2013

Human-Centered Automation as Effective Work Design

Amy R. Pritchett

Karen M. Feigh

HUMAN-CENTERED AUTOMATION AS EFFECTIVE WORK DESIGN

Amy R. Pritchett and Karen M. Feigh
Georgia Tech
Atlanta GA, USA.

This paper describes how the challenge of human-centered automation can be recast as the challenge of, first, designing the work performed by a team of agents and then, second, allocating this work amongst all the agents, human and automated, in support of their own needs and capabilities and to foster team goals. The paper starts by formally describing the construct of *work* as a structure which can be formally analyzed and around which other design decisions can be made. It then reviews the requirements of effective function allocation within a team to enable their collective taskwork, and to provide the appropriate teamwork. An example is given that highlights key tradeoffs in designing and allocating work in teams of human and automated agents: no one design can maximize all the desired attributes of human-centered automation.

Work is defined by Webster as "effort directed to some purpose or end." Thus, it is purposeful activity directed at goals established by a concept of operation. Here, we view work as a construct applied at the team level. Further, the notion of work is an ecological perspective: work is achieved by acting on a dynamic environment in response to its demands. This environment can be defined as the aggregation of physical and social/cultural/policy constructs required to describe, constrain, regulate and structure the dynamics of the work; thus, the environment may have inherent dynamics which agent actions need to mirror, may provide affordances which need to be sensed and capitalized upon, and may constrain behavior.

Thus, what overall taskwork needs to happen, and its overall structure and dynamics, is driven by the team goals and by the environment. The team's work emerges out of the collective behavior of all agents in the team, human and automated, even when some of the agents may not see how their activities contribute. The allocation of functions within the team creates the need for additional work: teamwork. This teamwork also requires its own constructs and resources, such that each individual's perception of the environment includes both part of the overall environment and the teamwork aspects created by his/her team members.

From this viewpoint, two things may be designed: the concept of operation defining the goals and structure of the overall task work, and specification of teamwork, including the allocation of functions. The concept of operations is specified at the team level, and establishes core goals for the work; it is constrained by key structures in the environment which the work needs to mirror. The specification of teamwork then brings in the notion of the agents, seeking to allocate functions and identify the constructs within the team that can establish effective human-automation interaction as seen from the humans' perspective.

This paper first summarizes the construct of work as a structure which can be formally analyzed and around which other design decisions can be made. It then reviews the requirements of effective function allocation within a team to enable their collective taskwork, and to provide the appropriate teamwork. An example is given that highlights key tradeoffs in designing and allocating work in teams of human and automated agents: no one design can maximize all the desired attributes of human-centered automation.

Modeling Work

Despite a common tendency to focus on technology design, designing the work can be the more important concern in establishing effective human-automation interaction. Indeed, work design is a harder task than technology design, as specifications of work, such as concepts of operation, must intrinsically integrate the economic and safety metrics by which the total system will be evaluated, the potential contributions of (or constraints on) technology and human performance, and the regulatory, policy and procedural considerations in allowing access to - and defining interaction within - the collaborative functioning of the system.

Thus, the foundation of human-centered automation is laid in the work design. At such an early stage, human in the loop evaluation is not possible – the training, procedures and technology are only specified in terms of the functions required. Instead, the important construct is: *If everything, and everyone, in the system performs their functions perfectly, what will emerge?* The answer is created by the interplay of the work environment (as defined by physics and regulations) and the team acting upon the environment. Concepts of operation can be constructed poorly when they are sensitive to small variations in how the work is performed, or where they assume actions will be performed with a speed or detail that is not possible or for which information is not available. For example, air traffic concepts of operation applying optimized profile descents must create the work activities that regulate the physics of an aircraft descending in a fuel-efficient manner, while recognizing that key variables – the aircraft performance, the wind profile through the descent, constraints on the aircraft to fit within the traffic stream – are known (or only partly known) at different locations and at different times, yet key decisions to descend earlier or later can have profound impacts on the aircraft's ability to follow its individually-optimized descent and fit within the broader traffic flow.

To ensure the overall specification of work is sound, work activities can be first modeled in detail without requiring detailed models of the agents who may perform them. Conceptually, this analysis is best conducted with simple models of human performance such that any problems can be clearly linked to the feasibility of a concept of operation. Further, once the broader dynamics of the work are established, a concept of operation can be examined for its robustness and resilience: *What if something doesn't go perfectly?* Here, the system's response to unexpected events can be modeled and simulated. These unexpected events may stem from several sources: exogenous inputs to the system (for example, on an air traffic system an unexpected tailwind or thunderstorm); technology (for example, the failure of a radar system); or from human performance (for example, limit on the number of simultaneous activities that can be performed). The work involved in responding to these events will be emergent and dynamic, and a concept of operation can be designed to be more (or less) robust and resilient).

Work can be analyzed in several ways. Approaches such as Contextual Design (Beyer & Holtzblatt, 1988) and Cognitive Work Analysis (Bisantz and Roth, 2007) provide qualitative and visual presentations of the work that are intended to guide and inform designers. Our own recent efforts have established a computational framework that enables work to be computationally modeled and simulated early in design (Pritchett, 2013), first to support analysis of the concept of operation (i.e. the required taskwork of the entire team) and then to examine the design of the team itself (i.e. allocation of functions within the team, and their teamwork).

Requirements for Effective Function Allocation

Function allocation distributes work between agents, human and automated, within a team. The following requirements for effective function allocation can be noted from first principles and the literature. The following summaries are a review from more extensive discussions by Feigh, Pritchett and Kim (in review).

Requirement 1: Each agent must be allocated functions that it is capable of performing.

Every agent in the team must be capable of each of the functions assigned to him/her/it, viewing each function in isolation. In a very coarse sense, such a strategy is supported by assessments of what "Men Are Better At" to what "Machines Are Better At." From this perspective, automation can serve to provide functions that a human cannot perform at all or with sufficient reliability. However, the automation must not be brittle such that, when placed outside its boundary conditions, such automation appears to its operator to fail. *Thus, a prediction of whether the automation will be placed outside its boundary conditions is itself a valuable metric that implies potential concerns with the resilient performance of the team.*

A further consideration in creating effective human-automation interaction examines responsibility and authority. Except when automation is proven to provide safety in all foreseeable operating conditions, humans remain vested with the responsibility for the outcome of automation's actions, a situation termed the "responsibility-authority double-bind" (Woods, 1985). If the human cannot knowledgably oversee the automation, they are forced to 'trust' the automation. However, without a concrete basis for assessing if the automation is correct, humans often over- and under-trust the automation (Parasuraman & Riley, 1997); either way, incorrect trust is viewed as human error, despite its basis in the function allocation. *Thus, identification of mismatches between responsibility and authority is itself a valuable metric that implies potential concerns with trust and reliance, and that requires monitoring by the human.*

Requirement 2: Each agent must be capable of performing its collective set of functions.

The metric for success for this requirement is whether each agent can perform his/her/its collective set of functions under realistic operating conditions. *Thus, prediction of the taskload placed on the human operators – or, where possible, workload experienced by the human operators – is a valuable metric of function allocations.* To fully address known issues with taskload corresponding to human-automation function allocation, such assessments must consider the full range of activities required, including underlying cognitive activities around information gathering and judgment, and requirements to monitor automation, in addition to explicit manual activities. Further, metrics of workload should consider not only aggregate or average workload, but also workload spikes and periods of complacency.

Further, human-centered automation requires that the function allocation establish coherent roles for agents. One attribute of a coherent function allocation can be viewed from the bottom up – within each agent its functions share (and build upon) obvious, common constructs underlying all their activities, such as a shared information and knowledge basis, and the allocation prevents conflicts between the actions of different agents. Another attribute can be

viewed from the top down - the functions collectively contribute towards work goals in a manner that is not only apparent to the human, but that can be purposefully coordinated and adapted in response to context. Thus, the coherence of the functions allocated to each human is itself intrinsically an important construct warranting its own analysis.

Requirement 3: The function allocation must be realizable with reasonable teamwork.

Each different function allocation of the same taskwork demands its own unique set of teamwork functions, including functions for human-automation interaction and for human-human coordination. The impact of this teamwork must then be considered from the perspective of the previous two concepts - can each agent perform each of his/her/its teamwork activities in isolation, and can each agent perform its assigned set of both task work and teamwork functions?

Members of good teams are able to anticipate each other's information needs and provide information at useful, non-interruptive times. However, too often automation is 'clumsy:' it unduly interrupts its human team members because, whereas humans can implicitly sense information about whether other team members would benefit from an interruption, automation historically cannot. *Thus, the potential for a function allocation to cause agents to interrupt each other is an important construct to be analyzed.* In some cases, such as poorly-timed output from automation, such interruptions may be unwarranted; in other cases, different function allocations may require agents to interrupt each other more or less depending on how their functions are allocated and, perhaps, inter-leaved.

Requirement 4: The function allocation must support the dynamics of the work.

Analysis of a function allocation should identify situations where, for example, the interleaving of functions assigned to disparate agents requires significant co-ordination or idling as one waits on another, or where workload may accumulate, or where one agent will be unduly interrupting another, or where executing prescribed procedures may conflict with other work demands, or where automation may be placed outside its boundary conditions. These issues were discussed in the preceding sections, but are repeated here to note their dynamic nature.

Further, resilience is fostered when a human agent may select strategies (courses of action) appropriate to the state of the environment and their own capabilities. The ability of each human in the team to adapt to immediate context has been found to reflect a good balance between the demands on the human and the resources available to them in terms of information, knowledge and time available (Feigh & Pritchett, 2006). However, such adaptation can be constrained or eliminated by an overly prescribed (or proscribed) function allocation, particularly where human-automation interaction dictates a specific sequence of activities from the human. The adverse effects of such overly prescribed function allocations have been found to manifest in work-arounds or dis-use of automation (Feigh & Pritchett, 2010; Parasuraman & Riley, 1997). *Thus, the ability to which a function allocation can accommodate a reasonable variety of human adaptations to context should also be analyzed and fostered.*

Likewise, human-centered automation should foster the humans' ability to maintain a stable work environment. A function allocation may aggravate inherent environmental

unpredictability by, for example, limiting human agents' ability to view important aspects of the environment or by distributing functions in a way such that one agent will trigger the requirement for another to act. In addition, a trade-off exists when designing function allocations between maintaining predictability vs. dynamically allocating functions (Miller & Parasuraman, 2007). *Thus, humans' ability to predict their activities has intrinsic value and should be fostered.*

Requirement 5: The function allocation should be the result of deliberate design decisions.

Changes in operational concepts may be incremental and constrained by current-day technologies, procedures, personnel and/or policies; in other cases, changes in concepts of operation may represent significant innovations in which constructs such as common work practices and relationships between tasks and tools must be significantly altered. Either way, designers need to simultaneously consider the economic and safety metrics by which the total system will be evaluated, the potential contributions of (or constraints on) technology and human performance, and regulatory, policy and procedural considerations. *Thus, the design of human-centered automation should consider not only each agent's experience, but also simultaneously consider the cost and performance of the combined efforts of the human-automated team.*

Conclusion: Perfect Human-Centered Automation is Impossible

In an earlier study we examined four function allocations using computational simulations of work, ranging from full autoflight with datalink (FA1) through progressively 'less automated' conditions to pilot control of the trajectory by setting immediate autopilot targets (FA4) (see Feigh, Pritchett and Kim, in review). In these simulations we also assumed that the human agent (the pilot in this case) might exhibit three different behaviors, as represented by the Opportunistic, Tactical and Strategic cognitive control modes (CCM).

Figure 1 reflects a subset of the metrics collected to examine the ability of the concept of operation and function allocation to meet the requirements noted above and to meet the mission goals as measured by metrics such as time to land. In this figure, the metrics are normalized such that 100% represents the ideal: perfect human-centered automation would have 100% on each of these metrics. Instead, each function allocation scores higher on some metrics and lower on others. The more automated function allocations required better (less) interaction with the pilot but were less predictable to the pilot, made for a lower coherency role for the pilot and interrupted the pilot more. The less automated function allocations provided a more coherent role for the pilot and more predictability, at the expense of requiring them to do more of the work. Further, all of the function allocations assumed the pilot would perform monitoring activities that we predict the pilot would shed in the opportunistic and tactical CCM.

In the end, all of the function allocations met the mission goals in this case. This reflects a situation common in aviation – the agents can adapt and respond to the environment to get things done. The challenge in designing human-centered automation is identifying how to design the work – the concept of operation and the function allocation within it – that strikes the right balance between key trade-offs inherent to divvying up the work to reduce workload, yet maintain coherency, predictability and reduce interruptions.
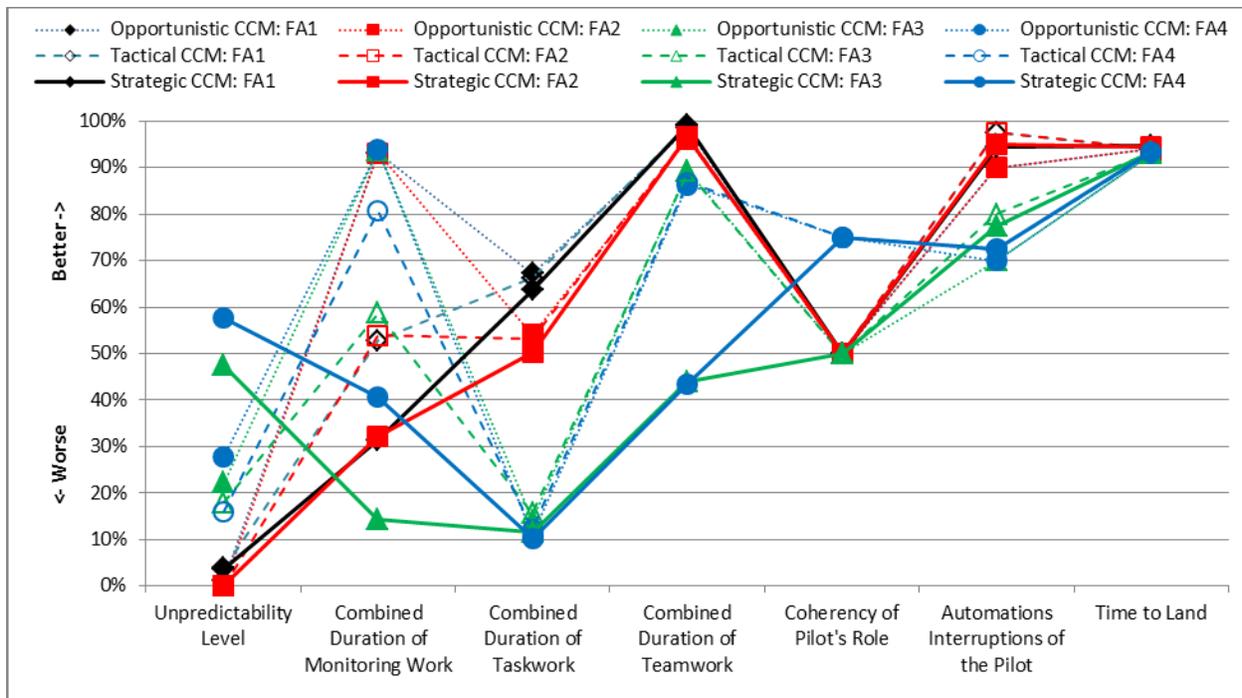
Figure 1. Key metrics of four function allocations between pilot and autoflight system, from the most automated FA1 to the least automated FA4, in conditions where the pilot behavior follows opportunistic, tactical and strategic cognitive control modes (CCM).

References

Beyer, H., & Holtzblatt, K. (1998). *Contextual Design*. San Diego, CA: Academic Press.

Bisantz, A.N., & Roth, E.M. (2007). Analysis of cognitive work. In D.A. Boehm-Davis, (Ed.), *Reviews of human factors and ergonomics, Vol. 3* (pp. 3–35). Santa Monica, CA: Human Factors and Ergonomics Society.

Feigh, K.M., & Pritchett, A. (2006, October). Contextual control modes during an airline rescheduling task. In *Annual Meeting of the Human Factors and Ergonomics Society*, San Francisco, CA. Santa Monica, CA: Human Factors and Ergonomics Society.

Feigh, K.M., & Pritchett, A.R. (2010). Modeling work for cognitive work support system design in operational control centers. *Journal of Cognitive Engineering and Decision Making*, *4*, 1–26.

Feigh, K.M., Pritchett, A.R. & Kim, S.Y. (in review) Modeling human-automation function allocation. Submitted to *Journal of Cognitive Engineering and Decision Making*.

Miller, C.A., & Parasuraman, R. (2007). Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Human Factors*, *49*, 57–75.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, *39*, 230–253.

Pritchett, A. (2013). Simulation to assess safety in complex work environments. In J.D. Lee and A. Kirlik (Eds.), *The Oxford Handbook of Cognitive Engineering.* Oxford University Press.

Woods, D. (1985). Cognitive technologies: The design of joint human-machine cognitive systems. *AI Magazine*, *6*(4), 86-92.