

# The effect of strategic conformance on acceptance of automated advice: concluding the MUFASA project

Carl Westin and Brian Hilburn

Center for Human Performance Research, CHPR BV  
Voorburg, The Netherlands  
carl@chpr.nl, brian@chpr.nl

Clark Borst

Control and Simulation Division  
Delft University of Technology  
Delft, The Netherlands  
c.borst@tudelft.nl

**Abstract**—Whereas there has been a great deal of empirical and theoretical work into ATM automation in relation to reliability and traffic complexity, much less has been done in the area of similarities and differences between human-machine interaction and human-human interaction. Paradoxically, this could in the future become the most critical issue of all, as mismatches between human and machine could threaten acceptance of advanced automation. Through a series of human-in-the-loop simulations, the work described in this paper examined the interacting effects of air traffic complexity and strategic conformance, i.e., the fit between human and machine strategies, on automation acceptance in a conflict detection and resolution task. An experiment with 16 professional air traffic controllers showed that strategic conformance is a potentially important construct in human-automation interaction. Conformal resolution advisories were more accepted, led to higher controller agreement, and also reduced response time to proposed automated advisories. The results also indicated, however, that participants disagreed with conformal advisories in close to one quarter of the time. This might be evidence of “dispositional” automation bias, a reluctance to use automation under any circumstance.

**Keywords**— automation, acceptance, strategic conformance, decision aiding, SESAR, complexity, air traffic

## I. INTRODUCTION

Future Air Traffic Management (ATM) will have to rely on more, and increasingly sophisticated, automation to accommodate predicted air traffic. This entails automation growth in terms of types of tasks it can perform, and the level of authority and autonomy it can assume. What we do not yet fully appreciate is how to achieve this growth in automation, let alone which philosophical path to pursue. Should we opt for a completely automated ATM system that replaces human control? Or should the human retain some sort of final authority? This has been a central question in the SESAR Work Package E (WP-E), Long-Term and Innovative Research program. Whereas the scope of SESAR is built on a human-centred philosophy, where the human constitutes the core of the ATM system, the human role in the future ATM system is not necessarily as central. The HALA! research network, itself a progeny of WP-E purposed to specifically address ATM and automation, suggests a system-centred philosophy. In this view, neither the human nor automation is given full authority. Instead, human and automation are designed to cooperate for the best of the ATM system [1]. But

This work is co-financed by EUROCONTROL acting on behalf of the SESAR Joint Undertaking (the SJU) and the EUROPEAN UNION as part of Work Package E in the SESAR Programme. Opinions expressed in this work reflect the authors' views only and EUROCONTROL and/or the SJU shall not be considered liable for them or for any use that may be made of the information contained herein.



achieving true and successful teamwork between humans and automation, corresponding to that of human-human teamwork, is not easy. Klein et al [2] outlined ten challenges for successfully creating human-automation teams that satisfy the basic requirements of joint activity central for teamwork. A key characteristic of these challenges is the requirement for automation to better acknowledge and respond to individual differences. If automation fails to meet the needs and expectations of the human team member the team's performance may suffer as human acceptance and trust in automation decline.

Studies across various domains have shown that user acceptance of automation decreases when the authority of decision-making automation increases [3][4]. Research has also shown that the predominantly algorithmic approaches to automation seldom fit well with the more heuristic processes employed by humans. For example, in their exhaustive survey of conflict detection and resolution (CD&R) modelling methods, [5] concluded that CD&R automation correlates poorly with how controllers prefer to work. Bekier, et al. [6] suggested that there is a “tipping point” for any automated tool, above which controller acceptance of that tool quickly drops. Consequently, low user acceptance of an automated tool may lead to disuse (and abuse) of that tool which could severely undermine the intended safety and performance benefits of the tool [4].

The Multidimensional Framework for Advanced SESAR Automation (MUFASA) project argued that strategic conformance (i.e. the fit between the individual human and machine strategy) is essential for acceptance of decision support systems. It was hypothesized that a controller would more readily accept machine advice if the machine “thought” and “behaved” like the controller himself. This paper intends to answer the following questions: are controllers more likely to accept automated advisories when the advisory mimics their own solution? Will acceptance of automation vary with air traffic complexity, regardless of strategic conformance? Will workload increase with non-conformal resolution advisories? Does response time vary with strategic conformance of the advisory?

The next section expands on the notion of acceptance, providing a brief overview of the existing literature (sections II and III) and its implications for the research objectives of

MUFASA (section IV), namely using the human as a template for automation design. Sections V, VI and VII deal with the experimental design, results, and discussion, respectively, of the final large human-in-the-loop experiment in the MUFASA project. Section VIII concludes with implications for future research.

## II. ACCEPTANCE RESEARCH

Of pivotal importance for the future ATM system, is that the technologies developed are adopted and used by the end-user it intends to benefit and support. This requirement is so obvious, yet so often seems to be forgotten. Perhaps it can be attributed to the variety of approaches used, and the lack of clear and good frameworks translating theory to practice. Although acceptance and usage behaviour in relation to technology has been researched broadly across research communities, both interaction and knowledge transfer across communities have been sparse [7]. In general, acceptance research can be divided into three main approaches:

- Innovation-diffusion theory (IDT) – sociology, psychology, communication studies
- Acceptance research – information systems (IS)
- Usability research and human-machine interaction – cognitive ergonomics (CE)

The relevance and contribution of these research approaches for the purpose of this paper will be discussed in order.

### A. Innovation-diffusion theory

The innovation-diffusion theory (IDT) is considered one of the oldest and most influential theories in sociology and psychology research devoted to explaining why only a few innovations (i.e. ideas, practices or items) gain widespread adoption in a population while the majority of innovations fail. As such, IDT provides a model describing a generalised adoption process, typically with the individual adopter in focus, although complementary models have been developed extending the scope of IDT to, for example, organisations [8]. Different characteristics of the innovation and user are believed to influence the rate of the adoption process. Although innovation as a term is not limited to technological applications, IDT research has almost exclusively focused on technologies as the innovation. In IDT, one of the most influential characteristics of an innovation affecting adoption is *compatibility* [9]. Compatibility addresses how user perceived similarities between the technology and user influence acceptance and adoption. Aspects of compatibility are related to the values and norms of the user, similarities with previous technology used, and compatibility with user needs. As the early researchers of IDT were anthropologists, its methods and approaches have been much influenced by those used by anthropologists, such as qualitative immersion methods utilizing observations, interviews, case studies and field studies [8].

### B. Acceptance models

In the IS community, several acceptance models have been proposed addressing this willingness or intent to use (see [10] for an overview). Compared to IDT, which addresses acceptance at a higher level with a focus on issues of social diffusion, these acceptance models concentrate more directly on determinants and importance of user acceptance. Next to the IDT model, the IS community has developed the other predominantly used model addressing technology acceptance. The Technology Acceptance Model (TAM) relies on two parameters: the perceived usefulness (PU) of the technology, and the perceived ease of use (PEOU), which combined predicts behavioural intention to use (BI) and actual use [11]. The model focuses largely on measures of user psychological characteristics and technological characteristics for facilitating acceptance. Normally the methodology relies on questionnaires where the user responds to a series of statements. The variance in acceptance is derived through regression analysis of the underlying factors identified. TAM has been extensively adopted in a variety of domains, providing consistent and reliable results on technology acceptance. Over the years, researchers have modified the TAM for their specific purposes. Recently the TAM was synthesised with research on human-machine interaction to create the Automation Acceptance Model (AAM) [12]. Using the same architecture as the TAM model, [12] added factors of compatibility and trust. Influenced by IDT, compatibility in the AAM refers to how similarities between the automation and user's past experience can support acceptance. The contribution of trust has not been given much credit in TAM or other IS-originated models, but is considered a key characteristic of acceptance in human-machine interaction research. Importantly, AAM addresses the dynamic nature of acceptance, highlighting the importance of feedback mechanisms.

### C. Usability research and human-machine interaction

In contrast to methods originated in the IS community which provide little guidance to system design, the CE community has provided methods offering a closer link between theory and practice. Acceptance research in the CE community has long focused around the paradigm of user-centred technology. An important derivative therefrom is the concept of usability [13], and the assumption that more usable technology will be more readily accepted. However, as Dillon and Morris [7] note, although acceptance and usability are linked, they are not equivalent. Usability research has not addressed the concept of acceptance and adoption explicitly, rather made the assumption that usability is a prerequisite of acceptance. Further research is however needed to clarify the relationship between usability and acceptance. Usability research tends to focus on technology such as consumer products, computer software, and office applications. Common methods are user participation, focus groups, and rapid iterative interface development. These methods typically do not address the requirements for technology used in more complex and dynamic environments, such as ATM.

To address the challenges of acceptance in human-machine interaction within such high-risk organisations as ATM, the CE community has focused on issues related to the influence of levels of automation [12]. Controller attitudes toward new automation with higher authority have been identified as a key factor in how successfully we will be able to introduce new automation in the future. One critical challenge in designing automation, particularly if that automation is capable of offering strategic advice, is how to ensure that users will trust the system enough to use it, yet remain discriminating and effective [14]. An operator's trust, it seems, must be calibrated to the automation. The issue of automation trust has been studied for years, in fields such as medicine, ATM, combat identification, uninhabited aerial operations and others. Much of this work was built on the notion of human-to-human trust [14]. Trust is related to acceptance of automation, but is generally considered distinct. Both trust and acceptance are also intertwined with automation usage, and trust is seen as one of the key determinants of whether a human operator will choose to use optional automation [4]. As a means to address acceptance and trust issues in automation, researchers have suggested that human-automation interaction should be modelled on human-human interaction. This area of research is briefly reviewed next.

### III. THE HUMAN TEMPLATE

The Computers Are Social Actors (CASA) research provided evidence that humans interact socially with computers (think household computer of the mid 90s) similar to how they interact with other humans [15][16]. Beyond simple computers, [17] showed that more complex machines, characteristic of automation used in high-risk environments, adhering to human social "rules" (defined as *etiquette*) have significant positive effects on both performance and trust. In a review on trust, both human-human and human-machine, [18] argued that "*automation users would benefit if machines were designed to incorporate characteristics of humanness that would, in turn, elicit social responses from the human user.*" (pp. 295). The human response described is more commonly known as *anthropomorphism* (i.e. a conscious process at which a human attributes human characteristics on a non-human item). Kiesler et al. [19] showed that participants anthropomorphised more with a humanoid robot than a robot-like agent. In a social interaction task, the humanoid robot was engaged with more (measured in minutes of interaction), liked better, and considered more lifelike. Waytz et al. [20] demonstrated that the effect of anthropomorphism is important for granting responsibility and developing trust in an agent, the affordance of moral care and concern to the agent, and extent to which the agent influences oneself.

Research has also challenged the premise that human-automation interaction should be modelled after human-human interaction. In a series of studies involving a visual detection task [21][22], it was shown that initial trust in a decision support aid tends to be higher than trust in humans because of preconceived expectations that the automation aid will perform "perfectly." This positive bias toward automation

makes operators more sensitive to errors made by the automation, which can cause large drops in trust and use, leading to disuse. Madhavan and Wiegmann [23] showed that, even when an automated advisory behaves like a human advisor, user perceptions of the two are not equal, and trust and reliance in automated advisors and human advisors varies. In their study, perceived reliability, trust and acceptance of the advisor depended on the perceived expertise of the advisor. They concluded that humans are judged more by subjective traits, whereas automation is judged more by its performance.

The description of human-automation cooperation as teamwork has recently stirred up a debate regarding whether automation can actually function as a teammate. Groom and Nass [24] argue that the "team" label implies expectations and qualities that automation, and specifically robots, are unable to meet, resulting in a negative response from the human when these expectations and qualities are not met. Instead, research should aim for a human-automation organizational structure that recognizes and better utilizes the different unique qualities and abilities of humans and machines. Others disagree, arguing that the capability of automation to take initiative and delegate orders qualifies automation as a team member [25][26]. While we leave this debate open, we acknowledge that indeed, we do not yet know how to create automation as creative, adaptive and innovative as a human. Until that day, there will always be differences between human-automation and human-human collaboration. But what if automation could equal human innovativeness and creativity in decision making? Would controllers accept such automation? This question is what inspired the MUFASA project.

### IV. THE MUFASA ACCEPTANCE MODEL

IDT research has received little, if any, attention in ATM. This is unfortunate but not surprising given how IDT has developed and been applied throughout the years. In contrast to contexts of predominant IDT research, ATM is a high-risk organisation characterised by complexity and a highly dynamic environment. The technologies used are advanced and characterised as highly coupled open systems. Controllers are highly specialised and experts at their work. IDT, however, generally focuses on the novice public and closed-system technologies. Acceptance research as defined in the IS community has predominantly focused on the intention to use, and not actual usage, although notable exceptions exist [10]. However, we know that determinants for intentions to use often do not translate well to determinants of actual use (e.g. [22][23]). Further, approaches such as TAM do not address characteristics of acceptance unique to ATM, and they do not provide a clear link between theory and design. Usability approaches provide a clear link between theory and practice but do not adequately address acceptance. While the CE community has implied use and acceptance by studying unique aspects of ATM, such as complexity, levels of automation, and trust, it has not explicitly studied acceptance, nor provided an explanatory model.

The purpose of our study was to examine how the similarity between decision making processes of the human and automation influences acceptance of an advisory from a decision support system. Although none of the reviewed acceptance approaches fully satisfies the purposes of MUFASA, aspects of each have been influential in the development of an acceptance model suitable for our study. Both the IS and CE communities have devoted considerable effort into researching acceptance, but done so independently. Recognising the potential benefits of combining the two, [12] proposed the AAM. This model however does not address ATM specifically and is more intended for consumer products. Our model of technology acceptance draws on both views. It incorporates the unique aspects of trust, level of automation, and complexity found to influence acceptance in ATM. From IDT and the approaches developed by the IS community the notion of compatibility is relevant. But the compatibility construct does not explicitly consider similarities in decision-making strategies. Influenced by the compatibility construct, and research on similarities between human-machine and human-human interaction, we derived the notion of *strategic conformance*, i.e., the extent to which automation's performance and underlying processes are similar to those of the human. In two important ways, this paper extends previous research that has explored how human characteristics superimposed on automation affect acceptance, trust, and performance. First, strategic conformance is unique in that it explores the *apparent* similarities between automation and human decision making processes underlying behaviour. Previous research has been limited to studying *overt* similarities between automation and human behaviour and communication. Second, strategic conformance addresses specific *individual* behaviour in contrast to *general* applicable human behaviour according to cultural norms.

## V. EXPERIMENTAL DESIGN

The method applied in this paper for studying strategic conformance involves presenting controllers, through means of a decision support system, unrecognisable replays of either their own solution (i.e., a *conformal*) or a colleague's different solution (i.e., *non-conformal*) to a series of pending aircraft separation conflicts. The process of identifying conformal and non-conformal solutions consisted of comparing one controller's solution to a conflict with all the other controllers' solutions to the same conflict. The degree of conformance was then determined using a scale. A conformal advisory was qualified in terms of aircraft choice, clearance type (e.g., heading change only), and clearance direction (e.g., heading change to the left). A non-conformal solution therefore always featured a different aircraft choice and/or clearance type and direction.

Controllers were instructed that all conflict resolution advisories were generated by automation. Assuming controllers remain consistent in their solution choice over time, this allows us to experimentally manipulate strategic conformance between human and (simulated) machine solutions. Note that by replaying human actions we are in fact

simulating automation capable of providing conflict resolutions of the same calibre as a human controller would be able. This replay procedure is inspired by one used many years ago, with which it was shown that operators might be more likely to find fault in automation than in themselves, even when "automation" is an unrecognisable replay of their own performance [27].

We conducted a series of two human-in-the-loop Air Traffic Control (ATC) simulations of increasing automation authority. The first simulation series, considered to be a *prequel* to the main conformance experiment, was designed to capture controllers' manual performance in maintaining safe separations between aircraft. Air Traffic controllers needed to use an advanced separation assistance tool, the Solution Space Diagram (SSD) developed by Delft University of Technology (Figure 1), to vector aircraft and solve conflicts by issuing speed and/or heading clearances to aircraft [28]. In its most succinct form, the SSD is a tactical decision-support tool that visualizes 'go' and 'no-go' areas on a circular heading ring around an aircraft. When vectoring an aircraft into a 'go' area, it will remain free of conflict with all other aircraft and a loss of separation will essentially never occur. In the second simulation series, i.e., the conformance experiment, the controller solutions (and those of their colleagues) to specific conflicts were replayed as automation advisories and plotted within the SSD, after which the controller could either accept or reject the advisory. By plotting the advisory within the SSD controllers could always inspect the validity of the advisory, but also "look around" for better alternatives.

### A. Participants

Sixteen professional air traffic controllers voluntarily participated at the Shannon Area Control Centre, Ireland. Experience ranged from zero to ten years ( $\bar{x} = 2.5$  years). Twelve controllers currently worked en-route and one controller worked the tower position. Three were students currently undergoing en-route training.

### B. Apparatus

The ATC simulator ran on a portable computer connected to an external 21" monitor. Participants interacted with the simulator through an externally connected computer mouse and keyboard. Aircraft could be controlled in the horizontal plane through heading and speed commands, or combinations thereof. The simulator was a Java-based application (using OpenGL extensions) that allowed participants to control short traffic scenarios. No wind conditions were taken into account, all aircraft remained on the same flight level and could not be changed, and the aircraft velocities (and speed clearances) were given in knots Indicated Airspeed (IAS). Further, the aircraft motion was simulated by first order, linear kinematic equations. To keep traffic scenarios sufficiently short and interesting, it was decided to run the simulator four-times faster than real time. Speeding up traffic scenarios in ATC simulators is a common technique to serve this purpose. Finally, the aircraft plots on the display were updated every second to simulate a 1 Hz radar update frequency.

### C. Traffic Scenarios

Each series consisted of 16 traffic scenarios, each based on a squared airspace equal in size (Figure 1). Four baseline scenarios were each rotated in different angles to create three variants, resulting in four scenario groups with four scenario variants in each group. This reduced potential confounding factors, and ensured that initial complexity was the same across scenarios, facilitating comparison between low and high complexity conditions. We aimed to make each traffic scenario repeatable, yet unrecognisable to participants. We maintained sector geometries through scenario rotations in which the relative trajectories and closure angles of aircraft were kept constant, but the entire sector was rotated, and sector entry/exit points renamed.

To guarantee exact replays of controller solutions, each baseline scenario featured only one designed conflict between two aircraft. As a result, any other conflict occurring in a scenario was the consequence of controller intervention. The geometry of the designed conflict was only varied between baseline scenarios. The conflict pair was initially aligned to the exit points and thus required no initial controller interaction. The other aircraft in the sector were considered “noise” aircraft to distract the controller from the conflict pair. Some of the noise aircraft were misaligned with their exit point and displayed in grey, whereas aligned noise aircraft were displayed in green such that the controller could immediately see which aircraft had not yet been cleared to their exit point.

In designing the scenarios it was very important that noise aircraft did not interfere with the designed conflict, such that in the conformal experiment the controller would have the same set of solutions as observed during the prequel experiment. In the prequel experiment each scenario lasted roughly two minutes. In the conformal experiment each scenario lasted less than one minute in order to reduce the likelihood of significantly changing the traffic situation and thus not being able to guarantee automation advisories (i.e. solution replays) that would solve the conflict. Further, we decided to introduce the automation advisory early in each scenario to prevent participants from solving the designed conflict proactively.

### D. Independent Variables

The experimental design followed a within-subjects design with two independent variables, degree of strategic conformance of the resolution advisory (conformal vs. non-conformal) and traffic complexity (low vs. high). A conformal advisory was qualified in terms of aircraft choice, clearance type (e.g., heading change only), and clearance direction (e.g., heading change to the left). A non-conformal solution therefore always featured a different aircraft choice and/or clearance type and direction. Further, non-conformal solutions were derived from solutions provided by other controllers. Complexity was varied through means of aircraft count, and calibrated in a series of test trials. Finally, presentation order of traffic complexity and strategic conformance was balanced

between participants and traffic scenarios using a Latin Square design.

### E. Dependent Variables

Dependent variables focused on acceptance of the advisory, the controller’s agreement with the advisory (measured on a zero to one-hundred scale), response time, and subjective workload ratings (measured on a zero to one-hundred scale). Response time was measured from presentation of the resolution advisory to pressing the ‘accept’ or ‘reject’ button.



Figure 1. The ATC simulator showing: 1) the current performance score, time, and runs to go, 2) the SSD of the currently selected aircraft in magenta, 3) the SSD with an automated resolution advisory in an amber colour (suggesting a speed decrease to resolve a future conflict), and 4) the advisory dialog window that the controllers could use to either accept or reject the advisory.

### F. Experiment Procedure

The whole experiment took four weeks in total. In the first week, the prequel experiment was conducted with the aim of capturing controller resolutions to the designed conflicts. Following briefing and consent procedures, we conducted 16 training runs and 16 measurement runs. Participants were given two main tasks to be performed using the SSD, namely resolving conflicts and clearing aircraft to their intended exit points. A continuously updated performance score reflecting these two task parameters was included to keep participants focused and motivated, and more importantly to prevent scenario recognition and early detection of the designed conflict. To warn the controllers of short-term conflicts, an auditory alert was triggered and the aircraft involved in the conflict were displayed in red.

In the second and third week, prequel data were analysed and a set of eight conformal (i.e., a replay of controller’s own

solution) and eight non-conformal (i.e., a replay of a colleague's different but workable solution) advisories were created for each individual participant. Finally the conformance experiment was conducted in week four. Following a simulator briefing the experiment started with eight training runs, followed by 16 measurement runs. The same scenarios as in the prequel were used but the order varied according to a Latin Square design. Participants performed the same task as in the prequel experiment, but were now assisted by a decision support system that would once in a while provide resolution advisories by pro-actively auto-selecting an aircraft in conflict. The decision support system always preceded the short-term conflict alert system.

The resolution advisory consisted of either a heading vector, speed vector, or combination thereof, plotted inside the SSD of that aircraft. The resolution advisory was accompanied by a beeping sound and the appearance of a dialog window that the controller had to use to either 'accept' or 'reject' the advisory (see Figure 1). Upon pressing the accept button, the advisory would be automatically implemented. Upon pressing the reject button, the advisory was discarded and the controller had to implement his own workable solution using the SSD. A 15 second time window was provided to inspect, accept, or reject the advisory. After 15 seconds, the advisory expired and participants had to implement their own workable solution. Participants were told that an advisory would always solve the conflict, but it would not always suggest the most optimal solution. As such, participants were encouraged to find better alternatives at their own discretion. After each scenario, participants were given performance feedback in terms of an average performance score. Second, participants were asked to give ratings on their experienced workload and their agreement with the automation advisory. After the experiment, participants were asked to complete a questionnaire containing information of demographic value and statements (in five-point Likert scale format) querying their opinions of the simulator, the SSD interface, and the automated advisories.

## VI. RESULTS

Experimenter observations and an analysis of questionnaire data indicated that participants did not recognise scenarios, nor, more importantly, that the automated advisories were, in half the cases, replays of their own prior conflict resolutions. Questionnaire data showed that controllers enjoyed the SSD tool and simulator, but did not find scenarios very realistic.

### A. Advisory Acceptance and Agreement with Resolution Advisory

Cumulative accept/reject scores for strategic conformance and complexity can be seen in Figure 2(a). Participants accepted more resolution advisories in the high complexity scenarios. Analysing strategic conformance, it can be seen that participants accepted more conformal solutions than non-conformal solutions. A 2 x 2 repeated measures ANOVA showed a significant main effect for both complexity ( $F(1,15) = 11.139, p < 0.01$ ) and strategic conformance ( $F(1,15) =$

10.624,  $p < 0.01$ ) on acceptance. The interaction between complexity and strategic conformance was not significant.




Figure 2. Proportion of advisory acceptances (a) and error bar plots of normalized agreement ratings (b), by complexity and strategic conformance.

Participant ratings of agreement with the resolution advisory revealed a difference between both the complexity levels and strategic conformance levels (Figure 2(b)). The 2 x 2 repeated measures ANOVA revealed that both complexity ( $F(1,15) = 7.735, p < 0.05$ ) and strategic conformance ( $F(1,15) = 18.095, p < 0.01$ ) had a significant main effect on agreement ratings. Agreement with resolution advisories varied positively with increasing complexity and strategic conformance.

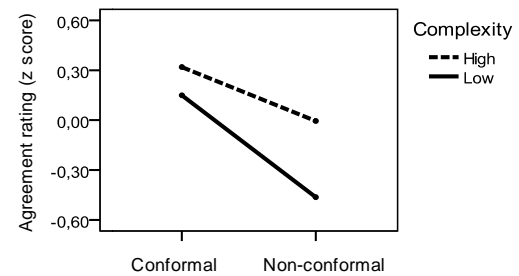


Figure 3. Agreement with resolution advisory ratings (normalised), by complexity and strategic conformance.

Although a trend was observed in the interaction between complexity and strategic conformance (Figure 3), this was not significant ( $F(1,15) = 3.186, p > 0.05$ ). This trend suggests that non-conformal solutions tended to be less agreed with than conformal solutions and that this effect was more pronounced under low complexity, perhaps because controllers were under less time pressure and would have had more time to evaluate candidate solutions.

### B. Workload and Response Time

The 2 x 2 repeated measures ANOVA showed that workload significantly increased with more complex scenarios ( $F(1,15) = 179.950, p < 0.01$ ). Neither strategic conformance ( $F(1,15) = 0.397, p > 0.05$ ), nor the interaction between the two factors was significant ( $F(1,15) = 0.266, p > 0.05$ ). For response time a main effect was measured by performing a 2 x 2 repeated measures ANOVA, with conformal scenarios having a significant faster response time ( $F(1,15) = 9.557, p < 0.01$ ). Although showing a trend, scenario complexity did not have an effect on response time ( $F(1,15) = 4.182, p > 0.05$ ). The interaction between complexity and strategic conformance was also not significant.

## VII. DISCUSSION

The main objective of this study was to investigate the effect of strategic conformance and complexity on controller performance and acceptance in the context of higher levels of automation. A significant main effect of strategic conformance was observed on acceptance, agreement with resolution advisory, and response time. Conformal resolution advisories were accepted more often, received higher agreement ratings, and were acted upon faster than were non-conformal advisories. Results could, however, not confirm the hypothesis that conformal advisories reduced participant workload. These findings suggest that controllers not only discriminate between resolution advisories, but more importantly, prefer resolution advisories that match their own way of working in comparison to advisories generated by their colleagues. The increased response time measured for non-conformal resolution advisories may indicate that controllers questioned advisories before making the decision to either accept or reject, and that it took longer to conclude that the advisory did not fit with the controllers' own solution. Possibly, the longer response time can also be attributed to the time taken for the controller to mentally construct their own solution. These results support further research initiatives into heuristic-based automation.

Similar to strategic conformance, complexity significantly affected acceptance and agreement with resolution advisory. In high complexity scenarios, acceptance and agreement with resolution advisory was higher than in low complexity scenarios. The lower acceptance and agreement observed in low complexity scenarios could be explained in that controllers had more time to question the advisory and consider more optimal options. With increasing workload there are less resources and time available to consider options before making a decision. But these results are also interesting in that they show that controllers disagree with themselves when complexity is lower. Overall, participants disagreed with themselves in close to one quarter of the time. Why would controllers disagree with themselves to that extent? This might be evidence of "dispositional" automation bias, a reluctance to use automation under any circumstance. In this study, controllers were told that all resolution advisories were generated by the automation. How would controllers have reacted if the presumed source was varied between themselves, colleagues, and automation? It is also possible that controllers

are inconsistent over time. Traditionally and currently, both machine design and research thereof have assumed that controllers are homogenous as a group. While neglecting individual differences perhaps has been sufficient for achieving the ATM system we have today, the envisioned sophisticated automation needed for successful human-machine teamwork in ATC may need to better acknowledge individual differences.

It is interesting to see that response time varied with strategic conformance but not complexity, and that workload varied with complexity but not strategic conformance. It could be argued that workload and response time would somehow be connected. However, it is important to underline that response time is applicable to the designed conflict and resolution advisory only, whereas workload considers the entire scenario. The specific workload pertaining to the designed conflict and resolution advisory was not pursued.

The unique experimental design required careful consideration in designing dynamic traffic scenarios allowing exploration of strategic conformance. For reasons of experimental control we sought to limit the number of scenarios and through variations of scenario rotation angles increase the measurement span. This experiment showed that it is possible to create "fake" resolution advisories based on participant's own solutions, without participants later recognising it as their own. It opens up new possibilities in researching not only strategic conformance, but also other aspects such as within-participant consistency and reliability over time.

Previous research exploring the link between automation design and acceptance, performance and use, has generally assumed user homogeneity and applied general human social traits to the automation used in experiments [9][17]. In contrast, the concept of strategic conformance is subjective and explicitly linked to an individual's own behaviour. If following the notion that the more socially "correct" automation behaves and communicates, the more positive the influence will be on human-automation interaction, then automation capable of strategic conformance should benefit human-automation interaction even more than automation merely following a generic social behaviour. But, on the downside it could potentially mask poor automation. Given that human-automation teamwork should benefit the ATM system, the contribution of the "best" decisions will likely fluctuate between humans and machines. Similarly, [17] noted that there may be risks associated with manipulating acceptance without addressing the abilities and reliability of the automation. The resulting action may not be the best decision, but chosen simply because the automation is "tuned" for acceptance. Future research is needed to address the implications of this potential manipulation of acceptance, trust, and use of automation.

Acceptance can be studied over a longer time perspective, focusing on how adoption of the technology can be influenced by viewing it from a macro-perspective where acceptance is a process. Conversely, acceptance can be viewed from a micro-

perspective, analysing how to influence acceptance by looking into specific instances of decision making. Time is an important aspect of acceptance and trust. Therefore, time should be incorporated in models of acceptance. Further, measures of longitudinal changes on acceptance should ideally measure acceptance at various cross segments. This is a weakness in this study, that time, and the longitudinal effects were not measured. However, we did not aim to measure adoption as a long-term construct, but the derivative of acceptance - the acceptance decision taken in a specific situation. We assume that if we can increase the acceptance of automation in specific instances, the long-term acceptance (i.e. stable disposition to accept) increases. Future research should investigate the long-term effects of conformance on acceptance, preferably through a longitudinal study in which acceptance of a specific technology is measured before (i.e. intent to use), during, and after adoption. A combined approach of the one used in this study together with IDT model can be used.

### VIII. CONCLUSION

These results suggest that strategic conformance is a potentially important construct underlying automation acceptance, and that it can benefit both acceptance of automation, agreement with automation and response time. A trend toward a strategic conformance and complexity interaction also suggests that the effect on controller agreement with automation advisories is more pronounced under low complexity situations. Though this experiment was considered an important first step, there are several remaining questions to be answered, with respect to both the definition of inherent automation bias, and the nature of such bias as it relates to presumed source of strategic advice.

### ACKNOWLEDGEMENT

The authors acknowledge the inspiration from EUROCONTROL and the SESAR Joint Undertaking. We would like to thank the staff at the Shannon Area Control Centre, Ireland, for all their amazing support in hosting, facilitating, and participating in the experiment.

### REFERENCES

- [1] HALA! "HALA! position paper", Edition 2.04, 2012.
- [2] G. Klein, D.D. Woods, J.M. Bradshaw, R.R. Hoffman, and P.J. Feltovich, "Ten challenges for making automation a "team player" in joint human-agent activity," *Intelligent Systems*, IEEE, 2004, vol. 19, no. 6, pp. 91–95.
- [3] D.B. Kaber, and M. Endsley, "The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task," *Theoretical Issues in Ergonomics Science*, 2004, vol. 5, no. 2, pp. 113-153.
- [4] R. Parasuraman, and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human Factors*, 1997, vol. 39, no. 2, pp. 230-253.
- [5] J.K. Kuchar, and L.C. Yang, "A review of conflict detection and resolution modeling methods," *IEEE Transactions on Intelligent Transportation Systems*, 2000, vol. 1, no. 4, pp. 179-189.
- [6] M. Bekier, B.R.C. Molesworth, and A.M. Williamson, "Tipping point: The narrow path between automation acceptance and rejection in air traffic management," *Safety Science*, 2012, vol. 50, no. 2, pp. 259-265.
- [7] A. Dillon, and M. Morris, "User acceptance of new information technology: theories and models," In M. Williams, Ed. *Annual Review of Information Science and Technology*, Medford NJ: Information Today, vol. 31, 1996, pp. 3-32.
- [8] E.M. Rogers, *Diffusion of Innovations*, 3rd ed., New York: Free Press, 1983.
- [9] E. Karahanna, R. Agarwal, and C.M. Angst, "Reconceptualizing compatibility beliefs," *MIS Quarterly*, 2006, vol. 30, no. 4, pp. 781-804.
- [10] V. Venkatesh, M. Morris, G. Davis, and F. Davis, "User acceptance of information technology: Toward a unified view," *MIS Quarterly*, 2003, vol. 27, no. 3, pp. 425–478.
- [11] F.D. Davis, R.P. Bagozzi, and P.R. Warshaw, "User acceptance of computer technology: a comparison of two theoretical models," *Management Science*, 1989, vol. 35, no. 8, pp. 982- 1003.
- [12] M. Ghazizadeh, J.D. Lee, and L.N. Boyle, "Extending the Technology Acceptance Model to assess automation," *Cognition, Technology & Work*, 2011, vol. 14, no. 1, pp. 39–49.
- [13] J. Nielsen, *Usability Engineering*, Boston: AP Professional, 1993.
- [14] B.M. Muir, "Trust between humans and machines, and the design of decision aids," *International Journal of Man-Machine Studies*, 1987, vol. 27, no. 5-6, pp. 527-539.
- [15] C. Nass, B.J. Fogg, and Y. Moon, "Can computers be teammates?" *International Journal of Human-Computer Studies*, 1996, vol. 45, no. 6, pp. 669-678.
- [16] C. Nass, and Y. Moon, "Machines and mindlessness: Social responses to computers," *Journal of Social Issues*, 2000, vol. 56, no. 1, pp. 81-103.
- [17] R. Parasuraman, and C.A. Miller, "Trust and etiquette in high-criticality automated systems," *Communications of the ACM*, 2004, vol. 47, no. 4, pp. 51-55.
- [18] P. Madhavan, and D.A. Wiegmann, "Similarities and differences between human-human and human-automation trust: An integrative review," *Theoretical Issues in Ergonomics Science*, 2007, vol. 8, no. 4, pp. 277–301.
- [19] S. Kiesler, A. Powers, S.R. Fussell, and C. Torrey, "Anthropomorphic interactions with a robot and robot-like agent," *Social Cognition*, 2008, vol. 26, no. 2, pp. 169-181.
- [20] A. Waytz, J. Cacioppo, and N. Epley, "Who sees human?: The stability and importance of individual differences in anthropomorphism. *Perspectives on Psychological Science*, 2010, vol. 5, no. 3, pp. 219–232.
- [21] M.T. Dzindolet, L.G. Pierce, H.P. Beck, L.A. Dawe, and B.W. Anderson, "The perceived utility of human and automated aids in a visual detection task," *Human Factors*, 2002, vol. 44, no.1, pp. 79-94
- [22] M.T. Dzindolet, S.A. Peterson, R.A. Pomranky, L.G. Pierce, H.P. Beck, "The role of trust in automation reliance," *International Journal of Human-Computer Studies*, 2003, vol. 58, no. 6, pp. 697-718.
- [23] P. Madhavan, and D.A. Wiegmann, "Effects of information source, pedigree, and reliability on operator interaction with decision support systems," *Human Factors*, 2007, vol. 49, no. 5, pp. 773-785.
- [24] V. Groom, and C. Nass, "Can robots be teammates?: Benchmarks and predictors of failure in human-robot teams," *Interaction Studies*, 2007, vol. 8, no. 3, pp. 483-500.
- [25] J. Langan-Fox, J.M. Canty, and M.J. Sankey, "Human-automation teams and adaptable control for future air traffic management," *International Journal of Industrial Ergonomics*, 2009, vol. 39, no. 5, pp. 894-903.
- [26] N. Wijngaards, M. Kempen, A. Smit, and K. Nieuwenhuis, "Towards sustained team effectiveness," In: *Coordination, Organizations, Institutions, and Norms in Multi-agent Systems*, vol. 3913, Springer, Berlin/Heidelberg, 2006, pp. 35–47.
- [27] R.B. Fuld, Y. Liu, and C.D. Wickens, "The impact of automation on error detection: Some results from a visual discrimination task," *Human Factors and Ergonomics Society Annual Meeting Proceedings*, 1987, vol. 31, no. 2, pp. 156-160.
- [28] G.A. Mercado Velasco, M. Mulder, and M.M. van Paassen, "Analysis of air traffic controller workload reduction based on the solution space for the merging task," *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Toronto, Canada, 2010.