Available online at www.sciencedirect.com

ScienceDirect

IFAC-PapersOnLine 49-19 (2016) 207–211

A Hybrid Approach to Training Expert Skills in Highly Automated Systems: Lessons from Air Traffic Management

Brian Hilburn

Center for Human Performance Research, CHPR BV, Voorburg, the Netherlands (brian@chpr.nl)

Abstract: In air traffic management (ATM), as in many other domains, automation is increasingly capable of performing more strategic and "cognitive" aspects of system performance. This paper sets out a potential hybrid approach to automation design, which assumes qualitatively different challenges at the introductory and mature phases of automation implementation. Whereas operator acceptance seems the critical issue at the time of automation introduction, skills development and maintenance seem most significant as expertise accrues. This proposed hybrid marries the notions of strategic conformance and adaptive automation to achieve a design approach in which, over the span of the skill acquisition cycle, automation is fitted to the novice, and the expert is fitted to the automation. Further, this approach assumes that training can (and indeed must) ultimately extend the training criterion from that of (heuristically based) expert operator performance, to that of (algorithmically based) optimized system performance.

Keywords: Training, adaptive automation, human centered automation, air traffic management.

1. INTRODUCTION

Over the next 20 years, global demand for air travel is predicted to increase roughly 5% annually, and the in-service jet fleet to grow by nearly 90% (JADC, 2015). Meanwhile, commercial and environmental concerns will mean increasingly complex routing, to minimize fuel burn and delays. Together, these factors are driving the need for improved planning and coordination functions, which in turn will likely require increasingly sophisticated automation.

In air traffic management (ATM), as in many other domains, automation has historically been used to control more mundane "housekeeping" tasks, with higher level tasks left to the human operator. This view was captured many decades ago in Paul Fitt's *Machines-are-better-at, Men-are-better-at* (MABA MABA) approach to function allocation, which assigned individual tasks to the more capable agent, either human or machine (Fitts, 1951). For instance, whereas humans are more adept at perception, judgment, improvisation and induction, machines excel at speed, power, computation and short term memory tasks. The MABA MABA view has fallen into disfavour over the years, as technological capabilities have evolved, and the line between human and machine capabilities has blurred (Bye et al, 1999; Hoffman et al., 2002). Automation is increasingly capable of assuming greater authority and autonomy, and performing more strategic and "cognitive" aspects of system performance.

One overarching challenge currently facing various work domains is how to design advanced automation in such a way that it is both used, and used in a beneficial manner (Parasuraman & Riley, 1997), to balance the potentially competing demands of greater automation and operator skill retention. This paper sets out, based on theoretical and empirical evidence, a potential hybrid approach to automation design that supports the training and maintenance of expert skills. As laid out in the following sections, this paper assumes that future automation design faces qualitatively different challenges at the introductory and mature phases of implementation.

2. HUMAN- VS TECHNOLOGY-CENTERED APPROACHES TO AUTOMATION DESIGN

The technology-centered approach (TCA) to automation design views the human as a source of potential error, and starts from the position that tasks should therefore be automated whenever possible. In a sense, this view dates backs to a paradox recognized by Fitts in the 1950s: If we understand how a human performs a task, we can construct a mathematical model of that task that should allow us to create a device, program or computer to perform the task at least as well as the human. To the extent that the human can be compared to a machine, he can be replaced, and designed out of the system. TCA argues that keeping the human in the loop is, by definition, impossible if system performance is extended through fundamentally new tasks, or tasks that cannot be overseen or performed by the human. If the system is to perform tasks that the human is incapable of performing (or performing say at the same rate or accuracy) then the underlying process should be a black box to the operator. According to this view, it is sufficient that input / output

2405-8963 © 2016, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.
Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2016.10.522

relationships are clear to the operator, and design need not consider transparency of the intervening processing. After all, the argument goes, "dumbing down" automation to the level of the human will limit system performance, and risks recreating human error modes.

The notion of human-centered automation (HCA) traces its roots to the work of Billings (1997). Based on empirical evidence from the flight deck and ATM, his seminal (and lengthy) treatise argued that operators in an automated system must be keep informed, active and in-the-loop. The risks of an out-of-the-loop operator include: Situation awareness problems, reversion-to-manual difficulties under off-nominal conditions, and mis-calibrated trust (either insufficient or excessive).

The aim of this paper is not to wade into the debate between the technology- and human-centered automation schools, but to note that these two viewpoints draw a clear contrast between the problem solving styles typically employed by human and automation—the former tend to rely on heuristic "rules of thumb," the latter on optimized algorithms. In the case of ATM, a mathematically-optimized solution (e.g., that minimizes total flightpath distance across a traffic pattern) might not fit with that of the human (Nantanen & Nunes, 2005; Prevot et al., 2012). Even if the human could derive the optimized solution, it might be too mentally demanding to implement and monitor. This is evident in controllers' "set-and-forget" strategy of turning aircraft on parallel headings, to ensure separation, or tending to turn slower aircraft behind faster ones (Kirwan & Flynn, 2002). Although these strategies might sometimes be sub-optimal from a mathematical standpoint, they ease cognitive burden and safeguard against failures to detect a future loss of separation.

This difference in human vs machine problem solving styles, and underlying mechanisms, is potentially important as we consider how to design advanced automation systems so as to best develop and maintain expert skills. As the following section discusses, this issue might be especially critical at early stages, when such automation is first introduced.

2. INTRODUCING ADVANCED AUTOMATION: THE ACCEPTANCE PROBLEM

Acceptance has been identified as one of the greatest obstacles to introducing new ATM automation (JPDO, 2011; Hilburn, 2003). Trends in ATM suggest that automation will likely become more strategic in both timescale and control authority, less transparent to the controller, and act via resolution advisories. Picture a system that advises the controller, say 20 minutes in advance, to resolve medium-term conflicts which the controller might have difficulty evaluating. In a real sense, automation would become an agent in the ATM process, much like a human colleague—and as with a human colleague, its advice can be ignored. This is especially likely to happen if its benefits are not perceived. Herein lies a paradox: a controller might only rely on such an advisory system if its benefits are obvious, yet those benefits will not be obvious until it is used. EUROCONTROL's CORA project set out to build a prototype advanced advisory system for strategic de-confliction (Kirwan & Flynn, 2002). The project recognized that controller initial acceptance was critical to its introduction (Hilburn, 2000), and tried to ensure that the system would solve problems like a human would. Although the project offered some promising results, it (like similar efforts before) was hindered in one important regard: one cannot guarantee similarity between human and machine solutions in a built system. More recently, Westin and colleagues (Westin et al, 2013) explored the impact of controller acceptance in a more fundamental way, by asking: If automation were to perform in a way that perfectly matches that of the controller, would controllers accept such automation? The concept of "strategic conformance" was defined as

…the degree to which automation's behavior and apparent underlying operations match those of the human.

Westin et al assessed controller acceptance of "automated" air traffic scenarios that were in fact unrecognizable replays of either a given controller's own previous performance (by definition, this was "conformal" with his / her own strategy), or that of a colleague (who had chosen a slightly different solution strategy). They found that acceptance of automated advisories was significantly higher for conformal solutions, 76% vs 57% ($F(1,15)=10.6$, $p<.01$). Similarly, agreement was significantly higher, and response time significantly lower, for conformal vs non-conformal advisories.

If, as these results suggest, strategic conformance can help foster acceptance of new automated advisory systems, the benefit of this would appear to lie in the initial deployment phase, when controllers are first introduced to new advisory automation. Ultimately, though, if we are to realize benefits of such advanced automation, it is not enough that it simply matches the controller's way of working. It must extend the controller's capabilities, as discussed above in section 1. But how do we design automation in such a way that it helps develop and maintain expert skills? One promising approach, as outlined in the following section, draws on the rich body of evidence on adaptive automation and intelligent tutoring system concepts.

3. OPTIMIZING ADVANCED AUTOMATION: THE EXPERTISE PROBLEM

The potential challenges that experts would face under this type of envisioned automation would be slightly different from the preceding. Automation would have to work hand-in-hand with the controller, but priority would now shift from matching the controller, to handling complex traffic flows. This would require high performance automation assuming control of some higher level functions. Much has been written over the years about the potential costs of static automation, in which a system operates at a fixed level of automation (LOA), and task allocation remains fixed between human and machine. Potential human performance costs include problems relating to monitoring and supervisory