



A framework for describing interaction between human operators and autonomous, automated, and manual control systems

Jonas Lundberg¹ · Björn J. E. Johansson²

Received: 16 August 2019 / Accepted: 9 June 2020 / Published online: 18 June 2020
© The Author(s) 2020

Abstract

This paper addresses how to describe critical episodes of interaction between human operators and autonomous, automated, and manual control systems. The first part of the paper poses three questions: (1) what levels of cognitive control are important to include in a descriptive framework for joint human-autonomy in process control; (2) how should one describe temporal developments in joint socio-technical systems; and (3) how does one analyse communication and control at the system joints. The paper proceeds by proposing a new framework for description and analysis, the Joint Control Framework (JCF), with a simple notation, the Score (JCF-S). It allows descriptions of the three previously mentioned aspects through three analytical activities: process mapping (PM), analysis of Levels of Autonomy in Cognitive Control (LACC), and temporal descriptions of human-machine interaction (T-HMI) through the Score notation. This facilitates analyses across cases and domains. The framework is discussed based on an analysis of two episodes; one work episode (from an air traffic control tower simulator); and one work procedure (from an unmanned traffic management system design concept).

Keywords Human-automation collaboration · Automation transparency · Levels of automation · Cognitive work analysis · Cognitive systems engineering

1 Introduction

Understanding critical episodes in which there is contention for control between erratic external processes, human operators, and automated or autonomous systems is a major issue for designers of safety-critical real-time applications. The human operator may be in the loop due to a design that includes human-automation teaming as a goal. However, the operator may also be in the loop due to the suspicion or knowledge that the cognitive autonomy of the automation or performance reliability is bounded (i.e., restricted to certain conditions, which may or may not be fully known). Further, most real-world systems comprise nested automated systems of varying and potentially changing levels of automation. This renders simplistic analytical frameworks insufficient for modelling the increasingly complex control relations

presented to operators today. The common characteristic of all examples above is that the end result is a situation where a human is left responsible for processes in which control is completely or partially delegated to automation, and where this delegation can change dynamically. That change in delegation occurs at the joints¹ of the human-machine system and is the result of communication or interaction between the components.

In this paper, we address how to understand episodes with potentially changing allocation of control between human and non-human actors (e.g., technical systems with varying degrees of autonomous capabilities) in joint socio-technical systems. We propose a new framework, the Joint Control

¹ The usage of the term *joints* is inspired by the concept of “joint cognitive systems” (Hollnagel, E., Woods, D.A., 2005. Joint cognitive systems: foundations of cognitive systems engineering. CRC Press, Boca Raton, FL.), which applies a systems view on human-machine constellations, suggesting that the performance of a human-machine system cannot be understood by studying its parts in isolation. The usage of perception, decision, and action points introduced in this paper is a way to describe *where* in space and time the jointedness of the human-machine system actually occurs and how it is *manifested*. However, Hollnagel and Woods’ argument that system performance only can be understood from the point of view of the joint system remains valid.

✉ Jonas Lundberg
jonas.lundberg@liu.se

¹ Department of Media and Information Technology,
Linköping University, Norrköping, Sweden

² Department of Computer and Information Science,
Linköping University, Linköping, Sweden

Framework, and discuss it in relation to episodes from case studies on drone traffic management systems and air traffic control tower work.

1.1 Understanding automation and control, rifts between related approaches

The challenge of implementing automated systems is as daunting today as it was when Bainbridge wrote her classic paper ‘The ironies of automation’ (Bainbridge 1983). This does not mean that we haven’t been able to solve a number of practical problems concerning automated systems, but the challenge remains, as there is a seemingly ever-increasing development of automated functions, which suggests that the potential for ‘automation surprises’ remains.

Within the field of human factors, the issue of how to divide tasks between humans and machines has been debated since Fitts’ famous examples of ‘Machines are better at.’ and ‘Humans are better at.’ (de Winter and Dodou 2011; Dekker and Woods 2002; Fitts 1951). This view was also reflected in the now (in?)famous *levels of automation* (Sheridan 1976; Sheridan and Verplank 1978), wherein automation increasingly (at “higher” levels) becomes more autonomous, but also less communicative, gradually excluding the human until only a silent black (autonomous) box remains. Control can also be adaptive/adaptable, shifting the trading and sharing of control over time (see, e.g., Sheridan 2017). Consequently, there is a need, not only to be able to describe these how these shifts in control occur, but also to describe when they occur. This calls for new approaches to analysis and design of human–machine interaction.

Currently, operators face an environment with nested automated systems of varying and potentially changing levels of automation. Understanding the potential interactions between these systems is a major challenge for existing analytical frameworks, as they depart from the comparatively simplistic assumptions of earlier research. The debate has gotten new fuel due to the increasing interest in artificial intelligence and automation, which has initiated new areas of focus such as Man-Unmanned-Teaming (MUM-T), Human-Autonomy-Robot-Teaming (HART), Joint Cognitive Systems (JCS), Coactive design, and similar initiatives that aim to improve the integration between human operators and the machines they are using or interacting with. In many cases, engineering psychology-oriented research has continued to investigate function allocation between humans and machines (see for example Miller and Parasuraman 2003; Parasuraman et al. 2000). However, a range of researchers have noted the potential pitfalls associated with dividing control between humans and machines according to unidimensional and categorical models (Bradshaw et al. 2013). We denote the former approaches as Extended Fitts List (EFL) and the latter as Man–Machine-Teaming (MMT)

approaches. Somewhere in between these two extremes, we find ‘good old-fashioned human factors’ (Goof-HF), which applies a dualistic view on humans and machines, focusing on interface design and automation as a way to overcome or support the limitations of human information processing.

In the MMT focus area, an approach emerged during the 1980s that focused on control rather than cognition, Cognitive Systems Engineering (CSE). The field emerged as a reaction to the dualistic, information processing-based tradition of human factors and cognitive psychology. CSE suggested that man–machine systems, especially from the point of view of performance, could only be understood as a whole (Hollnagel and Woods 1983). However, CSE was divided into two avenues of research: one based on structural models of process control and human decision making, represented by the Rasmussen tradition, and one focusing on a socio-technical view of man–machine systems, or Joint Cognitive Systems view, stemming from the Hollnagel and Woods tradition. The first CSE avenue, represented by the Rasmussen tradition of analysis (Rasmussen 1986) and design (Fay et al. 2017; Vicente and Rasmussen 1992; Westin et al. 2016) covers control through structures, from the constituent physical objects of the control system to its functional purpose through the abstraction hierarchy (AH). The second CSE avenue, represented by the Hollnagel and Woods tradition, covers control through continuous enforcement of constraints, from keeping processes within physical bounds to meeting of overarching targets (Hollnagel and Woods 2005, 1983), in a context of e.g., time, resources, and control guidance (Hoffman et al. 2014, 2013; Hollnagel 1998, 2001; Hollnagel and Bye 2000; Johnson et al. 2014; Klein et al. 2004). In the midst of the research field, we find (operator) situation awareness (Klein et al. 2010; Lundberg 2015; Parasuraman et al. 2008; Sorensen et al. 2011; van Westrenen and Praetorius 2014) and sensemaking (Weick 1993) as the focus of research, i.e., how the limitations of human information processing can be overcome by presenting information to operators in different ways.

In sum, there are several views on human–machine interaction that have all contributed to our understanding of how automated systems should be implemented. These views, or avenues of research, depart from different theoretical assumptions, which shapes the focus of research as well as the conclusions drawn from case studies and experiments. However, we believe that all approaches have more similarities than differences. This paper discusses and aligns some of the more influential theories and models into one analytical framework. This facilitates discussion and analysis of human–machine interaction in systems comprising automated functions. Considering differences and similarities in the three strands of theory, and also addressing issues with the EFL tradition, we have made an effort to create a synthesis of these different approaches, which raised a series

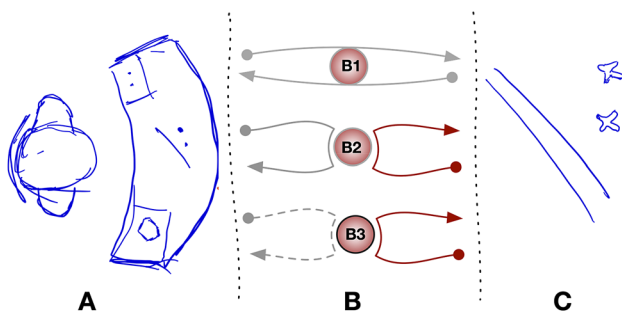


Fig. 1 Supervisory environment and process (a), automation/autonomous system/process (b), external process/system (c)

of questions and resulted in a new framework, the Joint Control Framework (JCF), along with a notation for temporal description, the Score (JCF-S). Our objective is to:

- Identify the *levels of cognitive control* that should be included in a descriptive framework for joint process control.
- Determine how to describe and understand temporal developments, regarding
 1. changes in *control and function allocation* between human and non-human actors in joint socio-technical systems.
 2. *communication and control at the system joints* over time.

2 What is control?

To understand, and then mend, the rift in theory on human–machine teaming, in particular in the area of CSE, we need to revisit the foundation of these approaches, and consequently the issue of what control is. Control has been described as the use of one process (the supervisory control process) to control another process (the external process) (Brehmer and Allard 1991). In modern systems, there is usually also a mechanical, automated, or even autonomous control structure, containing (and automating the control of) the external process. It may also automate the management and presentation of the process construct, that is, of what goes on in the process.

Many supervisory environments are quite similar in layout, in general looking much like the sketch in Fig. 1. Starting from the right (C), there is an external process (which may vary widely between domains). In this example, there is a layer of automation (B), which includes three possible scenarios. In the first scenario, there are some basic enhancements to communication, i.e., highlighting of objects in a video stream, or a direct view and direct control of the

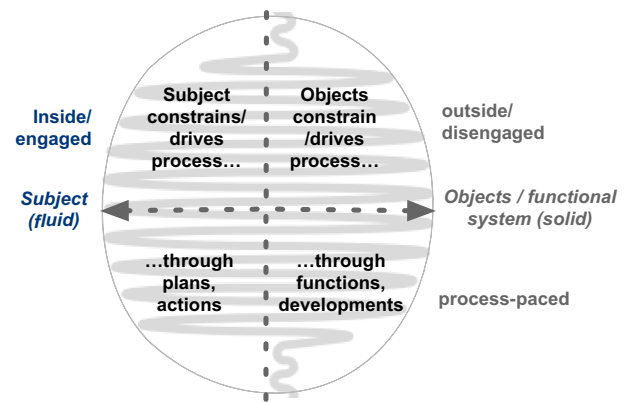


Fig. 2 Control by embedding the process in a system (right half), versus active control by continuous enforcement of constraints (left half)

process (B1). In the second (B2), the system takes a role in decision making and control, making control more indirect/passive. In the third scenario, the operator’s view/ability to control the process is reduced, perhaps cutting the operator out of the loop entirely and making the operation partly or fully autonomous (B3). See Sheridan (2012) for in-depth coverage of these basic human-automation relations.

A further complication is that in modern operating environments, many automations are present (i.e., mixes and variants of B1–3), sometimes as incremental add-ons. Rather than being well-integrated ‘systems-of-systems’, operators may have to work with several separate systems that do not always communicate with each other internally (or that are integrated, but only in a limited way). It becomes increasingly difficult to keep track of the status of the automated units as they increase in number and degree of automated behaviours (e.g., having various configurations of Figs. 1, 2, 3b). Note that this messy situation (potentially with several automations, potentially with an operator who can affect the process directly at the same time as the automation and regarding the same constraints) differs markedly from the neat divisions of human–machine work in older theoretical models (e.g., Sheridan 2012). In those older models, the assumption seems to be that there is only one automation with one automation level, although mode shifts are traditionally seen as a challenge (e.g., between B1, B2 or B3), and where it is either the human or the automation that acts on the process, although dynamic/adaptive allocations are seen as an opportunity (Hancock et al. 2013; Parasuraman et al. 1992; Sheridan 2011).

In contrast to home appliances (e.g., washing machines) that usually hide as much of the process from the operator as possible (an extreme version of B3), industrial safety-critical systems are much more open and expressive. The external process is presented on one or several screens/panels for the operator(s) in the supervisory control process (Fig. 1a).

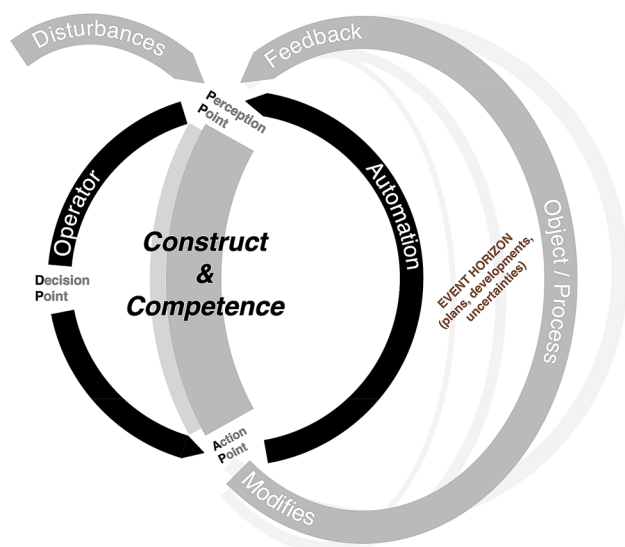


Fig. 3 Basic Joint Control Loop (JCL)

When introducing new automation into these environments, a side-effect can be more information (rather than less), on more screens or windows. It shows the status and output of the automation, potentially also showing the basis of the computation in transparent/explanatory interfaces, and/or status of equipment (e.g., sensors). For example, in Air Traffic Management, the automation of medium-term conflict detection (as a tool) adds information, e.g., in the form of a time-distance to conflict window, while the previously available information remains. The operator then needs to choose a strategy for selecting between information sources in time-critical situations (Lundberg et al. 2015). This use of large informative screens is representative of environments such as air traffic management, vessel traffic management, train traffic management, and nuclear power control.

In these messy situations, it becomes important to consider not only the basic control relations (Fig. 1b, Sheridan), but also the joints between human and machine(s), in terms of what is communicated and controlled (e.g., high-level goals at one extreme, low-level information such as the status of particular objects at the other extreme)—over time (is the interaction paced in a controllable manner for human operators?). In doing this, we need to consider an additional basic complication: the extent to which the process is controlled by erecting and monitoring control structures (including delegation to automation), versus by continuous active enforcement of constraints.

2.1 Control by embedding the process in a system, versus active control by continuous enforcement of constraints

Figure 2 depicts two different kinds of control situations. The left part shows active control by continuous enforcement of constraints on the process. In CSE, this has been addressed by the Hollnagel-Woods tradition, e.g., through the contextual and extended control models (Hollnagel and Woods 2005). The right part of Fig. 2 shows control by delegation, wherein the process is embedded within a structure. In CSE, this has been addressed by the Rasmussen tradition, e.g., through the abstraction hierarchy and decision ladder (Rasmussen 1986).

However, being required to select a different framework depending on whether the process is one of active control or structural control can be a problem. This is because somewhere in between the two extremes of Fig. 2, the subject may be partly engaged in control, and partly engulfed by a system, something like going down a series of rapids in a canoe. Due to this duality, we need one framework that can describe processes that are more subject-driven, processes that are more object-driven, and those that fall in between.

In addition to these subject-object relations, Fig. 2 illustrates the need for descriptive power for shifts between plans/structures and self-paced/process-paced developments. Moreover, we need to consider autonomous systems, i.e., objects that can also be interpreted as agents acting with a purpose. Such objects pose a particular challenge to modelling and understanding as they exhibit behaviours that are seemingly rational, but (usually) lack the understanding and creativity of a human being.

2.2 Active control by continuous enforcement of constraints

A central theme in the CSE tradition is active control of dynamic on-going processes through continuous enforcement of constraints. A control loop is the basis of many models (see, e.g., Hollnagel and Woods 2005; Lundberg 2015; Lundberg et al. 2012; Neisser 1976). We present a basic control loop in Fig. 3, both as the basis of the more advanced joint control model presented later in this article (Fig. 6), and to highlight common ground with other well-known models. Highlighting common ground simplifies analyses from several perspectives (i.e., using more than one model/framework). We have included both an operator and an automated system in the loop to highlight joint action and delegation. This stands in contrast to other uses of control loops in CSE such as the Contextual Control Model (COCOM) or the Extended Control Model (ECOM) (Hollnagel and Woods 2005). Figure 3 describes joint actions

between a human operator and automation controlling a dynamic process/context. The figure contains:

- A control loop (black circle) shared by a human (left half) and automation (right half),
- The bold grey circle indicates the aspect that is presently being controlled, lighter circles show an event horizon of past and present.
- Uncontrolled disturbances that could affect the system are indicated by the top left arrow.

The model also contains the construct, or mental model (Salas et al. 1994), that is being constantly changed in response to information about the process to be controlled. In addition, it contains three key points:

- an action point, AP (where the operator can affect the process),
- a decision point, DP (where the operator can decide what to do), and,
- a perception point, PP (where the operator can receive information about the process).

These points are crucial to understand with regard to time (see Johansson and Lundberg 2017).

When considering processes with inertia/energy, the “event horizon” (Fig. 3, light grey circles to the right) becomes important. It represents the look-ahead time for events that are already in motion. The horizon can be seen as a combination of operator plans and process developments, as well as visible parts of the environment in relation to plans/processes (situatedness/structure horizon). An appropriate time for adjusting this process is constituted by the PP, DP, and AP of relevance to understanding and controlling the process on the event horizon. The process itself may have energy and momentum (e.g., a vehicle that moves), or it may execute in discrete steps. This effort can also be affected by disturbances (Fig. 3, top left). Therefore, it is important to consider how human action is tied to artificial or natural² leverage points to affect the process.

Furthermore, the coupling of on-going processes is important, for example when two aircraft need to use the same runway. This, then, is a dependency with a duration, i.e., the dependency ends when one aircraft is out of the way. Therefore, in more process-driven situations the AP,

² By “natural” we refer to leverage points that occur as a result of human–human interaction, occurring at a pace that can be expected based on common knowledge of human behaviour. By “artificial” we refer to leverage points that have been designed into technical systems, paced by the interaction between different technical components. Such leverage points can be more or less transparent and predictable from the human operator’s point of view.




Fig. 4 Basic interaction grid, for describing exchanges between supervisory process and external process, through a medium/interface

PP, and DP are often tied to the tempo and flow of the process, whereas in more self-paced situations, the subject is more in control of those points. The AP, DP, and PP may also be tied to organizational routines and processes, e.g., specific decision meetings at specific dates and times (see e.g., Johansson and Lundberg 2017).

2.3 Control by embedding the process in a system

The purpose of control is to achieve a desired state in a system. This is exercised through a *control process* acting on an *external process*, which can be mediated through some kind of structure (Fig. 2, right). Process control through adjustment of structures is a central theme in the Rasmussen CSE tradition.

A control process consists of a series of actions with the explicit purpose of influencing the state of the external process so as to keep it within certain boundaries that describe the desired state of the external process. The medium (usually a technical system/automation/autonomous agent) is what mediates such actions. The control process is exercised by the *subject* on the *object* (as described above). A control situation can be illustrated in a grid, Fig. 4. In the figure, a subject in the supervisory control process (horizontal line), attempts to control an external process (vertical line), through the interface of a technical system (the medium).

A core question, then, is at what *level of cognitive control*, from overarching goals to specific procedures/actions, can the subject interact with the system (and thus also what the

Table 1 Comparison of the levels of three frameworks

	Activity	Process/system	Construct
	Extended control model, ECOM	Abstraction hierarchy, AH	Situation awareness, SA
WHY			
6	–	–	What goes on (frames/potential)?
5	Targeting	Functional purpose	What goes on? (a specific frame)?
What			
4	Targeting	Abstract Function	What about it? (concerns, e.g., risk)
3	Monitoring	Generalized Function	What about it? (expected events)
How			
2	Regulating	Physical Function	What can we see? (objects, cues)
1	Tracking	Physical Objects	

operator might want to know and affect)? For instance, can the operator directly see and/or set the goals of the automation (i.e., that an aircraft should land), or can the operator only see and/or affect specific object settings (i.e., the speed and direction of an aircraft)? A follow-up question has to do with what those levels are, or as is the focus here, what levels are most important to include as a starting point, in a description focusing on human-automation collaboration?

3 Theories of levels of autonomy in cognitive control

To identify a set of suitable analytical levels to describe cognitive control, we can revisit previous frameworks and models in CSE (e.g., ECOM, AH, SA), as shown in Table 1. Although the frameworks and models differ substantially in other ways (i.e., in assumptions regarding how control is achieved, Fig. 5), the levels used in the models are quite similar, which is illustrated in Table 1.

The Abstraction Hierarchy, together with the Decision Ladder (Rasmussen 1986), describes control by delegation, wherein the process is embedded within a structure (Fig. 2, right). The Extended Control Model (Hollnagel and Woods 2005) describes active control through continuous enforcement of constraints on the process (Fig. 2, left). Although these models differ in philosophy and focus, they intersect at two points: Firstly, they have similar analytical levels (see Table 1); Secondly, three functions are central: perception, decision making, and action. These well-established elements are therefore reused in our model (Tables 2, 3 and Fig. 6). Additionally, they lack a modelling framework for time and developments on the event horizon, which is required to model episodes, operating procedures, or envisioned interactions in design concepts. We remedy this issue in our model.

3.1 Structures

Firstly, the Abstraction Hierarchy (AH) is at the core of Cognitive Work Analysis (CWA), which is the Rasmussen CSE tradition (Naikar 2017). The AH describes the process embedded in the control structure (Fig. 2, left) as seen at various levels of abstraction, ranging from high-level system purposes down to specific physical processes in technical systems (Table 1). It describes process status versus system constraints, from constituent object parts to the functional purposes (i.e., the effect on an external process) of parts and of the system as a whole. The description makes dependencies between layers of abstraction traceable, in the best case all the way from object properties to the degree of achievement of the functional purpose.

Although the number of levels of the AH are not strictly prescribed (Vicente and Rasmussen 1992), the use of five levels is very common (see Table 1). Visualizing these layers is the foundation of what is called Ecological Interface Design (EID) (Vicente and Rasmussen 1992)—to show the process status and limits, in a way that humans can follow, through layers of abstraction.

To understand how operators maintain control through structures, CWA uses a second model in addition to the AH (Table 1): the decision ladder (DL). An outline of the DL is

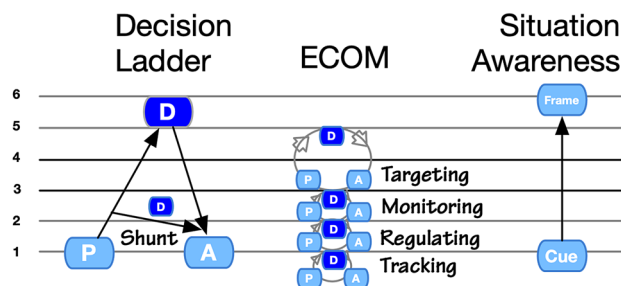


Fig. 5 Analytical levels as described through conceptual models by Rasmussen, Hollnagel and Endsley

Table 2 Levels of autonomy in cognitive control (LACC)

Levels of cognitive control (autonomy and interaction)	Level description	
WHY	6 Frames	Problem frames; what the system is (ascribed system type), situations and contexts. Recognition of a situation that is associated with well-known developments can result in a well-specified event horizon (see level 3)
	5 Effects	Functional purpose (why the system exists) and overarching goals (what effects we want to achieve with the whole system/activity)
WHAT	4 Values	Abstract system or value functions, such as calculations, qualities and tradeoffs (weighed criteria)
	3 Generic	Generalized/generic plans, typical states, patterns, and functions, i.e., non-situated (determinable before execution/implementation) aspects. Having engaged in or predicted a generic level activity, sets up an event horizon of expected or planned developments
HOW	2 Implementations	Assignment of resources to functions, constraints due to implementation properties. The adjustment and timing of constraints in situated (aspects that can be determined during but not before execution) planning Implementation decisions that can be made ahead of execution constitutes a more specific event horizon than on the level above
	1 Physical	Object properties and the enforcement of constraints through physical action. The ability to see objects ahead gives a situated basis for establishing an event horizon

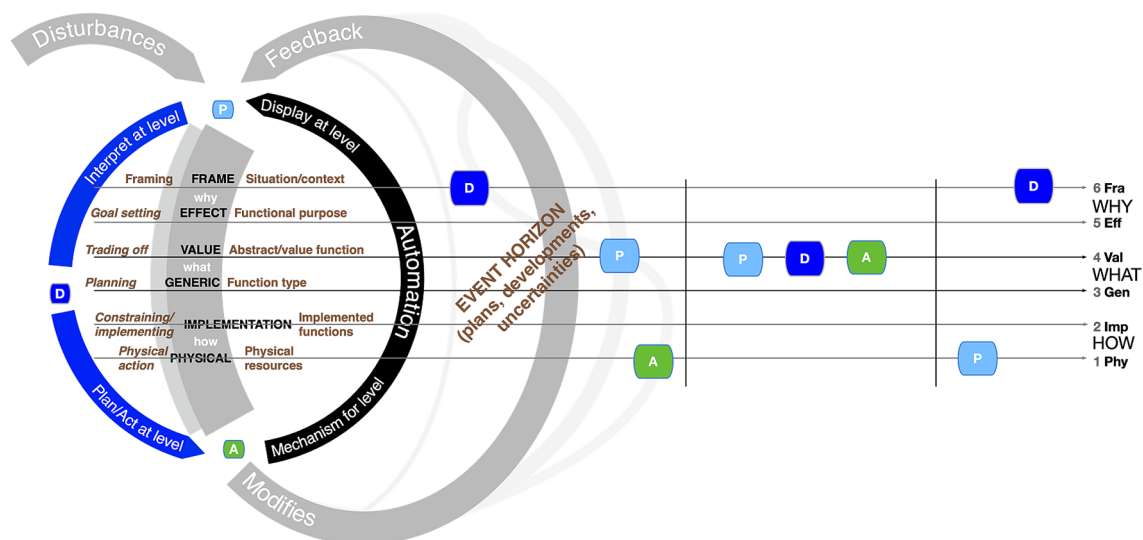


Fig. 6 Joint Control Framework

presented in Fig. 5. It describes how operators first attempt to understand (undesired) process deviation and control structure state through examining the system (left “leg”). It then describes how they attempt to devise a control task to get the process to a desired state by adjusting the control structure (right “leg”).

The DL also includes the notion of an elaborate set of “shunts” between the legs of the ladder, to take shortcuts from a partially understood situation to a pre-formulated control task. Thus, on the way “up” (from PP at the lowest level and DP at the highest level), a DP can occur at any lower level based on incomplete information (e.g., recognition). This shunt can then go down all the way to one specific action, or end up higher in the model by, for instance, having

a plan (procedure), or by knowing/recognizing the desired target state.

3.2 Active enforcement of constraints

In response to procedural models in which the operator/controller is curiously absent,³ Hollnagel (1993) proposed the use of contextual control models, focusing on which factor(s) control the next action (e.g., environmental cues). The Extended Control Model is at the core of CSE,

³ See also Hollnagel, E., 1993. Human reliability analysis: context and control. Academic Press, London, UK. for a comparison between contextual control and the decision ladder.

according to the Hollnagel tradition. The ECOM describes control as something that is achieved by the subject (i.e., the controller) through continuous action (Fig. 2, left). An outline of ECOM is presented in the centre of Fig. 5. It describes control, from continuous tracking of target values, to the setting of overarching targets. The description makes control traceable, in the best case from slips in tracking to the loss of control of overarching goals.

In ECOM, each level specifies the *kind* of knowledge or skills used for decisions and adjustments (action). For example, the lowest level (tracking) is knowing (or sensing) deviations from a track of target values (to make corrections). The next level up is concerned with issues that are central to deciding the track of target values (which are then used in tracking). For instance, to overtake a cyclist ahead involves: knowing that tracking of the own bike is close to failing before having reached the target speed; and knowing that this is due to a road that becomes slippery while attempting to overtake the bike ahead.

At the level above that (monitoring), the concern is with matters that are central to planning and re-planning decisions. ECOM can model decisions that are simultaneously made on all levels, i.e., it does not prescribe a sequential process. It can also be used to describe how the various cognitive functions are allocated to different components (human or machine) in the man–machine system depending on the configuration of the system (Aminoff et al. 2007; Johansson and Stenius 2015).

3.3 Grasping what is going on to make decisions

Central to decision making (regardless of whether it is through adjusting a structure or by enforcing constraints) is the need to know and understand (something of) what goes on. This can be tied to particular theories of control (as above), but has also been discussed in more generic terms, e.g., in terms of sensemaking and situation awareness. Lundberg (2015) relates SA to sensemaking, but also to ECOM as well as to modern systemic SA models (see, e.g., Sorensen et al. 2011).

The activity of grasping what goes on (sensemaking) is perhaps the most critical part of SA. Referred to as *framing*, it is central to naturalistic decision making (NDM). In NDM, framing represents the current understandings or hypotheses of what is going on (Nemeth and Klein 2011). It is the ability to generate new frames (how to make sense of this?), or to recognize situations. Frames represents the way one thinks about a specific issue in a context, i.e., our frames reflect how we understand and act upon the world in different situations (Lakoff 2014). Without frames, there would be no values, no norms, no expectations, etc. However, this is a two-way relationship (sometimes referred to as reflexivity), as our understanding shapes the way we act

upon the world, and the way that our actions change the world in turn shapes our understanding of the world further. Thus, frames are not static, although they are usually built up over a long time. Frames represent the most stable part of the cognitive levels, and largely guide the sensemaking process (Klein et al. 2006). Framing can be seen as the top (most abstract) level of understanding of what goes on; in a framework of SA (Fig. 5, Table 1) roughly consisting of (see e.g., Lundberg 2015):

1. what (framing),
2. what about it (specific details),
3. what can we see (cues and important objects),
4. when.

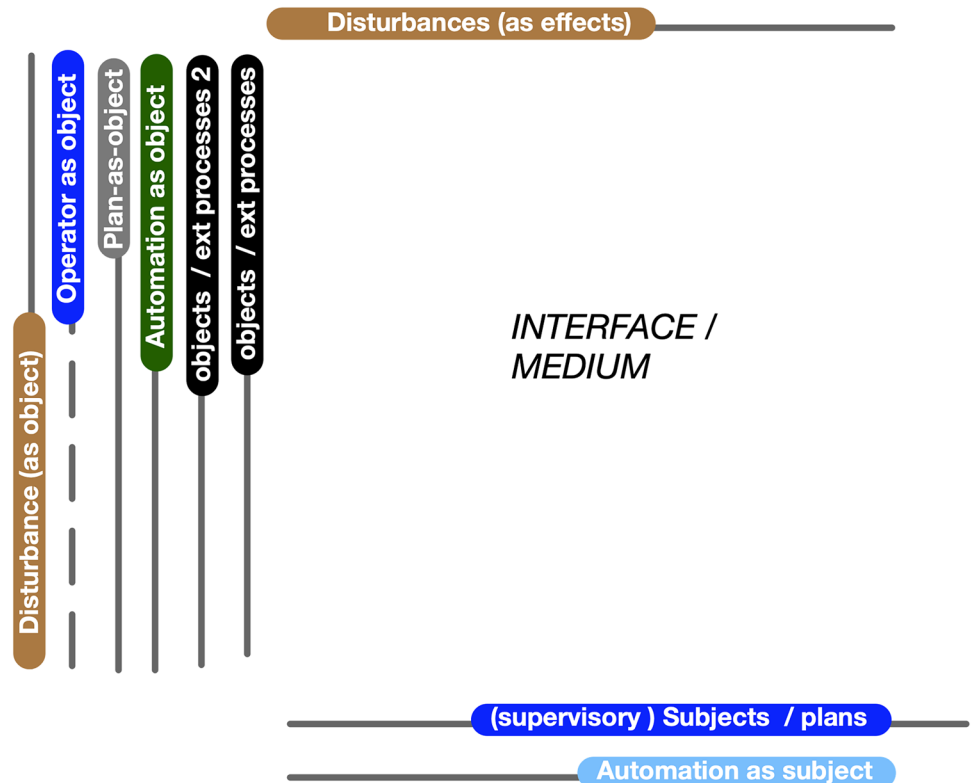
4 The Joint Control Framework (JCF)

In the following sections, we propose a framework for understanding how control processes float between humans and technical systems in different situations, forming joints of interaction, and how the context influences the way in which the subject-object relation in control can be analysed and modelled, and hence how it should be allocated. A framework for describing interaction between human operators and autonomous, automated, and manual control systems, the Joint Control Framework, is illustrated (showing the Joint Control Loop, the Levels of Autonomy in Cognitive Control, and the Score notation) in Fig. 6 and is described in the following section.

We propose an approach to modelling joint control that consists of three steps, explained below. Although each step will need to be revisited during analysis, we suggest that the first iteration progresses as follows:

1. Process mapping (PM) is the starting point. This step is needed initially to decide which processes and agents (subjects and agents) to include in the analysis. In particular, it is important to determine which subject(s) to focus on, as they control the external process, and perhaps also each other, jointly or disjointly. Also, it is useful to map which process(es) should be modelled. It is also useful to bear in mind that some processes may overlap in time when deciding on which episodes the analysis should focus on.
2. Examining the Levels of Autonomy in Cognitive Control (LACC) is the second step, to get an initial idea of the levels at which joints occur. This could have to do with the levels that exist in current systems, or that would be desirable in a future system (that, e.g., is being designed). Automation may also be limited in the ability to work (at all) at certain levels, and in certain situations. Although humans may have limits as well (e.g., in train-

Fig. 7 Objects and subjects in the decision situation



ing), all of these topics could be relevant to address in the LACC.

3. Finally, a Human–Machine Interaction Temporal analysis (HMI-T) can be performed to understand the joint cognitive system interactions over time. We propose the Joint Control Framework Score, (JCF-S) notation for HMI-T analyses. This, e.g., shows the consistency or variation of interaction, e.g., if there are mainly high joints in the score, or a variation that mixes joints at various levels. It also shows distances between joints at various levels that must, e.g., be covered by operator competence. Furthermore, it reveals overlaps in control of simultaneous processes, which can be problematic to manage, and highlights temporal constraints in the control process.

The ordering of the steps should be seen as a useful starting point, from which the analyses can then be revisited in a suitable order.

The Score can also highlight the basic shape of interactions. This can be useful for in-depth analyses using other approaches (i.e., if the Score resembles the core models of the approaches, see Fig. 5).

5 Process mapping

When addressing real-world control problems, we first need to consider that the real world is often more complex than managing a single process. Therefore, we first have to engage in the analytical activity of process mapping.

For instance, an ATCO with two aircraft to manage initially deals with two processes. In Fig. 7, this could be represented by showing several vertical lines to the left side of the grid. A central part of the ATCO work is then to integrate them into a larger process/plan. Taking a mundane example, when going downhill skiing, other people going downhill nearby constitute separate processes. Each person then needs to coordinate their movement with respect to that of others (without relying on centralized control, but rather on “rules of the road”).

Once again, there is a stark contrast between coordinating with humans and with automated systems. In such a case the context and purpose of the automated system are important enablers of both coordination and trust. Unless the human part is able to clearly understand the abilities and limitations of the automated system, control, both in terms of how it is experienced and in terms of how it is exercised, is going to be limited or non-existent (Bradshaw et al. 2013).

An important point to consider here is that (in addition to recognizing familiar patterns) the ability (or inability) to integrate several processes into one is a major part of what




Fig. 9 Temporal modelling, the Score (JCF-S) for HMI-T analyses

control/adjust as part of the external process. In contrast, it sometimes acts more like a team member (Klein et al. 2004), as a subject in the supervisory process.

5.3 Disturbances as objects and as effects

Furthermore, as illustrated in Fig. 7, disturbances may be seen both as an effect (on the process) and as an object—the operator may attempt to control the disturbance per se, or to control the process as affected by the (uncontrolled) effects of the disturbance. Especially when a system is facing disturbances, it may be hard to figure out what is actually going on, an activity usually referred to as sensemaking (Lundberg et al. 2012, 2014; Weick 1993), in which framing is central. At times, the disturbance may also affect the controller and the control structure (as an object).

6 Understanding the joints: Levels of Autonomy in Cognitive Control (LACC)

As recognized in previous theory (Table 1), it is often useful to view control in terms of levels of control. It can be useful to describe what the whole system (human and machine) needs to do, as well as function allocation between human(s) and machine. We may need to describe the performance limits of automation, the levels it is able to work on: for instance, the ability of the system to recognize known situations; its ability to plan for various contexts, and then implement the plans; function allocation and collaboration in evaluation high-level performance (e.g., efficiency) versus goals; or the required plans and actions needed for adjustments.

Based on frameworks that are similar and have generally been seen as useful in previous theory (Table 1), we suggest a new basic framework, the Levels of Autonomy in Cognitive Control, LACC (see Table 2). It simplifies analysis using one set of labels for the levels (instead of three sets). It covers three aspects that are central to control: control of

processes contained in a structure (AH), active control of processes through continuous adjustment (ECOM), and having a construct/awareness of the process for decision making (SA)/naturalistic decision making (NDM).

The top third of the six levels (Tables 2, 3) is concerned with the question of WHY with regard to humans and systems (overarching goals and functional purposes), the middle third is concerned with WHAT the system does (abilities, functions), whereas the bottom third concerns more specific detail regarding HOW it is realized (implemented functions, constrained plans, objects involved). These six main levels correspond to what has been used in AH and ECOM analyses for many years (see Table 3 in Appendix A; Figs. 6, 8, 9).

This definition raises the question of what is considered to be a success, which can also be described using the LACC levels, versus system performance. Limitations to autonomy on various levels can be expected (even for humans), resulting in “bounded autonomy”. A core question is then what happens at the boundaries, and how the system can monitor and manage its boundaries. This is a core concern regarding supervisory control in highly automated or autonomous systems. Furthermore, even for a system that can, in principle, act autonomously, humans might want to intervene and inform the system about what is going on, e.g., an emergency transport (level 6), set/alter goals (level 5), change priorities (level 4), adjust what types of plans and functions that are used (level 3), how they are implemented (level 2), and adjust the resources that are available to the system (level 1). The framework also relates to previous notions (Sheridan 2012) of programming or scripting of automation, by establishing the level of competence to which the automation has been programmed—e.g., has it been scripted specifically for one particular movement (level 2)? Has it been given a generic script that is adaptable to a variety of situational constraints (level 3), or can it even generalize from particular plan instances to generic functions (increasing its competence from level 2 to level 3)?

In sum, this new framework (Table 2) simplifies our work (compared to using the somewhat misaligned frameworks in Table 1) when describing information exchange in automation-mediated human-centred process control. The framework harmonizes the previous three frameworks (ECOM, AH/DL, SA) by:

- a) placing framing on a sixth level, outside of the activity or system,
- b) tentatively splitting the Targeting level of ECOM into two levels (effects and values), and,
- c) tentatively splitting each of the three SA levels into two (frames/effects; values/generic; implementations/physical).

7 Temporal description notation, the Joint Control Framework Score (JCF-S)

In many domains, it is important to understand how interactions unfold over time, e.g., considering parallel process developments and time pressure. These interactions occur at joints between human, machines, and processes. The joints emerge in a physical medium of some sort. Between humans and machines, the mediating artefact is a Human–Machine Interface, HMI. At the interface, the challenge becomes one of displaying part-whole relations in a way that includes time, process/sub-process, and relations between information about the processes at various levels (Figs. 8, 9).

Although when it is seen as an object (e.g., lever, button) the interface is always at level 1 (physical objects), we are concerned here with what the interface object represents or controls about the process/objects that the joint system is attempting to manage. For instance, a lever could be used to instruct the automation to carry out a process of balancing between goals. That joint would then be modelled at level 4, as an action point. If, instead, the lever controlled some physical object status (e.g., speed of a moving object), then the joint would instead be placed on level 1. In that case, the decision would still be about balancing between goals, at level 4. This would then be modelled as a decision point at level 4. The analysis would then show the level discrepancy between decision point and action joint. The analyst would also need to model the information required for decision making, examining the level at which it was presented in the interface, e.g., at level 1 (object status) or perhaps level 4 (calculated measures). The analysis could then both show level discrepancies, and also the sequence and amount of information at various levels that the operator would have to collect to make the decision. The interaction patterns can also be examined by the analyst to understand whether they are typical or atypical of operator work and training.

For this analysis, the LACC could be placed in a grid (Fig. 8) between the subject (planner) and the object (structure), similar to Figs. 4, 7. This ‘LACC grid’ could then be used for describing snapshots of single interactions (joints) frozen in time.

However, the grid lacks a temporal dimension, just like one depiction of the decision ladder or one depiction of the ECOM (Fig. 5). Each specific exchange is part of an ongoing flow of exchanges, with both a history and a future. Thus, we need to conduct an analysis of temporal descriptions of human–machine interaction (HMI-T). To conduct an HMI-T (to describe episodes / plans / the event horizon), we need to use a simplified layout (Fig. 9) of the JCF (Fig. 6) for a temporal description, the JCF Score.

In Fig. 9, Joint Control is represented as six parallel lines (the Score, JCF-S), to enable the modelling of temporal (sequential and parallel) developments at various levels of human-automation interaction. Each line represents one LACC level. The DP, PP, and AP of the joint system can then be placed on these lines, forming an event horizon of plans and developments at different cognitive control levels. The Score can indicate actual occurrences/duration or intended procedures, as well as potential/leverage points. We provide two examples below of how the analysis can be conducted.

8 Interaction episodes—joints in human-machine systems

As a basis for our discussion, we use the Score to describe short episodes from two cases. Our first case represents analysis of empirical data, and of rather direct control of a process, with a low degree of automation. It shows how to model a landing clearance, using an Air Traffic Control Tower simulation (see also Lundberg 2015).

Our second case represents analysis of a new design, and of highly automated work. It shows how to model a tentative episode in a future highly automated system for management of intense drone traffic in cities. In the first half of this episode, traffic follows the airspace structure and is adjusted through automation. Manual process control occurs in the second half (see also Lundberg et al. 2018).

We refer to the six LACC levels by number (indicating the analytical focus corresponding to cognitive levels) and the first three characters of the label of the level (making it easier to recall the labels, e.g., 1phy to indicate the first LACC level, 1 physical).

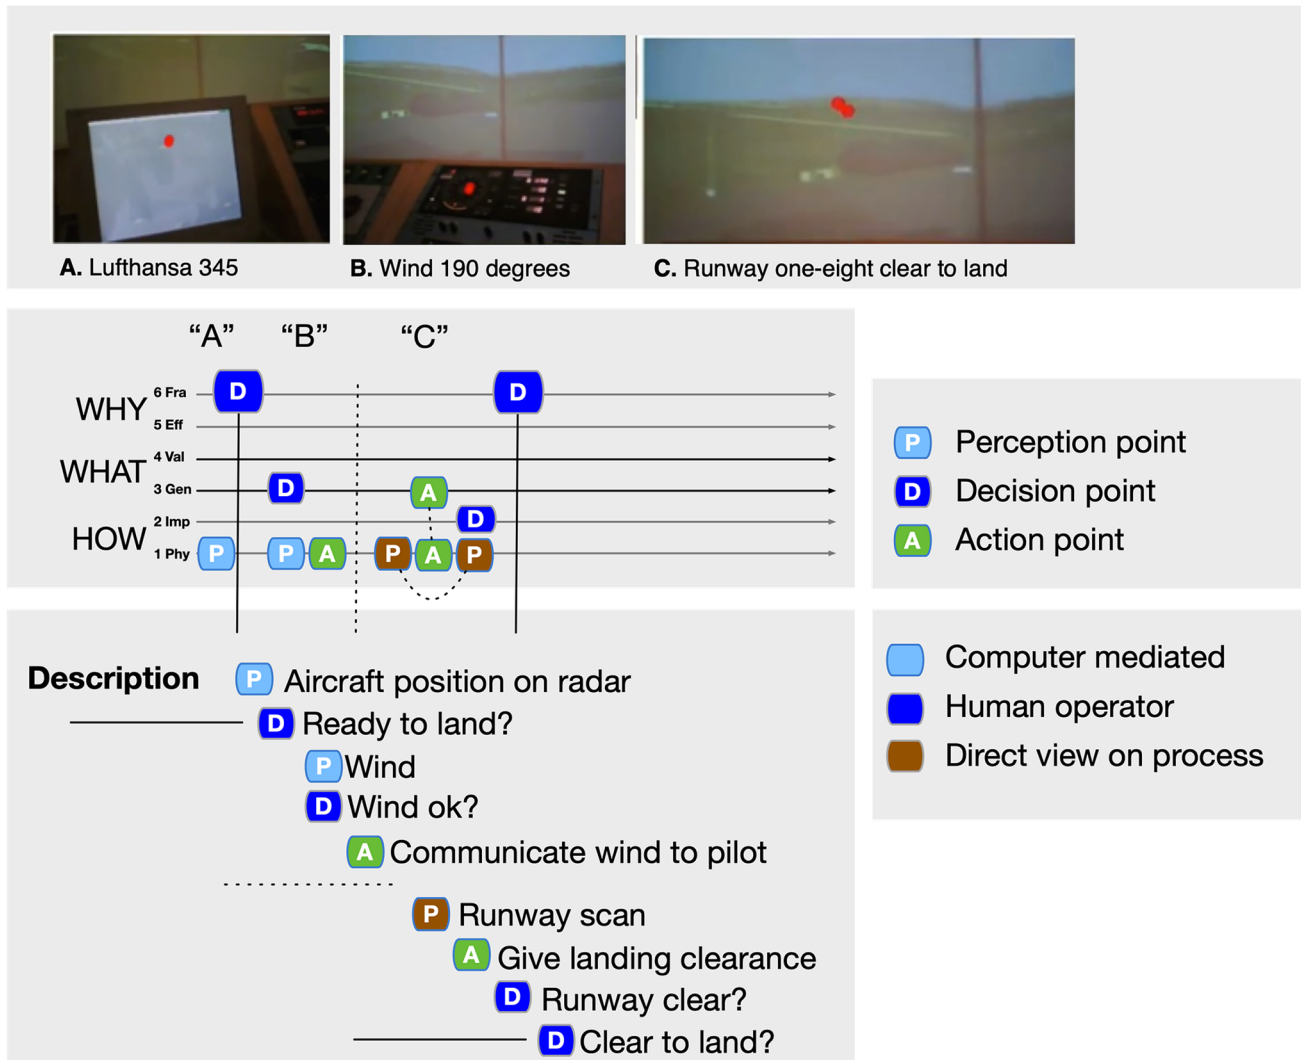


Fig. 10 Episode 1. Landing clearance score

8.1 Episode 1, manual control, in air traffic management


In Episode 1 (Fig. 10), three steps (A–C) are performed in a landing clearance. In step “A”, the ATCO checks the aircraft position and call sign on the radar display. This is modelled as a PP on level 1phy, and a decision regarding what is going on (ready to land) on level 6fra. Step “B” is a wind check, again with a PP on level 1phy. This time the DP is on level 3gen (does this adhere with the current plan for the landing?), and an AP that communicates the information to the pilot at level 1p. In step “C”, the ATCO continues to the runway scan, checking to see if the runway is clear (DP at level 2imp), using a scan of the runway (several PPs along the runway, at level 1p). The result is a decision at level 6fra about what is going on here (that the aircraft is clear to land), which is communicated (at level 3gen, monitoring

the process/on-going plan, giving an ok to proceed) together with information on which runway is clear (at level 1per).

This example shows a very low degree of automation (with all PPs at the lowest level) but with DPs at higher levels. It also shows a low level of digitalization with the ATCO communicating information verbally and at the lowest level, and also using a direct view of the process in step “B”.

8.2 Episode 2, highly automated system, in drone traffic management

Episode 2 (Fig. 11) uses the Score to describe an example from a prototype of an air traffic management systems for drone city traffic. Although it is a prototype of a design concept, we can analyse it nevertheless (see also Lundberg et al. 2018). This system works at a much higher level of autonomy (of cognitive control) than the system in Episode



Description

- P** observes a high even level of congestion or noise on the heat map
- D** decides to introduce a traffic limit in the area
- A** sets a traffic limit, the system automatically enforces it
- P** monitors the development on the heat map
- P** monitors the list of upcoming flights to be grounded
- D** **A** excludes selected flights from the limit
- P** monitors the development on the heat map

Fig. 11 Episode 2. Management of unmanned traffic, high level of automation, score

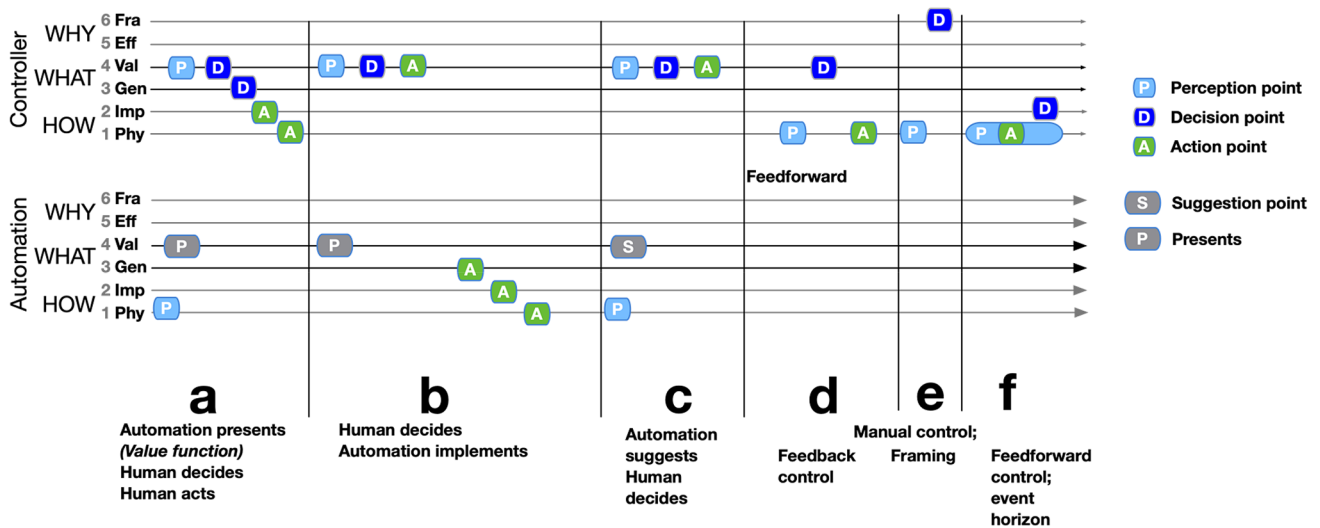


Fig. 12 The Score notation, showing some basic interaction patterns with different function allocations. **a** a human action on a lower level than the automation presents information on; **b** the automation imple-

menting plans, in supervisory control; **c** the automation suggesting plans; **d** manual feedback-based control; **e** framing; **f** manual feedforward control

1, with monitoring of an abstract measure shown as a heat map (e.g., congestion or noise, PP at level 4 val). In this concept, a limit can also be set directly on the priority measure (AP at level 4 val), resulting in an automated adjustment of the amount of traffic (e.g., fewer drones are granted permission to fly). The operator can then monitor the effect of this limit using the heat map (level 4 val).

The operator could also monitor the automation’s decisions about the flights that are to be denied take-off clearance, overruling the decisions on a drone-by-drone basis (at level 1per). This would shift the character of the work to more direct manual process control, by manually checking each flight, deciding, and then acting (approving/cancelling).

9 Score analysis and discussion

Having described how two episodes at low versus high automation could play out, we can now discuss key points in the episodes in more detail. In this discussion we refer both to the particular episodes and to abstract Score patterns (Fig. 12) highlighting more general issues. These Score patterns (Fig. 12) vary in function allocation between operator and automation. By the Score, important episodes, or *joints*, in a human–machine control process can be highlighted. A “joint” represents the place in time and space where human and machine-allocated control processes interact with each other, defining the subject-object relationship at the current joint. The Score thus facilitates HMI-T analyses.

9.1 Analysis of the system joints

Analysing the order of the different points in the Score (AP, DP, PP), we can see differences in control tactics. We can see a DP before the PP (Episode 1, Fig. 12f), with a tentative/projected event horizon (for the continued landing, not drawn in the figure), indicating feed-forward control. Conversely, we can see a DP after the PP (Episode 2, Fig. 12d) (perhaps in some cases after each PP, in really capricious control), indicating feedback control.

In Fig. 12e (Part A in Episode 1) we describe an example of framing, where an operator is cued by the properties of an object (aircraft, level 1phy), and recognizes that its process (2imp) makes it ready to land (6fra). This also (most likely) prepares the operator for what comes next, the procedure for managing the landing (level 3gen), i.e., the approximate placement of points in Fig. 12d (part B of Episode 1).

In addition, we can identify joints between control processes. We could, for instance, identify a control point (AP) that is a communication or delegation to another human agent. (Episode 1, Fig. 12f).

9.2 Episode scores

One basic question has to do with the cognitive control levels at which communication and interaction should or does take place. We can firstly describe the levels of cognitive control at which the system operates at each joint, and then analyse the episodes. For instance, in Episode 2 (Fig. 12b), the system displays a performance indicator, congestion, at level 4 val, with a decision point regarding congestion at the same level. In contrast, in Episode 1 the operator observes particular objects

through the tower window (1phy), with decision points at level 6fra, 3gen, and 2imp. The transitions between the joints on and between LACC levels can then be analysed, as follows:

9.3 Transition episodes I, on the same level: the coincidence of having the right information at the right time

The AP and PP can converge on the line (i.e., be in the same spot as the DP). This corresponds to a situation where the human needs to decide (DP) and act on “level X”, and has information (PP) and a leverage point for action (AP) available at the interface on the same level. This is what we see in the first part of Episode 2 on unmanned traffic management, and in the first three steps of Fig. 12b. It corresponds to having “*the right information at the right time*” (and a means for action, as well). Although the DP, PP, and AP may coincide at times, in dynamic systems there is an element of unpredictability. This means that we cannot accurately anticipate which information will be relevant for each step in a process. If there are uncertainties in the environment—for example if there is an unexpected object on the runway, a system breakdown, or a sudden weather shift—then the timing of the event appearance may be impossible to accurately anticipate. Thus, designing a system that will always provide “*the right information at the right time*” would require solving the seemingly intractable problem of perfect anticipation, rendering this ambition a chimera, as pointed out by Johansson and Lundberg (2017). Therefore, we need to consider other alignments of AP, DP, and PP. Moreover, the operator may have cause to doubt the information that is presented, e.g., requiring exploration downwards to figure out how it has been derived from lower levels.

9.4 Transitions episodes II: upwards

We have already considered a situation in which there is a PP and AP at the same level (same point) as the DP. Less ideal situations can occur, for instance, when the human must monitor the work of the automation. Either going “downwards” (from AP, DP, or PP) to check the lower-level implementation, or “upwards” to check higher-level consequences. In these cases the automation may be either transparent or a “black box” at lower or higher levels. Through the Score we can analyse these situations as episodes. The cognitive challenge is exemplified in Fig. 8.

Our episodes exemplify transitions upwards. Starting with manual control (abstracted in Fig. 12d), the human might need information at the 4 val level, but information is only available at the 1phy level. This would correspond to the situation in Episode 2, if the heat map was removed (in which case only the drone positions and movements would be visible). That would correspond to a re-allocation of the function of determining the congestion level from automation to human

operator. This is also similar to the situation in Episode 1 (Fig. 12f), in which the operator needs to check object positions and status, to determine whether an aircraft is clear to land (e.g., whether the wind is ok, and whether the runway is clear). In this case the human must infer the information at the 2imp level from the information on the 1phy level (upwards).

9.5 Transitions episodes III: downwards (automation transparency)

To understand how the automation arrived at higher-level decisions, suggestions, or presentations of situation status, downwards transitions must be considered. To examine downwards transitions further, Episode 2 (abstracted in Fig. 12b) provides an example in which the operator needs to decide whether the current congestion level is ok (and what to do about it), at the 4 val level. There is a PP at the same level in the system, displaying the current trade-off, which makes it easy to directly see the current state versus the decision. If the AP is on the same level (e.g., using a slider to set a congestion limit), the adjustment is also easy.

However (Fig. 12a), if the actual adjustment function were to be re-allocated to the human, and if it was directed at the external process at a lower level (e.g., 1obj, addressing the particular drones that are involved), then the operator would need to anticipate, know, or experiment with the potential effect that lower-level manipulation might have on the higher level. The operator would then also need to decide on or devise a procedure for the situation type (3gen) and then manage the particular situation (2imp) by addressing particular drones (1phy).

Furthermore, even if the operator could give directions to the automation at level 4 val, the operator might doubt the automation’s competence, in which case the operator might want to inspect the automation plan for implementation (Fig. 12b), or inspect the basis of a suggested suitable level of congestion by the automation (Fig. 12c). Both of these transitions would occur downwards. In this case, for the design of automation transparency, the operator would need to see not only the information at each level, but also the relations between the levels. See, e.g., Fig. 8 for examples of what relations the operator may need to see.

In sum, with low-level automation (low level of cognitive control), the problem is to project upwards. With high-level automation (high level of cognitive control) the problem is the opposite—to understand the lower-level grounding.

9.6 Transitions episodes IV: the automation monitoring the operator

In a case of reversed positions between automation and operator, the automation might attempt to infer (see, e.g., Fig. 8) what the operator is up to (automation in a supervisory control

position) based on observing operator actions. In Episode 1, the automation could monitor which objects the operator inspects (at level 1phy in the episode), and the unfolding order of perception points. An upwards inference would be required to figure out what the operator is actually doing (e.g., what function is performed, upwards to level 3gen). However, this inference may not be possible due to ambiguity between processes that exhibit similarities in visual behaviour.

Moreover, functions may not require strict ordering (weak linear functional dependencies, Fig. 8, 3gen), potentially resulting in variations in work patterns (Fig. 8, 2imp). Additionally, the automation could analyse action points, which in Episode 1 gives richer information, allowing downwards inferences and projections forward on the event horizon. In particular, the landing clearance gives information about what the operator expects will unfold next (e.g., a landing, with a work pattern at level 3gen). Also, if the automation has information on the pre-requisites of a landing clearance, then it could infer that the operator should have conducted a runway check, level 2imp. Then it could inspect its operator data backwards in time and examine whether the runway check has been conducted; what objects it involved (level 1phy); and compare that to objects that its own sensors directed at the external process have registered (1phy). The automation could also continue monitoring of these objects for as long as the landing clearance was relevant.

9.7 Re-considering function allocation in man-machine teaming

Regarding function allocation, and the ‘extended Fitts list’ (EFL) approaches, including “levels of automation”, we suggest using the LACC as part of the process of describing ability, and the Score to describe function allocation variation and autonomy over time. This takes a contextual control perspective (what-is-good-here), rather than taking the perspective of who-is-good-at-what (Fitts list). From a Good-old Human Factors Perspective the decision points (DP) are what require situation awareness (SA). The Score shows how SA can be affected by PP, and at what levels of cognitive control decisions are made (indicating what kind of SA is required). The framing level (6) is particularly central from a sensemaking perspective, with the Score showing how it relates to control.

As a final note, the abstract score patterns (Fig. 12) is a continuation of what was shown in Fig. 5. Figure 5 shows (through the Score) the extremes of those two conceptual models (for active contextual control versus structure-mediated control respectively). Both of these models are particular abstractions that can be described using the more generic (Joint Control) Score framework. Figure 12 continues what we outlined in Fig. 6, which is that such abstractions of control can be based on empirical data. The

underlying models by Hollnagel and Rasmussen in Fig. 5 are now around 20–30 years old. To be able to derive new analytical approaches from empirical data, or from an interplay between sketching and testing new control models, will arguably form a sounder basis for taking CSE further to tackle current and upcoming challenges and to remain relevant with a “living” and growing set of tools, rather than freezing old abstractions in time (matching challenges that were relevant in the times at which they were “frozen”). In that sense the models in Fig. 5 represent two starting points, particular scores matching the challenges of yesterday, that can now become exemplars for renewing CSE by rethinking and applying these concepts in the current context. This is central to developing the ‘man-machine-teaming’ (MMT) approach further.

10 Conclusion

Comparing previous theoretical strands, we found that each by itself is incomplete (focusing too strongly on one facet, e.g., the control structure, or the active enforcement of constraints), and when taken together the result is a mess of slightly different concepts (Table 1), which stands in the way of working with the facets together. Current socio-technical control systems are characterized by humans and autonomous/automated systems working together to control one or more processes, either directly by acting on the process or indirectly by erecting structures/boundaries containing the process. This can include both humans monitoring automation, and vice versa. To understand critical episodes, our review and analysis of episodes suggests that it should be useful to describe what occurs regarding level of cognitive control, function allocation, and communication at the system joints, over time.

This paper has presented a framework that brings these strands together, to describe joint control in systems with humans and autonomous agents as well as more basic automation, regarding function allocation and interaction at the joints. The JCF-S notation can be used to analyse the joints of the joint control system over time, e.g., as episodes (see, e.g., the analysis of Episodes 1 and 2), to describe and understand:

- Control tactics, including feedback/feedforward
- The level of cognitive control at the joints in the human-machine system
- Joints between control processes
- (potentially dynamic) function allocations over episodes
- The levels at which the system is transparent versus being a black box, versus the need to use the transparency/open the black box in particular scenarios and interaction patterns

- The levels at which interactions should/need to take place as a consequence of particular function allocations

In analysing interaction between human operators and autonomous, automated, and manual control systems, the Joint Control Framework (Fig. 6) uses four analytical steps:

- Process Mapping (PM)
- Levels of Autonomy in Cognitive Control (LACC)
- Analysis of temporal descriptions of human–machine interaction (HMI-T),
- through a notation for describing joint control over time (JCF-S)

The output of these analytical steps is the Joint Control Framework Score (JCF-S). Firstly, *process mapping* (PM, Fig. 7) describes the core processes and control processes that are on-going. Secondly, the Levels of Autonomy in Cognitive Control (LACC) (Table 3, Fig. 8) describe interactions between control processes (humans, automation/autonomous systems) and core (controlled) processes. It describes the interactions in a control process that correspond to various levels of cognition (summarized in one framework) based on previous theories in CSE and Human Factors. Thirdly, *the (joint control) Score*, (Fig. 9) is used to describe joint control over time in dynamic systems, the HMI-T. It describes exchanges between external processes (that is, processes that are to be controlled) and control processes *over time*. Thus it can also complement structural analyses in CWA (i.e., hierarchical functional constraint trees using the AH) with temporal control episode analyses.

Interactions are described in the JCF-S by placing these exchanges at different levels representing delegation of control and degree of autonomy in the joint system. A contribution of the framework is that it identifies the control joints where interactions between controlling agents and controlled processes occur. It describes control independent of whether it is exercised by a human being or a technical system. In this way, the notion of “joint cognitive system” (JCS) suggested by Hollnagel and Woods (2005) remains intact, while allowing for a detailed analysis that explains exactly how and when the “jointness” of the JCS takes place.

JCF analysis also allows a nuanced description of transparent versus black-box episodes on the Score, regarding both the human side (to the automation) and the automation (to the human). The analysis can also show changes in automation-human roles during control episodes (e.g., due to adaptive automation). However, even though the underlying model is quite complex (see Figs. 6, 8), the Score is quite a simple notation.

Moreover, the Score allows for abstraction of particular episodes into episode types (Fig. 12), which can potentially

facilitate generalizability and transfer of analyses and solutions across cases and domains. Although shorter episodes can be analysed manually, analysis tools could be needed to cope with longer episodes, and with episodes that have many simultaneous control processes and external processes. Further, we have not measured how difficult or time demanding it would be to apply the suggested framework compared to other frameworks. Nor has the framework been validated in the sense of applying it in a completed design cycle. Currently, several projects have been initiated with the purpose of studying these issues.

Practitioner Summary. The Joint Control Framework, through the Score notation, can be used to describe and analyse interaction between human operators and autonomous, automated, and manual control systems. Existing work systems can be analysed, as can new designs. The Score allows for description of the level of cognitive control at the joints in the human–machine system, joints between control processes, and (potentially dynamic) function allocations over episodes of interaction. This allows analyses of control tactics, including feedback/feedforward; the levels at which the system is transparent versus being a black box, versus the need to use the transparency/open the black box in particular scenarios and in scenario types; function allocations; and analyses of questions regarding the levels at which interactions should/need to take place as a consequence of particular function allocations.

Acknowledgements Open access funding provided by Linköping University. The study presented above was funded by the Swedish Transport Administration, and the Air Navigation Services of Sweden.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A Detailed description of the JCF levels

Table 3 provides a comprehensive description of the six JCF levels. In Table 3, the levels are described and explained in even greater detail than in Table 2, splitting each of the six levels into a structural (function) and an active (plans) control facet (as in Figs. 6, 8). The table also exemplifies how each level defines the event horizon (represented by the Score).

Table 3 Levels of Autonomy in Cognitive Control (LACC): Structure, action, and event horizon facets

Levels of cognitive control	Level description
WHY 6 Frames	<p><i>Functions (as situations and contexts)</i>. This level abstracts away the system per se, maintaining a division between a “situation” and its “context”. This division is dependent on framing, which provides a way of understanding that results in a specific division of what should be considered part of the situation and what should be considered part of the context. It is a way of “seeing-as”, where, e.g., a rock (level 1) can be seen as a means for sitting—becoming a kind of chair, within a context of other objects and activities. The situation defines which process is seen as going on, and thus also which system may be appropriate to engage with it</p> <p><i>Plans (as the frames per se)</i>. At this overarching level, situations and problems are “framed” through different perspectives. Within a frame, a functional purpose for a system may crystallize (to use, deploy, or design) at the level below. Taking emergency response as an example, a fire is seen as something that should be contained and put out (rather than, for instance, something natural that should be allowed to roam free, or a hazard to run away from, or something to use to get rid of undesired items, or maybe to harvest energy from). In some cases, especially when dealing with the known and expected, framing will be straightforward. In other cases, it will take effort to make sense of the situation—especially if the situation initially looked familiar but took a surprising turn. In a broad sense, the “frame” evokes or incorporates all the levels below, e.g., any associated goals, plans, or behaviour patterns</p> <p><i>Event horizon</i>. Recognition of a situation that is associated with well-known developments can result in a well-specified event horizon (see level 3)</p>
5 Effects	<p><i>Of functions (Functional purpose)</i>. At this level, the functional purpose or motive of the system is described—why does it exist in the first place? What overarching function does the system of functions serve, what effect does it have as a whole on the environment? In the emergency response example, one such overarching function might be to keep people safe</p> <p><i>Of plans (Overarching goals)</i>. Overarching goals are the reasons for why activities as a whole are carried out, or what effect the activity should have, such as the moving of object X from point A to B</p> <p><i>Event horizon</i>. The effects level sets expectations for outcomes (effects on the process) on the event horizon. It can also set generic expectations of what goals will be pursued in case of (complete/partial) failure or success</p>
HOW 4 Values	<p><i>Of functions (Abstract functions—abilities, qualities tensions)</i>. Functional potential and deployed functions can be represented in terms of more abstract terms, such as “preparedness” versus “deployment”. Preparedness, a value function, can for instance be calculated (using an abstract mathematical function), and then visualized as different areas of colour on a map, highlighting the potential to deploy functions over that area, or as a number summarizing the deployment potential of an area. Other abstract qualities of functions at this level are, for instance, workload, situation awareness, uncertainty, or resilience (as abstract measures)</p> <p><i>Of plans (Criteria/Tradeoffs)</i>. Regarding plans, tradeoffs are considered at level 4. For instance, efficiency versus thoroughness. This level also concerns priorities (e.g., safety over efficiency), and criteria (e.g., desired levels of preparedness or risk)</p> <p><i>Event horizon</i>. Values refer on the one hand to a value ladder, which guides trading-off (if I cannot achieve, a, then prioritise b), and on the other hand to value contexts on the event horizon, and which measures are present (e.g., changing preparedness levels over the course of a crisis, or value measures associated with being at a particular place at a particular time)</p>
3 Generic	<p><i>Functions (types, typical states)</i>. At the second level, abstraction strips away the actual resources, and keeps only the function being performed. For instance, a coolant function, or a smoke diving function</p> <p><i>Plans (generic plans)</i>. Planning is performed at this level. The current state of the plan may also be monitored, at whatever step is currently executing. Regarding plans, this level may specify a procedure of steps to perform, describing how to carry out a task. At this level, the plan is generic and reusable in similar situations. Local adaptations and particular constraints (implementation of the plan) belongs to level 2</p> <p><i>Event horizon</i>. Having engaged in or predicted a generic level activity, sets up an event horizon of expected or planned developments</p>
WHAT 2 Implementations	<p><i>Of functions (Implemented functions, physical system state)</i>. This level focuses on the constraints that are due to the specific implementation. At this level, objects are considered to embody functions. A pipe, which performs the job of keeping water within limits, may here be seen as part of a coolant function. Fire fighters and their gear are described as embodying a smoke diving fire fighter unit. The current configuration of the system constitutes its physical state</p> <p><i>Of plans (adjustment and timing of constraints)</i>. At this level, constraints are set, with respect to the timing and position of setting constraints, such as when overtaking a car while driving an emergency response truck</p> <p><i>Event horizon</i>. Implementation decisions that can be made ahead of execution constitute a more specific event horizon than on the level above</p>
1 Physical	<p><i>Functions (physical resources)</i> At level 1 functions are described in terms of the physical resources through which they are embodied. Properties like size, weight, and physical location are central, such as the location of emergency response vehicles and personnel</p> <p><i>Plans (physical action)</i>: Turning from functions to the execution of plans, level 1 is concerned with physical action, such as setting or holding constraints within levels. However, the constraints are enforced through human involvement rather than through physical objects on their own— adjusting performance through feedback and feed-forward, to hold values within bounds</p> <p><i>Event horizon</i>. The ability to see objects ahead provides a situated basis for establishing an event horizon, which can also serve as a basis for implementation (at level 2)</p>

References

- Aminoff, H, Johansson, B, Trnka, J (2007) Understanding coordination in emergency response, Proceedings of the European Conference on Human Decision making and Manual Control (EAM) conference, Lyngby, Denmark
- Bainbridge L (1983) Ironies of automation. *Automatica* 19:775–779
- Bradshaw JM, Hoffman RR, Woods DD, Johnson M (2013) The seven deadly myths of "autonomous systems". *IEEE Intell Syst* 28:54–61
- Brehmer B, Allard R (1991) Real-time dynamic decision making. Effects of task complexity and feedback delays. In: Rasmussen J, Brehmer B, Leplat J (eds) *Distributed decision making: cognitive models for cooperative work*. Wiley, Chichester
- de Winter JCF, Dodou D (2011) Why the Fitts list has persisted throughout the history of function allocation. *Cogn Technol Work* 1:1–11
- Dekker SWA, Woods DD (2002) MABA-MABA or abracadabra? Progress on human-automation co-ordination. *Cogn Technol Work* 4:240–244
- Fay D, Stanton NA, Roberts A (2017) Designing new interfaces for submarines: from cognitive work analysis to ecological interface design. In: Stanton NA, Landry S, Di Bucchianico G, Vallicelli A (Eds.). *Advances in human aspects of transportation: proceedings of the AHFE 2016 international conference on human factors in transportation, July 27–31, 2016, Walt Disney World®, Florida, USA*. Springer International Publishing, Cham, pp. 413–425
- Fitts PM (1951) Human engineering for an effective air-navigation and traffic-control system. In: Fitts PM (ed) *Human engineering for an effective air-navigation and traffic-control system*, vol 84. National Research Council, Div. of, Oxford, pp xxii–84–xxii
- Hancock PA, Jagacinski RJ, Parasuraman R, Wickens CD, Wilson GF, Kaber DB (2013) Human-automation interaction research: past, present, and future. *Ergon Des* 21:9–14
- Hoffman RR, Johnson M, Bradshaw JM, Underbrink A (2013) Trust in automation. *IEEE Intell Syst* 28:84–88
- Hoffman RR, Hawley JK, Bradshaw JM (2014) Myths of automation, part 2: some very human consequences. *IEEE Intell Syst* 29:82–85
- Hollnagel E (1993) *Human reliability analysis: context and control*. Academic Press, London
- Hollnagel E (1998) Context, cognition and control. In: Waern Y (ed) *Co-operative process management: cognition and information technology*. Taylor and Francis, Bristol, pp 27–52
- Hollnagel E (2001) Extended cognition and the future of ergonomics. *Theor Issues Ergon Sci* 2:309–315
- Hollnagel E, Bye A (2000) Principles for modelling function allocation. *Int J Hum Comput Stud* 52:253–265
- Hollnagel E, Woods DD (1983) Cognitive systems engineering: new wine in new bottles. *Int J Man Mach Stud* 18:583–600
- Hollnagel E, Woods DA (2005) *Joint cognitive systems: foundations of cognitive systems engineering*. CRC Press, Boca Raton
- Johansson B, Lundberg J (2017) Resilience and the temporal dimension—the chimera of timely response. *Theor Issues Ergon Sci* 18:110–127
- Johansson B, Stenius C (2015) Navigation support using minimal information as a supplement to a digital map. *Int J Inf Syst Crisis Res Manage (IJISCRAM)* 7:61–79
- Johnson M, Bradshaw JM, Hoffman RR, Feltovich PJ, Woods DD (2014) Seven cardinal virtues of human-machine teamwork: examples from the DARPA robotic challenge. *IEEE Intell Syst* 29:74–80
- Klein G, Woods DD, Bradshaw JM, Hoffman RR, Feltovich PJ (2004) Ten challenges for making automation a "team player" in joint human-agent activity. *Intell Syst IEEE* 19:91–95
- Klein G, Moon B, Hoffman RR (2006) Making sense of sensemaking 2: a macrocognitive model. *IEEE Intell Syst* 21:88–92
- Klein G, Wiggins S, Dominguez CO (2010) Team sensemaking. *Theor Issues Ergon Sci* 11:304–320
- Lakoff G (2014) *The all new don't think of an elephant!: know your values and frame the debate*. Green Publishing, Chelsea
- Lundberg J (2015) Situation awareness systems, states and processes: a holistic framework. *Theor Issues Ergon Sci* 16:447–473
- Lundberg J, Törnqvist E, Nadjm-Tehrani S (2012) Resilience in sense-making and control of emergency response. *Int J Emerg Manage* 8:99–122
- Lundberg J, Törnqvist EK, Nadjm-Tehrani S (2014) Establishing conversation spaces in hastily formed networks: the worst fire in modern Swedish history. *Disasters* 38:790–807
- Lundberg J, Svensson Å, Johansson J, Josefsson B (2015) Human-automation Collaboration Strategies. In: Schaefer D, (Ed.). *Proceedings of the SESAR Innovation Days*. EUROCONTROL, ISSN 0770–1268, University of Bologna
- Lundberg J, Arvola M, Westin C, Holmlid S, Nordvall M, Josefsson B (2018) Cognitive work analysis in the conceptual design of first-of-a-kind systems—designing urban air traffic management. *Behav Inform Technol* 37:904–925
- Miller CA, Parasuraman R (2003) Beyond levels of automation: an architecture for more flexible human-automation collaboration. *Proc Hum Factors Ergon Soc Annual Meeting* 47:182–186
- Naikar N (2017) Cognitive work analysis: an influential legacy extending beyond human factors and engineering. *Appl Ergon* 59:528–540
- Neisser U (1976) *Cognition and reality: principles and implications of cognitive psychology*. W H Freeman and Company, San Francisco
- Nemeth C, Klein G (2011) The naturalistic decision making perspective. In: Cochran JJ, Cox LA, Keskinocak P, Kharoufeh JP, Smith JC (eds) *Wiley encyclopedia of operations research and management science*. Wiley, New York. <https://doi.org/10.1002/9780470400531.eorms0410>
- Parasuraman R, Bahri T, Deaton JE, Morrison JG, Barnes M (1992) Theory and design of adaptive automation in aviation systems. *Naval Air Warfare Center, Warminster*
- Parasuraman R, Sheridan TB, Wickens CD (2000) A model for types and levels of human interaction with automation. *IEEE Trans Syst Man Cybern Part A Syst Hum* 30:286–297
- Parasuraman R, Sheridan TB, Wickens CD (2008) Situation awareness, mental workload, and trust in automation: viable, empirically supported cognitive engineering constructs. *J Cogn Eng Dec Making* 2:140–160
- Rasmussen J (1986) *Information processing and human-machine interaction—an approach to cognitive engineering*. Elsevier, Amsterdam
- Salas E, Stout RJ, Janis AC-B, (1994) The Role of Shared Mental Models in Developing Shared Situational Awareness. In: Gilson RD, Garland DJ, Koonce JM (Eds.). *Situational Awareness in Complex Systems—Proceedings of a CAHFA Conference*. Embry-Riddle Aeronautical University Press, Daytona Beach. pp. 297–304
- Sheridan TB (1976) Toward a model of supervisory control. In: Sheridan TB, Johanssen G (eds) *Monitoring behavior and supervisory control*. Plenum press, New York, pp 271–281
- Sheridan TB (2011) Adaptive automation, level of automation, allocation authority, supervisory control, and adaptive control: distinctions and modes of adaptation. *IEEE Trans Syst Man Cybern Part A Syst Hum* 41:662–667
- Sheridan TB (2012) Human supervisory control. In: Salvendy G (ed) *Handbook of human factors and ergonomics*, 4th edn. Wiley, Hoboken, pp 990–1015

- Sheridan TB (2017) Modeling human-system interaction : philosophical and methodological considerations, with examples. Wiley, New York
- Sheridan TB, Verplank WL (1978) Human and computer control of undersea teleoperators (No ADA057655). Massachusetts Institute of Technology, Cambridge
- Sorensen LJ, Stanton NA, Banks AP (2011) Back to SA school: contrasting three approaches to situation awareness in the cockpit. *Theor Issues Ergon Sci* 12:451–471
- van Westrenen F, Praetorius G (2014) Situation awareness and maritime traffic: having awareness or being in control? *Theor Issues Ergon Sci* 15:161–180
- Vicente K, Rasmussen J (1992) Ecological interface design: theoretical foundations. *Syst Man Cybern IEEE Trans* 22:589–606
- Weick K (1993) The collapse of sensemaking in organizations: the Mann Gulch disaster. *Adm Sci Q* 38:628–652
- Westin C, Borst C, Hilburn B (2016) Automation transparency and personalized decision support: air traffic controller interaction with a resolution advisory system*. *IFAC-PapersOnLine* 49:201–206
- Wiener N (1948) *Cybernetics; or control and communication in the animal and the machine*. Wiley, Oxford

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.