments and can be used for extra information later in the analysis.

4.4 Independent variables and approaches

Considering the time and resources available for this research, 10 custom approaches were flown during the human-in-the-loop experiments. It was chosen to design one benchmark approach and nine approaches for each of which a separate independent variable is changed with respect to the benchmark approach, see Table 1.



Figure 9. Cessna Citation II laboratory aircraft.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

Table 1
Independent variables

	Independent variable	Linked factor	TDL effect
APP01	Benchmark	Benchmark	
APP02	Short LOC intercept heading	-	+
APP03	Large LOC intercept angle	-	+
APP04	Low FAF (normal line-up distance)	Lower FAF speed Lower LOC int. speed Long distance IF-FAF	+
APP05	Short line-up distance	Lower LOC int. speed Lower IF altitude Short distance IF-FAF Lower FAF altitude Lower FAF speed	+
APP06	Short leg IF-FAF	Lower LOC int. speed Lower IF altitude	+
APP07	High LOC intercept speed	Large line-up distance Higher IF altitude Large distance IF-FAF	+
APP08	Not meeting constraints at WP's, but stabilized at 1000'	Horizontal GS intercept	+
APP09	Not stabilized due to high speed at low FAF (Eratio>1)	Higher LOC int. speed (Lower FAF alittude Higher FAF speed)	+
APP10	Not stabilized due to high energy at start of approach (Eratio>1)	Higher LOC int. speed Low IF altitude Low FAF altitude Higher FAF speed Short distance IF-FAF Short line-up distance	+

The factor 'Eratio' in the second column needs explanation: this is the energy rate demand, which is the ratio between the rate at which the trajectory requires the aircraft to dissipate energy, and the rate at which the aircraft can dissipate energy. Once this ratio becomes larger than one this means that the altitude and velocity constraints at the next waypoint will not be met.

If possible, the independent variable was the only changing factor between the benchmark and the respective approach. However, sometimes, due to changing the independent variable another aspect of the approach also had to be changed. These aspects, if there were any, are listed in the column labeled 'Linked factor'. For example, for APP07 a higher Localiser intercept speed was applied. There are then two options: (1) if the FAF altitude, line-up distance and required airspeed at the FAF are maintained equal to the benchmark approach, this will result in an energy rate demand larger than 1 between IF and FAF, (2) the other option is to keep the energy rate demand the same as in the benchmark approach, but in order to do this the line-up distance (and thus IF altitude and distance between IF and FAF) needs to be increased. While designing the approaches in Table 1 the objective was to keep the energy rate demand for all approaches equal to the benchmark approach (except when the energy rate demand was the independent variable), since previous research[12,13] showed that when the value of the energy rate demand becomes larger than one, this has a large influence on pilot TDL. Therefore, for APP07 the second option is chosen for the design.

4.5 Hypotheses

Regarding the influence of the independent variables the following was hypothesised (see also last column of Table 1):

- A short LOC intercept heading increases pilot TDL
- A large LOC intercept heading angle increases pilot TDL
- A low FAF increases pilot TDL
- A short line-up distance increases pilot TDL
- A short leg IF-FAF increases pilot TDL
- A high LOC intercept speed increases pilot TDL
- Not meeting constraints increases pilot TDL
- Not being stabilised at 1,000ft increases pilot TDL

5.0 RESULTS

Due to the total number of tests that is performed on the data set resulting from the flight simulator experiment and real flight tests, a *p*-value smaller than ·05 should be used (when applying a Bonferroni correction) as criterion for significance in order to control the Type I error rate (this error represents the situation that an effect is found using a statistical test, whereas in reality there is no effect). However, by correcting the p-value to smaller than ·05, statistical power is lost (meaning that the probability of rejecting an effect that does actually exist is increased (a Type II error)). Moreover, with the small sample size available for all comparisons ($N = 12$ at most) it will be highly unlikely to arrive at very small significance values. Therefore, it is decided to use a significance value of ·05 for all comparisons, while keeping in mind that this inflates the Type I error, and that we thus might classify factors as influencing pilot TDL while in reality they do not. For the purpose of this research however, this is deemed more favorable than discarding factors that actually do influence pilot TDL. In any regard, due to the small sample size, the results of the tests given below should only be considered as an indication of possible effects.

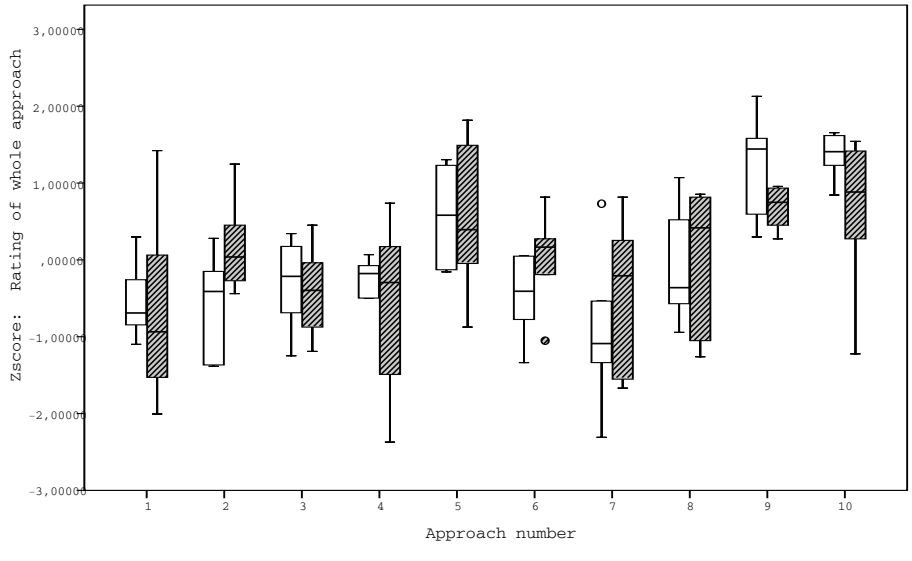


Figure 10. Boxplot of standardised ratings of entire approach, white boxplots are for flight simulator tests, striped boxplots for real flight tests.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

Table 2
Results of paired samples *t*-test (comparison of SIMONA research simulator test and Citation) of the *z*-scores of the RSME ratings of the entire approach

		Paired Differences							
		Mean	Std. Deviation	Std. Error Mean	95% Confidence Interval of the Difference		t	df	Sig. (2-tailed)
					Lower	Upper			
Pair 1	SIM1 - Cit1	.10609	1.32690	.54170	-1.28640	1.49859	.196	5	.852
Pair 2	SIM2 - Cit2	-.75148	1.12667	.45996	-1.93385	.43089	-1.634	5	.163
Pair 3	SIM3 - Cit3	.09902	1.15279	.47062	-1.11075	1.30880	.210	5	.842
Pair 4	SIM4 - Cit4	.36378	1.29521	.52877	-.99546	1.72302	.688	5	.522
Pair 5	SIM5 - Cit5	.03972	1.61328	.65862	-1.65331	1.73275	.060	5	.954
Pair 6	SIM6 - Cit6	-.50136	1.07567	.43914	-1.63021	.62748	-1.142	5	.305
Pair 7	SIM7 - Cit7	-.51029	1.51476	.61840	-2.09993	1.07936	-.825	5	.447
Pair 8	SIM8 - Cit8	-.13877	1.65862	.67713	-1.87938	1.60185	-.205	5	.846
Pair 9	SIM9 - Cit9	.56194	.63780	.26038	-.10739	1.23127	2.158	5	.083
Pair 10	SIM10 - Cit10	.73134	1.12780	.46042	-.45221	1.91489	1.588	5	.173

5.1 Comparison between flight simulator tests and real flight tests

The first goal of the flight simulator tests and real flight tests is to identify whether pilots classify the same factors as increasing or decreasing pilot TDL during both experiments. If there would be discrepancies between the results of the two experiments, for instance, pilots would identify a large Localiser angle as increasing pilot TDL during the simulator tests, and would identify this same factor as decreasing pilot TDL during the real flight tests, this would be apparent from the RSME scores for the approach in which this factor was tested.

To test whether this is the case the RSME *z*-scores for all pilots and all approaches are calculated for each of the two experiments. The RSME *z*-scores for the entire approach are plotted in Fig. 10 for both experiments, the RSME *z*-scores for the three approach parts are also calculated but not given in Fig. 10. If the same factors influence pilot TDL during both experiments, one would expect to see the same trend relative to the benchmark approach (approach 1) for both experiments.

To analyse whether the same trend can be observed, the RSME *z*-scores for both experiments are compared for each approach. If there is no difference between the RSME *z*-scores per approach this means that the factor that was tested during that approach had the same influence both during the simulator experiment and during the real flight tests. The RSME *z*-scores are compared using the paired samples *t*-test for the parametric cases, and using the Wilcoxon matched pairs test for the non-parametric cases. In total 40 comparisons were made (four RSME ratings per approach for 10 approaches).

The results for the RSME ratings for the entire approach are given in Table 2, there was no significant difference ($p < \cdot 05$) for any of the approaches, and hence for any of the TDL influencing factors tested during these approaches. Comparison of the RSME *z*-scores for the three approach parts proved that three comparisons out of the thirty possible comparisons were significant ($p < \cdot 05$), which

Table 3
Analysis of differences in effort between the approaches and APP01

	Independent measure	RSME *z*-scores entire approach	RSME *z*-scores part of approach	Run questionnaire flight sim tests	In-flight interview real flight tests	Conclusion
APP01	-	-	-	-	-	-
APP02	Short LOC intercept heading	No difference	More effort	More effort	More effort	More effort
APP03	Large LOC intercept angle	No difference	More effort	More effort	More effort	More effort
APP04	Low FAF, normal line-up distance	No difference	More effort	More effort	More effort	More effort
APP05	Short line-up distance	More effort	More effort	More effort	More effort	More effort
APP06	Short leg IF-FAF, normal line-up distance	No difference	No difference	No difference	No difference	No difference
APP07	High LOC intercept speed	No difference	No difference	No difference	No difference	No difference
APP08	Not meeting constraints at WP's, but stabilized at 1000'	No difference	More effort	More effort	More effort	More effort
APP09	Not stabilized due to high speed at low FAF	More effort	More effort	More effort	More effort	More effort
APP10	Not stabilized due to high energy at start of approach	More effort	More effort	More effort	More effort	More effort

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

Table 4
Overview of factors that increase the TDL during RNAV approaches

Factor	Approach	Run-question-naire flight sim test	In-flight interview real flight test	End-of-day question-naire flight sim tests	End-of-day question-naire real flight tests	RSME z-scores part of approach	Conclusion
Eratio more than 1, part 1 approach	APP08	More effort	More effort	More effort	More effort	More effort	More effort
Eratio more than 1, part 2 approach	APP10	More effort	-	More effort	More effort	-	More effort
Eratio more than 1, part 3 approach	APP09 & APP10	More effort	-	More effort	More effort	-	More effort
Short LOC intercept heading	APP02	More effort	More effort	-	-	More effort	More effort
Large LOC intercept angle	APP03	More effort	More effort	More effort	More effort	More effort	More effort
High LOC intercept speed	APP07, 09 & 10	More effort	Neutral effort	More effort	More effort	-	More effort*
Short leg IF-FAF	APP05, 06 & 10	More effort	-	More effort	More effort	-	More effort*
Short Line up distance	APP05 & 10	More effort	-	More effort	More effort	-	More effort
High Speed on FAF	APP09 & 10	More effort	More effort	More effort	More effort	More effort	More effort
Low FAF	APP04, 05 & 10	More effort	More effort	More effort	More effort	-	More effort
Not stabilized at 1000'	APP09 & 10	More effort	-	More effort	More effort	-	More effort
Horizontal intercept of GS	APP08	Less effort	Less effort	-	-	-	Less effort

(* = under specific circumstances - = No information available)

indicates that there was, in some way, a difference in effect for the same approach parts between the two experiments.

For approach 6 (short leg IF-FAF) there was a difference in the RSME *z*-scores for the first part of the approach. Compared to the benchmark approach, the RSME z-scores were lower for both experiments, but the RSME *z*-scores for the real flight tests were even lower when compared to the *z*-scores for the flight simulator tests. What caused this difference is unclear, since compared to the benchmark approach nothing was adjusted in this part (segment) of the approach.

For approach 2 (Short Localiser intercept heading) there was a difference in RSME *z*-scores for the final part of the approach. The RSME *z*-scores were both higher than for the benchmark, but for the real flight tests even more so. Approach 2 was the approach with the short Localiser intercept heading. Therefore, this might indicate that this factor had more effect on pilot TDL during the real flight tests than during the flight simulator tests.

For approach 9 (Not stabilised due to high speed at low FAF) the RSME *z*-scores for the final part of the approach were significantly different. Again, both RSME *z*-scores were higher than for the benchmark, but this time this effect was more pronounced for the flight simulator tests. Approach 9 was designed such that it was very difficult to achieve a stabilised approach at 1,000ft. This factor thus appears to have more effect on pilot TDL during the simulator tests than during the real flight tests.

Although some differences were found in the pilots' ratings of the approach during both experiments it can be concluded that the 'direction' of the effect, i.e., the trend, was always the same. That is, for both experiments the same factor always either increased or decreased pilot TDL, there was never an opposite effect between the two experiments for the same factor. The difference in RSME *z*-scores was caused by the fact that the effect was more pronounced in one of the two experiments. Additionally, this difference in effect was only found in the RSME sub-ratings (the rating for an approach part), it never affected the 'overall' RSME rating for the entire approach. Therefore it can be concluded that pilots classify the same factors as increasing or decreasing pilot TDL during both experiments.

Some differences, although they did not result in different factors for TDL, could be observed between the two experiments. It appeared that the duration of the landing checklist is longer during the real flight tests, than during the flight simulator experiment. The mean for the simulator data being 15·7 seconds and the mean for the Citation real flight tests is 20·9 seconds. During the flight simulator tests it was already observed that the landing checklist was often hastily completed. In some cases certain actions that needed to be performed by the PF during the checklist (physically checking the brake pressure by pressing the pedals, actually flipping the ignition switches, etc) were in fact not done. During the real flights in the Citation aircraft the checklist was taken more seriously for obvious reasons.

Another difference that was observed concerns the communication with Air Traffic Control (ATC). During the flight simulator tests all communication was standard and the same for all approaches (in order not to add yet another variable to the test). As a result pilots would know, after having flown a couple of approaches, what ATC was going to say. Consequently, after a while, they would continue performing checklists even if ATC was giving them instructions.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

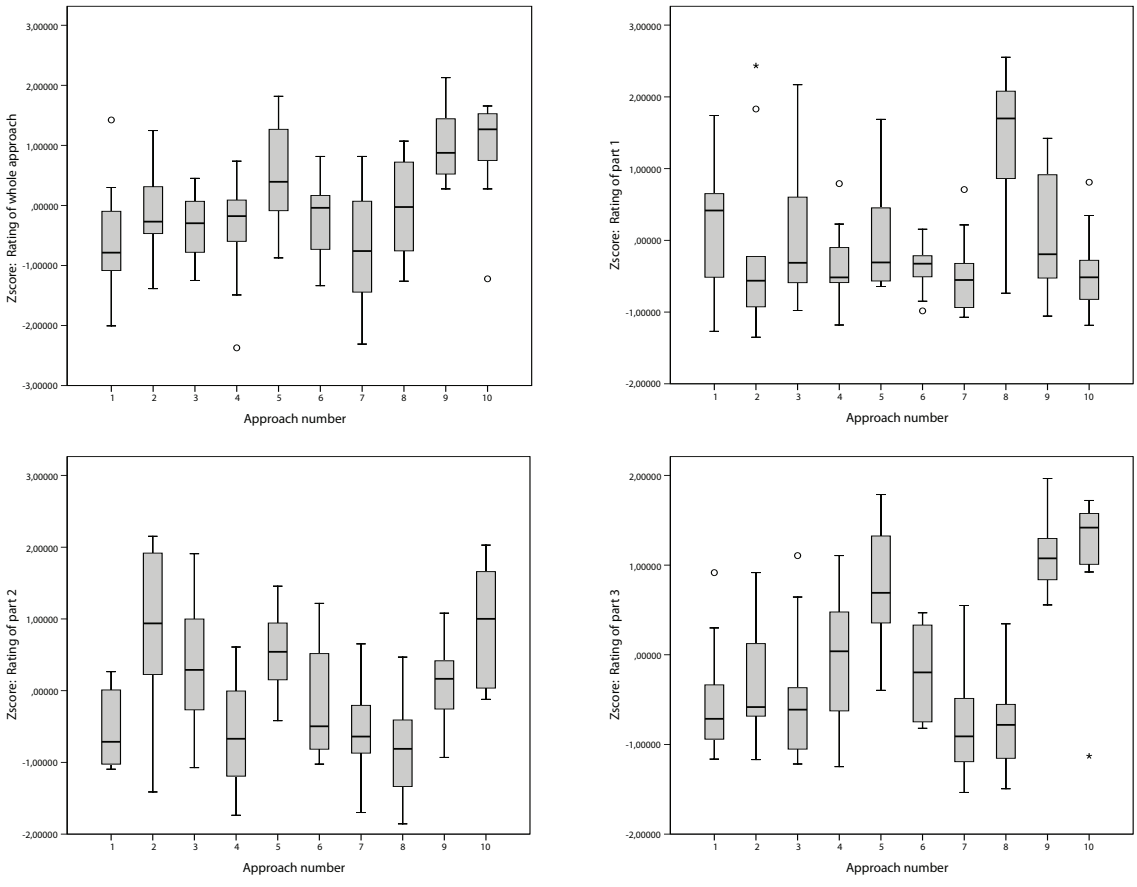


Figure 11. Boxplots of ratings approaches (SIMONA research simulator test and Citation combined).

During the real flights this was not the case: once the pilot monitoring received a call from ATC all attention was diverted to ATC contact and all other activities (such as performing checklists, selecting flaps, corresponding to calls from the pilot flying) stopped.

5.2 Factors that influence pilot TDL

Six sets of subjective data are available to determine which factors influence pilot TDL. These six sets are:

1. The RSME *z*-scores for the entire approach, and
2. the RSME *z*-scores for the three approach parts. Since the effect of the different approaches on the RSME *z*-scores was the same for both the flight simulator tests and the real flight tests, the RSME *z*-scores for both experiments are combined, in order to yield one larger data set. As a result each approach now has 12 RSME ratings. The boxplots of this combined set of ratings can be found in Fig. 11. To analyse whether there was a difference between the RSME *z*-scores for any approach (part) and the benchmark approach (approach 1), paired samples t-test were used when the RSME *z*-scores were parametric, and Wilcoxon matched pairs tests were used when they were non-parametric.
3. The run questionnaires from the flight simulator tests. The pilots' answers to the questions regarding the difficulty (see Fig. 8 for an example) are converted to z-scores. To determine whether there was an effect of a factor, the answers given for one specific approach are compared to the answers given for the benchmark approach. The comparison is again performed by using paired samples *t*-tests and Wilcoxon matched pairs test.
4. The answers from pilots to the brief in-flight interview during the real flight tests. To analyse these answers the majority rule is used: when more than three of the pilots have the same opinion on a matter, this opinion is used for further analysis.
5. The answers to the end-of-day questionnaires for the flight simulator tests, and
6. The answers to the end-of-day questionnaires for the real flight tests. These two end-of-day questionnaires are mostly identical. It is interesting to see whether the general opinion of the pilots changes between the two test series. Again, the majority rule is used to analyse the answers.

5.2.1 Overview per approach

Table 3 shows the results when comparing each approach to the benchmark approach, when using the results from the first four datasets.

To explain the idea behind Table 3 an example for APP02 is given here. The data for each approach are compared to the data of the benchmark approach (APP01). The first comparison is between the RSME *z*-scores for the entire approach for APP02 and APP01. The result of the *t*-test is that there is no significant difference ($t(11) = -1{\cdot}649$, $p > {\cdot}1$). For the second test the RSME *z*-scores of part 2 of the approach are used, since the independent measure for APP02 is a short LOC intercept heading. The result of the t- test is that there is a significant difference between the two ratings ($t(11) = -3{\cdot}003$, $p < {\cdot}05$). On average the ratings of part 2 of APP02 are significantly higher than the ratings of part 2 of APP01. Hence the table states 'more effort'. The next column is the flight simulator run-questionnaire. In this particular case only the question about the length of the LOC intercept heading is important. The standardised data for this question are normally distributed, so again a paired samples t-test is used. The result is that a short LOC intercept

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

heading significantly increases the amount of effort needed ($t(5) = 2{\cdot}965$, $p < {\cdot}05$). Finally the in-flight interview from the real flight tests is reviewed on this issue. All six pilots stated that this short LOC intercept heading increased the amount of effort needed.

Since three out of four sets of data indicate 'more effort' the conclusion can be drawn that APP02 requires more effort than the benchmark approach, APP01. For this reason the final column reads 'more effort' for approach 2. Using this method for each approach gives the results in Table 3.

5.2.2 Overview per factor

Table 3 thus gives the overview per approach. However, per approach there were sometimes more factors that were changed relative to the benchmark approach than just the independent variable, see Table 1. Therefore it is interesting to also consider the effects per factor. For this reason Table 4 is created. This table lists all the specific factors that are investigated in this research in the first column.

The second column lists the approach numbers where the factor mentioned in the first column occurred. The next four columns list the results of the flight simulator test run-questionnaires, the real flight test in-flight interviews, the end-of-day questionnaire for the flight simulator tests and the end-of-day questionnaire for the real flight tests respectively. The fifth column indicates the results according to the RSME z-scores for that particular part of the approach that is connected to the factor. There is only an entry in this column when there was no other linked factor (see Table 1) that was also changing in that part of the approach. The same procedure is used as before: when a certain factor is identified to increase the effort at least in threeof the data sets, it is concluded that this factor increases the effort.

An example: the first entry in the table is 'Eratio more than 1, in part 1 of the approach'. This only occurs in APP08. Using the data from APP08, statistical tests are performed on the answers to the flight simulator run questionnaires. A paired samples t-test is conducted on the answers of the question in the run questionnaire (regarding the ability to meet constraints at waypoints in the first part of the approach, see Fig. 8 for an example of the format). The results indicate that this significantly increases the effort ($t = 2{\cdot}939$, $p < {\cdot}05$). During the in-flight interview of the Citation tests four pilots answered that this aspect increased the amount of effort. The same answers are found from the two end-of-day questionnaires. Additionally, in APP08 the only factor that was changed in the first part of the approach relative to the benchmark approach was the Eratio in the first part of the approach, there were no linked factors. Therefore in the fifth column the result of the RSME z-scores for the first part of the approach is also incorporated (taken from Table 3). So, in the case of 'Eratio more than 1, in part 1 approach' all the data concur, this factor increases the effort.

In several parts of the table it states 'no information available'. This occurs frequently in the column of the Citation in-flight interview. These interviews were so short that only a limited amount of information could be gathered. In some cases the end of day questionnaires do not contain a question that specifically handles a factor. So these cases are also noted as 'no information available'. When this occurs twice in one row and the other two columns do state 'increase in effort', the conclusion is that that particular factor does increase the effort. An example is 'short LOC intercept heading'. Unfortunately there was no question incorporated in the end of day questionnaires that specifically targeted this aspect. But from the flight simulator test run questionnaire and the real flight test in-flight interview it was found that this factor does increase the effort. In this case only two out of four are needed for a positive conclusion.

Very interesting to note in Table 4 are the factors: 'High LOC intercept speed' and 'short leg IF-FAF'. From the comparisons of the RSME ratings of APP06 (which has a short leg IF-FAF) it was concluded that this approach does not require significantly more effort to fly. However, APP06 is not the only approach with a short IF-FAF leg. As can be seen in the second column of Table 4, APP05 and APP10 also have a short leg IF-FAF (among other factors). The results of the Wilcoxon test on the answers of the question in the run questionnaire regarding the length of leg IF-FAF of APP05 and APP10 showed a significant difference compared to the benchmark approach. This difference is not found on APP06. From this it seems that a short leg IF-FAF only increases the effort when an approach also has a short line-up distance (like APP05 and APP10). In APP06 the line-up distance is normal, so pilots have enough time to 'recover' from the short leg IF-FAF. However in the end of day questionnaires 5 out the 6 pilots answered 'more effort' on the question regarding the effect of a short IF-FAF distance. Taking this information into consideration it is concluded that in the researchers' opinion a short leg IF-FAF does increase the TDL, but the increase in TDL is limited and only really occurs when the line-up distance is short as well.

The same conclusion is made regarding the high LOC intercept speed. No increase in TDL is found when investigating APP07 (which has a high LOC intercept speed, but with a long line-up distance). But analysis of APP09 and APP10 shows that a high LOC intercept speed accompanied by a short line-up distance (little time to recover on final), does indeed increase the TDL.

Interesting to note is the last row, where the factor 'horizontal intercept of the GS' is stated. This aspect occurs in APP08. This approach is designed in such a way that instead of a CDA, this approach has several 'step-down's' (legs where the aircraft descends, followed by legs where the aircraft flies level). In APP08 the leg before the GS is flown level and as a result the GS is intercepted horizontally. From the flight simulator test run questionnaire and the Citation real flight tests in-flight interviews it can be concluded that this in fact decreases the effort during an RNAV approach.

5.2.3 Conclusions TDL factors

Analysing the factors from Table 4 and taking the observations made by the authors during all the tests (both flight simulator and real flight tests) into consideration, it can be concluded that especially the final part of the approach (the glideslope) has great influence on the TDL of pilots. An Eratio of more than one in the beginning of an approach (and not being able to meet constraints on waypoints as a result) increases the TDL to some extent. However, when this occurs in the final part of the approach the increase in TDL is much more significant.

The factors concerning the LOC intercept (length of localiser intercept heading, angle and speed) increase the TDL slightly. When pilots have a long line-up distance to 'recover', the increase in TDL is even less.

A short line-up distance and low FAF are of great influence to the TDL. Even when the Eratio on the glideslope is less than one (and a stabilised approach is possible), a short line-up distance (and a low FAF) results in very limited time available to perform the necessary pilot actions. This 'limited time available' (which is always a consequence of one of the above stated factors), is indeed the driving force on increasing the TDL. It was observed during both experiments that due to limited time, pilots were often late with SOPs and occasionally even forgot certain actions altogether.

To summarise the analysis of this section, the following list of factors increase the TDL during an RNAV approach:

Limited increase in TDL:

- Eratio more than one (before LOC intercept heading)
- Short LOC intercept heading
- Large LOC intercept angle
- High LOC intercept speed (in combination with other factors, e.g. short line-up distance)
- Non-horizontal intercept of the Glideslope

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

High increase in TDL:

- Short leg IF-FAF (in combination with other factors, e.g. short line-up distance)
- Short line-up distance
- High speed on FAF
- Low FAF
- Eratio more than one (after LOC intercept heading)
- Not stabilised at 1,000ft

The factors that are found to influence pilot TDL for the Cessna Citation were found to be identical to those found for the B747. Except for the fact that for the Citation experiments pilots indicated that a horizontal intercept of the Glideslope resulted in a decrease in effort, whereas for the B747 experiments the results showed that this did not have an influence on the effort. This difference can be explained by the fact that during the B747 experiments pilots had a VNAV mode available, guiding them correctly towards the Glideslope intercept independent of the altitude profile, whereas pilots during the Citation experiment had not. It is therefore stated that for aircraft with VNAV mode the Glideslope intercept does not influence pilot TDL.

The hypothesis that the list of factors that influence pilot TDL and the guidelines to keep pilot TDL at an acceptable level as found for the B747 are also valid for other aircraft types, thus is a reasonable one. It is noted again that the factors 'FAF altitude' and 'time available on Localiser intercept heading' are factors that originate from the fact that both for the B747 and for the Citation a number of actions need to be performed on Localiser intercept heading and between the FAF and 1,000ft. If, for other aircraft, these actions are performed at a different location in the approach, then there should be sufficient time at that particular location.

6.0 FLIGHT MECHANICAL TOOL

For the B747 a comprehensive Monte Carlo simulation and a Point Mass Model (PMM) simulation were developed which could, when combined with the regulations in the PANS-OPS[17], for a given approach predict whether the guidelines to keep pilot TDL at an acceptable level are met. It is now investigated whether these two simulation models can also provide predictions for another aircraft, in this case the Cessna Citation. To this end, first the general idea behind the Monte Carlo computer simulation[11] is repeated here, and the necessary adjustments to incorporate the Cessna Citation are explained. After that, a case study is considered for the Cessna Citation. To conclude the PMM simulation is briefly mentioned.

6.1 Monte Carlo computer simulation

When a (newly designed) approach is entered into the Monte Carlo computer simulation, the simulation predicts, amongst others, the percentage of flights that will meet the constraints at the waypoints, and the percentage of flights that will result in a stabilised approach at 1,000ft, both factors proved to have a significant influence on pilot TDL. It also predicts under what circumstances (e.g., wind conditions) this can be achieved. This section will describe the aircraft model, pilot model, SOPs, wind model and turbulence model that are used within the Monte Carlo computer simulation.

6.2 Computer simulation input

The input of the Monte Carlo computer simulation exists of a list of waypoints of a (newly designed) approach, defined by their lat-lon co-ordinates, and the altitude and speed constraints at these waypoints. Additionally, the user has to define which waypoint in the list is the Final Approach Fix (FAF).

6.3 Aircraft (Cessna Citation) and autopilot models

The Monte Carlo computer simulation[11] was set-up in a modular way. The B747 aircraft is replaced by the model for the Cessna Citation I (500)[18] which is exactly similar to the model used in the SIMONA flight simulator. Just as in the flight simulator models (and identical to the flight simulator models), the autopilot models are based on the autopilots developed for the model of the B747[11], which were derived from[19]. Autopilot modes included are: LNAV, Altitude hold, Vertical Speed Select, Glideslope, Heading Select and Localiser modes. The LNAV mode is based on the VOR modes described in Ref. 19. As the Cessna Citation II Laboratory aircraft does not contain a VNAV mode the aircraft is modeled to descend from waypoint to waypoint in the vertical speed mode where the selected (calculated) vertical speed depends on the constraints at the waypoints and the wind conditions.

The hierarchy in meeting the constraints at the waypoints is as follows: the Autopilot models will always aim to meet the altitude constraints at the waypoints, second to this, the Autothrottle controls the airspeed. This results in the situation that the altitude constraint at the next waypoint will always be met, while the speed constraint might not be met (airspeed might be higher than required).

6.4 Pilot model and standard operating procedures for the Cessna Citation laboratory aircraft

All pilot actions such as selecting flaps are modeled according to the SOPs for the Cessna Citation (see Fig. 6). Each of the pilot actions prescribed by the SOPs is modeled using a 'trigger' event (e.g. reaching 1,200ft) and a reaction time representing the time between reaching the trigger event and actually performing the action (e.g., 2 seconds after reaching 1,200ft, flaps LAND are selected). These reaction times are modeled as normal distributions and are based on the distributions of the reaction times as obtained from the flight simulator tests and real flight tests. The trigger events and corresponding reaction time normal distributions as used in the computer simulation are given in Table 5.

Table 5
Trigger events and reaction time distributions for pilot actions in Monte Carlo simulation

Pilot action	Trigger Event	Mean	Standard deviation
Switch to heading select mode	Waypoint capture at start Localizer Intercept Heading	16.4 sec	9.4 sec
Arm approach	Waypoint capture at start Localizer Intercept Heading	22.2 sec	10.1 sec
Disengage autopilot	Waypoint capture at start Localizer Intercept Heading	23.5 sec	10.6 sec
Flaps APPROACH	2.0 nm before reaching FAF	10.3 sec	13.7 sec
Gear Down	Reaching FAF	0.5 sec	6.9 sec
LAND checklist start	Gear Down	14.3 sec	12.8 sec
(Checklist duration	-	20.9 sec	8.7 sec)
Flaps LAND	Reaching 1,200' above airport elevation	1.5 sec	4.1 sec

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

It is important to note that if the airspeed constraints at the waypoints required an airspeed lower than the instantaneous flap speed mark (the IAS below which the next flap setting needs to be selected), the next flap setting is selected in the Monte Carlo simulation irrespective of SOPs. Also, when the airspeed exceeds the placard speed of a certain flap setting, this flap setting cannot be selected. Gear Down selection has no upper speed limit.

The Cessna Citation II Laboratory aircraft does not contain an autothrottle. In the model the airspeed is regulated using a simple proportional controller to simulate the pilots' manual throttle control. Additionally, a relatively simple pilot manual control model for the flight director task is added, consisting of only a time delay (equal to 0·3 seconds) and pure gain.

6.5 Turbulence and wind models

A patchy turbulence model is used within the Monte Carlo simulation. The intensity of the turbulence can be adjusted with a gain. During one Monte Carlo simulation run the turbulence intensity, wind direction and wind speed are constant throughout the entire approach, between different Monte Carlo runs these are varied.

6.6 Outputs of the computer simulation

The constraints at a waypoint were considered to be met when the actual Indicated Airspeed (IAS) at that waypoint (to be more specific: the point in the trajectory closest to that waypoint) was less than the required IAS plus 10kt, and the actual altitude at that waypoint was less than the required altitude plus 100ft. A lower boundary for these constraints is not necessary since the Monte Carlo simulation always regulates the airspeed and altitude towards the constraints at the next waypoint, when the required airspeed and altitude are attained, the Monte Carlo simulation maintains the required airspeed and altitude until the waypoint is reached. Therefore, in the Monte Carlo simulation, the altitude and airspeed will never be too low at a waypoint.

To determine whether a Monte Carlo simulation run of the approach resulted in a stabilised approach the criteria given earlier were quantified as follows:

- Heading change, and roll rate are within 5deg/s
- The IAS is not more than VREF + 20kt;
- Flaps LAND are selected, landing gear is down;
- Sink rate is not larger than 1,000ft per minute; and
- Localiser and glide slope are within one dot;

These criteria are evaluated exactly at 1,000ft above airport elevation.

6.7 Validation of Monte Carlo simulation with experiment data

To properly validate the predictions of the Monte Carlo simulation, a comparison is made with the results obtained during the two experiments (flight simulator test and real flight test). During these experiments flight data were recorded and analysed to evaluate how many flights were stabilised at 1,000ft, to determine how often the constraints were met at the waypoints and to gain insight into all pilots' actions. The Monte Carlo simulation will be run with exactly the same conditions as encountered during each of the runs during the two experiments, and while simulating the pilots' actions at exactly the same moments in time at which the pilots performed their actions during the experiments. This implies that the RNAV routes, aircraft weight, wind speed and wind direction, QNH pressure levels, airport elevation, runway heading and pilot action times (the moment that SOPs are performed) are simulated exactly identical to the experiment situation. This results in a total of 12 simulated Monte Carlo runs per approach (since each approach was flown only once by each pilot during each experiment).

The results of the Monte Carlo simulation are presented in Fig. 12 as the percentage of approaches that are predicted to be stabilised at 1,000ft.

Table 6
Waypoint constraints met during tests compared to the simulations with input parameters as measured during test runs (*N* = 6)

	2nd WP (APP08 only)		3rd WP (APP08 only)		2nd WP or 4th WP (APP08 only)		Begin LOC Intercept hdg		IF		FAF	
	SRS tests	Monte Carlo simulation	SRS tests	Monte Carlo simulation	SRS tests	Monte Carlo simulation	SRS tests	Monte Carlo simulation	SRS tests	Monte Carlo simulation	SRS tests	Monte Carlo simulation
APP01					100%	100%	100%	100%	100%	100%	100%	100%
APP02					100%	100%	**83%**	**100%**	**67%**	**100%**	**83%**	**100%**
APP03					100%	100%	100%	100%	**83%**	**100%**	**67%**	**100%**
APP04					100%	100%	100%	100%	100%	100%	**83%**	**100%**
APP05					100%	100%	**83%**	**100%**	100%	100%	**83%**	**100%**
APP06					100%	100%	100%	100%	100%	100%	100%	100%
APP07					100%	100%	**83%**	**100%**	100%	100%	**67%**	**100%**
APP08	100%	100%	0%	0%	100%	100%	0%	0%	100%	100%	100%	100%
APP09					100%	100%	**83%**	**100%**	100%	100%	**83%**	**100%**
APP10					100%	100%	100%	100%	0%	0%	0%	0%
	Citation flight tests	Monte Carlo simulation	Citation flight tests	Monte Carlo simulation	Citation flight tests	Monte Carlo simulation	Citation flight tests	Monte Carlo simulation	Citation flight tests	Monte Carlo simulation	Citation flight tests	Monte Carlo simulation
APP01					100%	100%	100%	100%	100%	100%	**83%**	**100%**
APP02					100%	100%	100%	100%	100%	100%	**67%**	**100%**
APP03					100%	100%	100%	100%	100%	100%	100%	100%
APP04					100%	100%	100%	100%	100%	100%	**83%**	**67%**
APP05					100%	100%	100%	100%	100%	100%	**67%**	**100%**
APP06					100%	100%	100%	100%	100%	98%	**83%**	**100%**
APP07					100%	100%	100%	100%	100%	100%	100%	100%
APP08	100%	100%	0%	0%	100%	100%	**50%**	**0%**	100%	100%	100%	100%
APP09					100%	100%	100%	100%	**83%**	**100%**	100%	100%
APP10					100%	100%	100%	100%	**17%**	**0%**	0%	0%

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

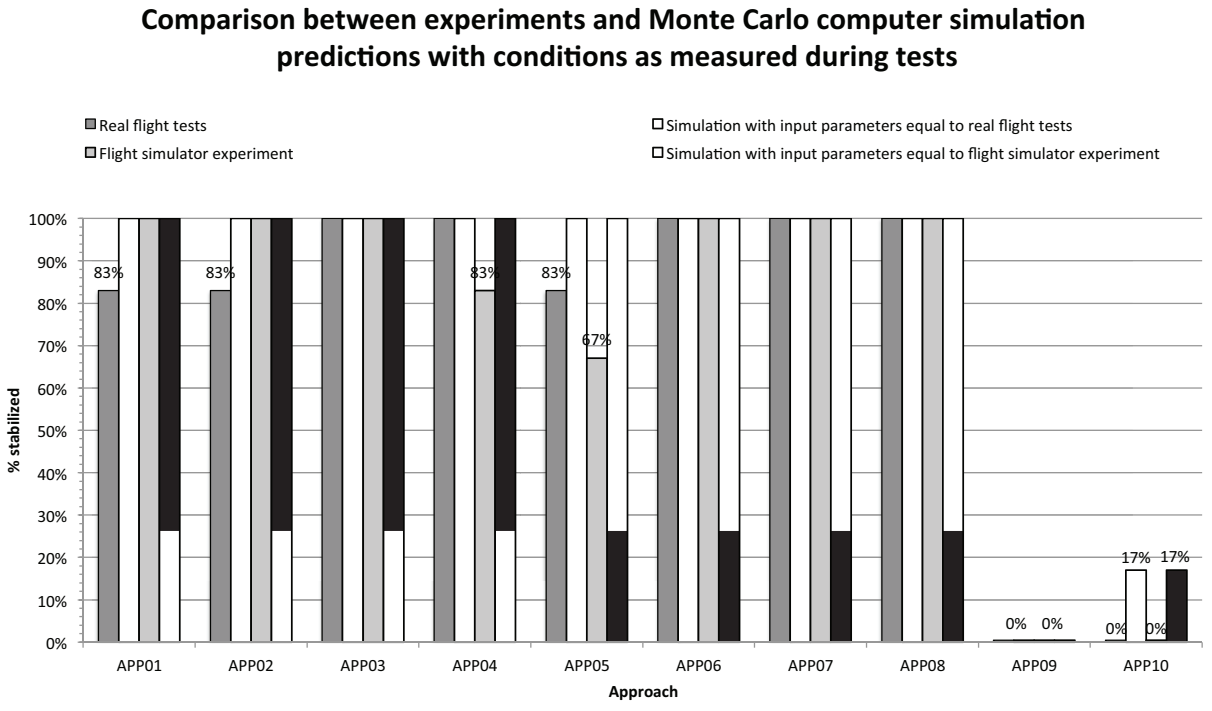


Figure 12. Stabilised approaches during tests compared to the simulations with environment input and reaction times equal to every test performed in the SRS and during the Citation flight tests.

As Fig. 12 shows, there is still a discrepancy between the predictions of the Monte Carlo computer simulation and the actual flight simulator test and real flight test results. Excessive airspeed was always one of the violated conditions in the unstabilised approaches in the experiments. Even though APP01, APP02, APP03 and APP05 are no high energy approaches, the airspeed at 1,000ft still exceeded the VREF + 20kt limit during the flight simulator tests and real flight tests resulting in unstabilised approaches. These large variations in airspeed can all be traced back to human performance in throttle control. Figure 13 shows the airspeed profiles during APP01 as an example to indicate the difference in variation between a throttle controlled by a human pilot and by the simple throttle model in the Monte Carlo simulation. It can be seen that the simple throttle model in the Monte Carlo computer simulation regulates the airspeed more strictly than the pilots did during the experiment.

In the case of APP10, a high-energy approach, the model simulates a stabilised approach, even though none of the approaches during the flight simulator tests and real flight tests were stabilised approaches (all due to an excessive airspeed at 1,000ft above airport elevation). This difference is related to the slightly larger variation in manual altitude control by the pilots compared to the altitude control of the model, resulting in slightly different altitude profiles. For this particular case the altitude profile during the flight simulator test was steeper than the three degree glideslope. If, for this case, flaps 'land' were selected at 1,500ft this occurred at a later moment in time for the steep approach during the flight simulator test than during the three degree glideslope approach simulated in the Monte Carlo computer

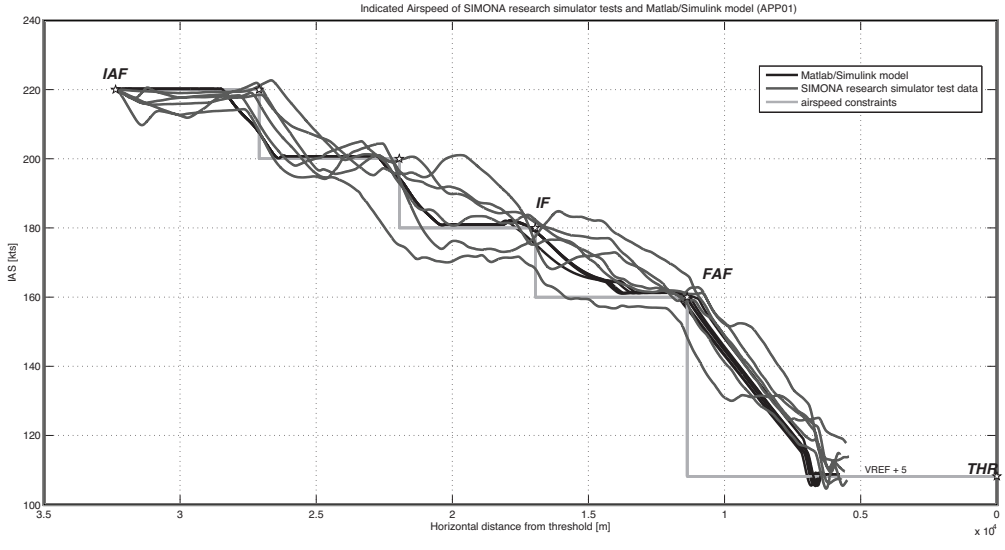


Figure 13. Airspeed profile of APP01.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

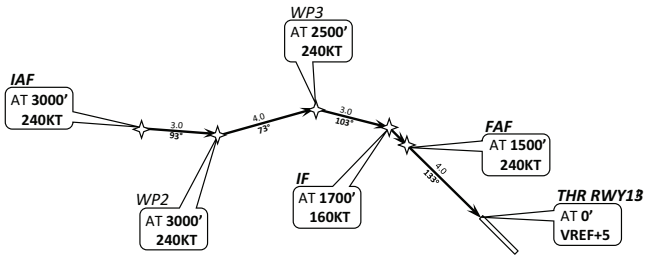


Figure 14. Example approach towards runway[13].

Table 7
Waypoints with altitude and airspeed constraints for example approach towards runway[13]

Waypoint	Altitude [ft]	IAS [knots]
IAF	3000	240
WP2	3000	240
WP3	2500	240
IF	1700	160
FAF	1500	130
THR RW13	0	VREF + 5

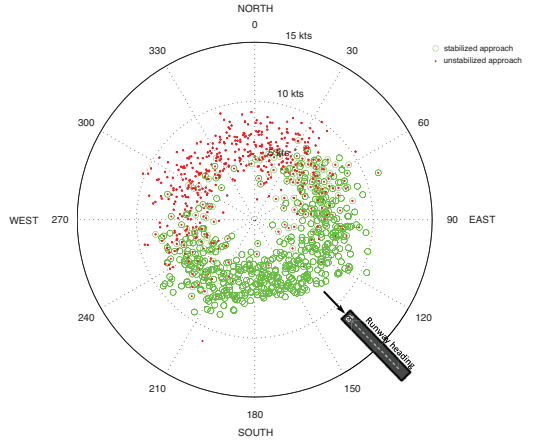


Figure 15. Example of the results of the Monte Carlo computer simulation with respect to the possibility of achieving a stabilised (grey circle) or unstabilised (black dot) approach, as a function of wind direction and wind speed.

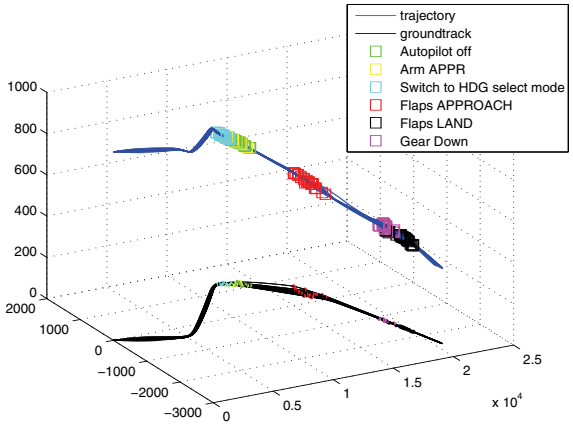


Figure 16. Example of the results of the computer simulation, providing insight into the locations in the approach where pilots are performing actions such as selecting flaps, gear, etc.

simulation. This results in a stabilised approach in the computer simulation and an unstabilised approach during the flight simulator tests. This is the reason for the discrepancy for APP10 in Fig. 12. Table 6 shows how many times the constraints at the waypoints have been met. The grey cells indicate the fact that a large amount of energy (altitude and/or airspeed) before that waypoint needed to be dissipated in order to meet the constraints. Again, there are discrepancies to be found in how many times waypoint constraints have been met during the experiments compared to the output of the simulation. However, as with the prediction of stabilised approaches, the reason for the differences lies in the unpredictability of the manual throttle control by the pilots.

It can thus be concluded that the predictions of the Monte Carlo computer simulation model are not always fully in agreement with reality, but that the trends are predicted rather well. The major cause for discrepancies is the lack of correct simulation of the pilots' manual control of the throttle. It should be noted that the method and computer simulation developed within this research will be used to give an indication of pilot TDL during RNAV approaches at airports. The majority of these aircraft will have an autothrottle, and therefore it is expected that better results will be obtained for these aircraft.

6.8 Case study

The Monte Carlo computer simulation can now be run many times, using all recorded reaction times, and applying different wind conditions. The computer simulation can then be used to predict under what wind conditions the constraints at the waypoints can be met and a stabilised approach can be achieved. As an example an approach towards runway 13 is considered, see Fig. 14 and Table 7. The results are depicted in Fig. 15 and Fig. 16. Fig. 15 clearly shows that the possibility of achieving a stabilised approach depends on the wind direction (a headwind on final results in a stabilised approach). Other reasons for ending stabilised or unstabilised are the moment in time at which flaps LAND and/or gear down are selected. A similar plot as in Fig. 15 can be generated for meeting the constraints at each waypoint. The results of the computer simulation also provide insight into the locations in the approach where pilots are performing many actions, thereby providing approach designers with an indication of the 'busy' parts of an approach, see Fig. 16 for an example.

6.9 Point mass model

The Monte Carlo simulation model generates reliable results regarding the percentage of flights that can achieve a stabilised approach and factors that influence pilot TDL, but takes a long time to produce these results (in the order of several hours per approach that is analysed). This is not very practical when the computer simulation is intended to be used as a tool during the design of approaches. Therefore it was investigated[13] for the B747 whether a much simpler model based on a point mass model, with a considerably shorter calculation time, can generate results as reliable as the highly detailed computer simulation. It was found for the B747 that a point mass model could indeed generate the same results as the more detailed computer simulation. It can now be stated that this is also the case for the Cessna Citation. This is true as long as the point mass model contains:

1. a detailed lift-drag polar for all flap settings and gear up/down setting,
2. a detailed model of the flight idle thrust,
3. an accurate model to simulate the lateral track, specifically the distance of turn anticipation since this influences the amount of trackmiles available between two waypoints, and
4. a model to simulate the pilots' actions according to the trigger events and reaction time distributions found in this research.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

6.10 Are all design guidelines predicted by the computer simulation?

The question now is, whether the output of this computer simulation provides sufficient information in order to assess whether the approach meets all guidelines for the contributors to pilot TDL as given earlier in this chapter. The simulation obviously predicts whether a stabilised approach can be achieved (guideline 3), and whether the constraints at the waypoints can be met (guideline 1). It also provides insight whether there is sufficient time on Localiser intercept heading since the moments in time at which all actions are performed is predicted (guideline 2), and it can easily predict the sink rate (guideline 5).

However, although the simulation can predict or calculate the numerical values for the (actual) Localiser intercept speed (guideline 6), the Localiser intercept angle (guideline 7), and the distance between IF and FAF (guideline 4), it does not give a qualitative indication of whether these numerical values are sufficiently high or low. Fortunately, the minimum or maximum values for these factors are very accurately prescribed in the Procedures for Air Navigation Services Aircraft Operations (PANS-OPS)[17]. The PANS-OPS prescribe a minimum straight distance between IF and FAF of 2nm with an additional turning distance (which depends on the airspeed and intercept angle), and recommend an interception angle at the Localiser not exceeding 30 degrees. Actually, the PANS-OPS and the predictions of the computer simulation complement each other very nicely regarding factors contributing to pilot workload, since what is not prescribed in the PANS-OPS is predicted by the computer simulation and vice versa.

The conclusion thus is that the computer simulation, combined with the regulations in the PANS-OP, provides sufficient information to assess whether the guidelines for the contributors to pilot TDL for the Citation are met.

6.11 Quantification of guidelines

In order to be able to assess whether, for instance, sufficient time is available on Localiser Intercept heading, the time required in this part of the approach needs to be quantified. Table 8 gives an indication of the reaction times for all pilot actions for the Citation based on the reaction times recorded during the two experiments. As an example: 41 seconds after starting the turn towards Localiser intercept heading, pilots had switched off the autopilot during 95% of all runs. From this it can be concluded that, when it is assumed that 95% of the pilots should be able to fly the approach, the time available on Localiser intercept heading should at least be 41 seconds. These values can also be used in the point mass model.

It is now interesting to compare the values found for the Citation (Table 8) to the values found for the B747[13], see Table 9. The only actions that are based on the same trigger event for both aircraft are the actions Flaps LAND and Gear down (the reaction time for the B747 stated to the left of the forward slash is of importance here). It is important to note that the reaction times given for each percentage of pilots in both tables are based on the actual recorded reaction times, not on the approximating normal distribution curves.

It can be seen that the reaction times for the B747 are in all cases larger than for the Citation. This might be caused by the following two facts: first, the sample size for the B747 was much larger (see also the histograms in Figure 17 as an example), and second, the pilots participating in the Citation experiments were pilots that were used to participate in scientific research and flight tests, the pilots participating in the B747 experiment were commercial airline pilots. For these reasons, the reaction times for the B747 given in Table 9 are regarded more reliable and more realistic.

To illustrate the difference between the reaction times for the B747 and Citation, Fig. 17 is presented. Mann-Whitney, Wald-Wolfowitz and Kolmogorov-Smirnov tests all indicated that the

Table 8
Maximum pilot reaction times in seconds for different percentages of pilots based on 120 samples for the Cessna Citation

Pilot action	Trigger Event	85 %	90 %	95 %	99 %
Heading Select	start turn to LOC int. HDG	23.1	25.3	29.6	55.2
ARM Approach	start turn to LOC int. HDG	30.5	33.6	38.3	57.0
AP OFF	start turn to LOC int. HDG	31.7	35.4	41.0	58.8
Flaps APP	2 NM before FAF	21.5	29.6	38.1	44.3
Gear Down	Reaching FAF	5.6	7.3	9.2	19.8
Flaps LAND	1200ft	5.5	7.5	10.6	14.5
Landing Checklist	Gear Down	55.2*	60.2*	70.2*	75.7*

* Time in seconds until the checklist was completed

Table 9
Maximum pilot reaction times for different percentages of pilots for the B747(13)

Pilot action	Trigger Event	85 %	90 %	95 %	99 %
Flaps 1	IAS	VREF+75.5kts	VREF+73kts	VREF+70.2kts	VREF+50kts
Flaps 5	IAS	VREF+57.3kts	VREF+56.2kts	VREF+54.2kts	VREF+43.3kts
Flaps 10	End turn to LOC int. HDG	33.1 s	38.7 s	62.9 s	86.1 s
Heading Select	End turn to LOC int. HDG	24.0 s	30.8 s	45.4 s	147.0 s
ARM Approach	End turn to LOC int. HDG	37.1 s	44.8 s	57.6 s	130.0 s
AP & AT Off	End turn to LOC int. HDG	30.8 s	38.8 s	47.3 s	65.2 s
Gear Down	Reaching FAF	9.6 / -8.5 s	11.3 / 1.3 s	15 / 2.5 s	27.8 / -7 s
Flaps 20	Gear Down	-9.6 / 14.9 s	-6.9 / 8.3 s	-15 / 12.8 s	-26.8 / 30.4 s
Flaps LAND	Reaching 1,200’	6.4 s	9.6 s	15.3 s	23.4 s
Approach CL	Transition Level	TL – 333ft	TL – 350ft	TL – 400ft	TL – 1,220ft
Duration		9.8 s	10.4 s	11.5 s	16.5 s
CL after GD/flaps20	Latest of GD or Flaps 20	23.6 s	34.2 s	35.7 s	65.6 s
Duration		13.9 s	15.3 s	16.7 s	21.8 s
CL after flaps 25	Flaps 25	8.4 s	9 s	12.3 s	13.0 s
Duration		9 s	9.3 s	11.5 s	12.1 s

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

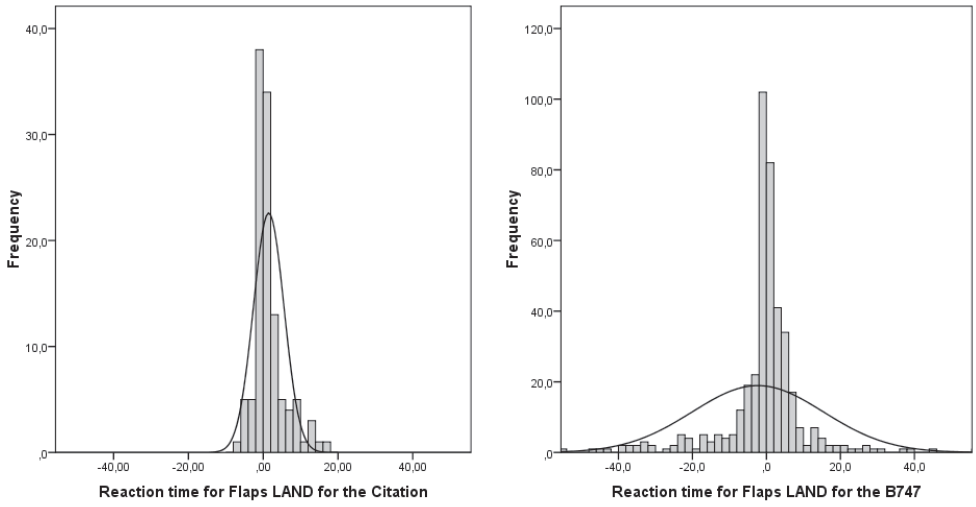


Figure 17. Histograms for the reaction times for selecting Flaps LAND for the B747 (left, *N* = 116) and Citation (right, *N* = 421).

distributions for the B747 and the Citation were significantly different.

For clarity, the exact quantification of the guidelines as used during this research is given below (numbers correspond to numbered guidelines):

1. The constraints at the waypoints are considered to be met when the airspeed is within 10kt of the required airspeed and the actual altitude is within 100ft of the required altitude.
2. Using the reaction times for the actions performed on Localiser intercept heading as given in Table 9 it can be determined (for a percentage of pilots defined by the approach designer) whether all actions that should be performed on Localiser intercept heading are actually performed on Localiser intercept heading. When this is the case this is regarded 'sufficient time' to perform all actions.
3. An approach is considered stabilised at 1,000ft when:
   - Heading change and pitch change are within 5deg/s;
   - The IAS is not more than VREF + 20kt;
   - Flaps 25 are selected, landing gear is down (whether this is the case for the defined percentage of pilots can be derived from the values in Table 9);
   - Sink rate is not larger than 1,000ft per minute;
   - Localiser and glide slope are within one dot ; and
   - All checklists are completed (whether this is the case can again be derived from the values in Table 9)
4. the distance between IF and FAF is considered sufficient when it meets the requirements in the PANS-OPS[17]
5. the vertical speed should be below the sink rate warning.
6. the Localiser intercept speed is considered 'not too high' when it meets the requirements in the PANS-OPS[17]
7. the Localiser intercept angle is considered 'not too large' when it meets the requirements in the PANS-OPS[17].

7.0 CONCLUSIONS

The list of factors that influence pilot TDL during RNAV approaches as found for the B747 are also applicable to the Cessna Citation. In other words: the same factors influence pilot TDL while flying an RNAV approach both for the B747 and for the Citation. This is concluded from the results from flight simulator tests and real flight tests with the Cessna Citation. Consequently, the guidelines to keep pilot TDL at an acceptable level as defined for the B747 also apply to the Cessna Citation. From this it is concluded that these guidelines are valid for aircraft types that fly approaches according to the assumptions set forth in this research, and that these guidelines can be used as such during the design of approaches.

It was also found that there were no discrepancies between the list of factors influencing pilot TDL resulting from the flight simulator tests and the list of factors resulting from the real flight tests.

Finally, the computer simulations that were developed for the B747 in order to predict whether the guidelines were met for a certain approach, were successfully adjusted to incorporate the Cessna Citation and produced reliable results. The same simulation technique can thus be used for different types of aircraft.

ACKNOWLEDGEMENTS

This research was financially supported by the Dutch Technology Foundation STW (DLR7071). We wish to thank the technical staff and pilots participating in the experiments, as well as Air Traffic Control at Malta International Airport for accommodating the test flights.

REFERENCES

1. STASSEN, H.G., JOHANNSEN, G. and MORAY, N. Internal Representation, Internal Model, *Human Performance Model and Mental Workload. Automatica*, 1990, **26**, (4), pp 811-820.
2. Flight Safety Foundation Approach-and-landing Accident Reduction (ALAR) Task Force. Killers in Aviation: FSF Task Force Presents Facts about Approach-and-landing and Controlled-flight-into-terrain Accidents, *Flight Safety Digest*, November 1998 – February 1999.
3. HEILIGERS, M.M., HOLTEN, TH. VAN and MULDER, M. Predicting pilot task demand load during final approach, *Int J Aviation Psychology*, October-December 2009, **19**, (4), pp 391-416.
4. VICENTE, K.J. *Cognitive Work Analysis, Towards Safe, Productive, and Healthy Computer-Based Work*, Lawrence Erlbaum Associates, Publishers, Mahwah, New Jersey, USA, 1999.
5. BARON, S., ZACHARIAS, G., MURALIDHARAN, R. and LANCRAFT, R. PROCRU: A model for analyzing flight crew procedures in approach to landing, NASA-MIT 16th Annual Manual, 1980, pp 488-520.
6. BARON, S., KRUSER, D.S. and HUEY, M.B. *Quantitative Modeling of Human Performance in Complex*, *Dynamic Systems*, National Academy Press, Washington DC, USA, 1990.
7. CORKER, K.M. Human Performance Simulation in the Analysis of Advanced Air Traffic Management, Proceedings of the 1999 Winter Simulation Conference, FARRINGTON, P.A., NEMBHARD, H.B., STURROCK, D.T. and EVANS, G.W. (Eds), 1999.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press

8. GORE, B.F. and CORKER, K.M. A Systems Engineering Approach to Behavioral Predictions of an Advanced Air Traffic Management Concept, 19th Digital Avionics Systems Conference, Entering the Second Generation of Powered Flight, Philadelphia, Pennsylvania, US, 2000.
9. SMITH, B.R. and TYLER, S.W. The Design and Application of MIDAS: A Constructive Simulation for Human-System Analysis, 2nd Simulation Technology and Training Conference, Canberra, Australia, 1997.
10. HEILIGERS, M.M., VAN HOLTEN, TH. and MULDER, M. Feasibility analysis of Achieving a Stabilized Approach, 26th International Congress of the Aeronautical Sciences, 2008.
11. HEILIGERS, M.M., van HOLTEN, TH. and MULDER, M. Flight mechanical evaluation of approaches, *J Aircraft*, **48**, (3), May-June 2011, pp 975-994, DOI 10.2514/1.C031188.
12. HEILIGERS, M.M., VAN HOLTEN, TH. and MULDER, M. Factors that influence Pilot Task Demand Load during RNAV Approaches, *J Aircraft*, **48**, (3), May-June 2011, pp 975-994, DOI 10.2514/1.C031188.
13. HEILIGERS, M.M., van HOLTEN, TH. and MULDER, M. Seven guidelines for limiting pilot task demand load area navigation approaches, accepted for publication in *J Aircr*.
14. HEILIGERS, M.M. Pilot Task Demand Load during RNAV approaches, PhD thesis, Ipskamp drukkers, The Netherlands, 2010.
15. HILBURN, B. and JORNA, P. Workload and Air Traffic Control. In HANCOCK and DESMOND, P.A. (Eds) Stress, Workload and Fatigue: Theory, Research and Practice, 2001, Hillsdale, New Jersey, US: Erlbaum.
16. GODLEY, S.T. Perceived Pilot Workload and Perceived Safety of RNAV (GNSS) Approaches, Australian Transport Safety Bureau, ATSB Transport Safety Investigation Report, Aviation Safety Research and Analysis Report 20050342, 2006.
17. International Civil Aviation Organization, *Procedures for Air Navigation Services Aircraft Operations, Volume II Constructions of Visual and Instrument Flight Procedures*, 5th ed, 2006.
18. BORST, C. Citast: Citation Analysis and Simulation Toolkit, Delft University of Technology, Delft, The Netherlands, 2004.
19. HANKE, R.C. and NORDWALL, DONALD R. The Simulation of a Jumbo Jet Transport Aircraft Volume II: Modeling Data, NASA report D6-30643, September 1970, Wichita, Kansas, USA.
20. Aircraft Operations Manual TU Delft Laboratory Aircraft.
21. ZIJLSTRA, F. and MEIJMAN, Th. *Het meten van mentale inspanning met behulp van een subjectieve methode, Mentale belasting en werkstress, een arbeidspsychologische benadering*, MEIJMAN, TH. (Ed), Van Gorcum, Assen/Maastricht, The Netherlands, 1989.
22. LODGE, M. Magnitude Scaling, Quantitative Measurement of Opinions, Sage university paper series on quantitative applications in the social sciences, 07-025, Newbury Park, CA, USA: Sage.
23. MEIJMAN, T., ZIJLSTRA, F., KOMPIER, M., MULDERS, H. and BROERSEN, J. *The Measurement of Perceived Effort*, Contemporary Ergonomics 1986, OBORNE, D.J. (Ed), Taylor & Frances, London, 1986.
24. HART, S.G. and STAVELAND, L.E. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research, *Advances in Psychology*, **52**, pp 139-183, HANCOCK, P.A., MESKATI, N. (Eds), Elsevier Science Publishers B.V., Amsterdam, The Netherlands, 1988.
25. JAMIESON, S. Likert Scales: how to (ab)use them, Medical Education 2004, **38**, pp 1217-1218, Blackwell Publishing, 2004.
26. KNAPP, T.R. Treating ordinal scales as interval scales: an attempt to resolve the controversy, *Nursing Research*, March/April 1990, **39**, (2), pp 121-123.

https://doi.org/10.1017/S0001924000006035 Published online by Cambridge University Press