26th Annual INCOSE International Symposium (IS 2016)
Edinburg, Scotland, UK, July 18-21, 2016

On the Automation of Intelligence, Sensing, and Reconnaissance Systems in Low-UDR Operations

Michael A. Schaffner
Sandia National Laboratories
505-284-8060
maschaf@sandia.gov

Abstract. Currently deployed Intelligence, Sensing, and Reconnaissance (ISR) systems fall along the spectrum of autonomy in operations, ranging from fully autonomous systems that collect and analyze data to those systems that require humans as integral to effective collection operations. This paper outlines two traits of collection systems that require relatively less human involvement in operations, introducing the measurement of an operation's *Understanding-to-Data Ratio* (UDR) to help determine the appropriateness of automating a system in a given operating context. Human and automated capabilities in existing systems are contrasted, with the strengths of each identified. Finally, a spectrum of system capability and human involvement is proposed for considering tradeoffs involved with new acquisitions in the ISR domain, and the key areas for design effort in automation are outlined, including a focus on the interface between human operators and automated systems.

Introduction

Intelligence, Sensing, and Reconnaissance (ISR) systems have been used since the beginnings of warfare, with humans traditionally performing all activities in the Tasking, Collection, Processing, Exploitation, and Dissemination (TCPED) chain. Throughout the twentieth century, the Collection and Processing stages began to involve a greater number of electronic systems, including wired communication systems, optical systems, radio, radar, infrared and multi-spectral systems, digital communication systems, and eventually satellite-based, sea-based, and air-based capabilities of all kinds. Some of these collection systems were simply an interface from a real-time human operator to an electronic medium (e.g., wiretaps), while other collection systems required little human activity or intervention beyond monitoring state of health and occasional updates to tasking (e.g., satellite-based film photography). With today's more capable collection systems collecting more data than ever before, with more demand for ISR systems (www.gao.gov), and with an increasing demand for solutions requiring less manpower (www.rand.org), a desire for more automated solutions is a natural response. This idea is often couched in terms of traversing left on the spectrum shown in Figure 1 below.

Figure 1. A simple spectrum implicitly used in discussions of system automation.

This paper approaches the need for "traversing left" by examining the following questions: Do certain system attributes necessitate lower human involvement in electronic collection systems? Are there common traits of those systems requiring less manpower in operations? In answering these questions, the present work addresses two traits that have enabled automated systems to succeed in certain modern environments: 1) Limited real-time feedback needed for effective collection; and 2) A high understanding-to-data ratio (UDR) of operation – in some tightly-defined contexts, large amounts of situational understanding can be derived from very specific data collects. Each of these attributes will now be examined vis-à-vis its contribution to reducing and/or eliminating human responsibilities in the operations of ISR systems.

Trait 1: Limited real-time feedback needed for effective collection.

A collection system's ability to adapt to real-time events in operation is a primary driver of the level of human involvement required for effective collection. The need for this ability can stem from a collection system's mobility as well as the desired level of dynamic tasking for a collection system. Mobile systems such as ground and air vehicles may need to respond to real-world events more quickly than a stationary radio antenna or a satellite with a fixed-pointing payload. The less mobile a collection system is, though, the less necessary it is to have real-time feedback in operations for effective collection. In addition, a UAV may be tasked with collecting pictures or video of stationary and/or moving targets, with tasking subject to change at a moment's notice depending on external needs or in-scene developments (i.e., dynamic tasking). A radio antenna, however, may only ever be tasked with receiving and recording certain frequencies around the clock (i.e., fixed tasking). Less dynamic tasking necessitates less real-time feedback in operations for effective collection.

The more mobile a system is, and/or the more its tasking is intended to adapt to developing situations, the more essential human operators' real-time feedback becomes. As Lowenthal writes, "…Technical collection is less than precise. The problem underscores the importance of processing and exploitation" (Lowenthal 2014, p. 71). Real-time (albeit partial) processing and exploitation by human operators ensures that dynamic collection systems operate collect the right signals in the right way at the right time. This incorporation of real-time feedback into a collection system precludes wholesale replacement of humans by automation, since real-time feedback in open environments – especially from exploitation, a particularly creative problem-solving activity – is a uniquely human skill. With respect to this adaptable ability, Sterman nicely summarizes the differences between (human) mental models and (automated) computer models: "A mental model is flexible; it can take into account a wider range of information than just numerical data; it can be adapted to new situations and be modified as new information becomes available...[Computer models] are unable to deal with relationships and factors that are difficult to quantify, for which numerical data do not exist, or that lie outside the expertise of the specialists who built the model" (Sterman 1991).

Thus the incorporation of real-time feedback into a collection system's tasking and operations requires efficient and effective real-time processing and exploitation of collected data, which in turn requires humans to be "in the loop", with a *high level of situational awareness*. Lowenthal comments on the necessity of analysts' judgments in identifying more valuable collections: "…the analysts' expertise should be an integral part of collection sorting" (Lowenthal 2014, p. 114). The necessity of humans for processing and exploitation activities brings into view the second trait of interest to the present analysis, which is the understanding-to-data ratio (UDR) of an operation.

Trait 2: High Understanding-to-Data Ratio (UDR)

A second trait of ISR systems requiring less human involvement is the high Understanding-to-Data Ratio (UDR) of an operation with respect to the system's collection ability and its operating environment. The concept of UDR is analogous to the measure of Signal-to-Noise Ratio (SNR) commonly used in engineered systems such as antennas and visual detectors. The amount of *understanding* (i.e., the Signal), for the purposes of the present concept, is proportional to the fidelity of high-level situational questions being answered by an ISR system in operation (vague or imprecise ideas imply a lower understanding of a problem, while narrowly defined questions and targets imply a high degree of understanding). The *data* (i.e., the Noise) is the low-level (usually electronic) collections from that same ISR system. UDR is then the measure of how much the data can be extrapolated to answer a high-level question related to situational understanding. In "From Data to Wisdom", Ackoff defines a hierarchy of human sensing and thought which relates the concepts of data, information, knowledge, understanding, and wisdom (Ackoff 1989). This hierarchy is shown in Figure 2 below.

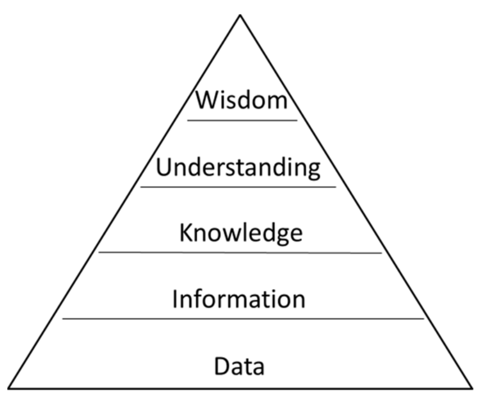

Figure 2. The heirarchy of human sensing and thought proposed by Ackoff (1989).

Each ascending level of the hierarchy distills the level below it, resulting in general, top-level models (i.e., understanding and wisdom) which humans use to make sense of and predict events in the world at large (note how the idea of distillation and integration parallels the CPE activities of the TCPED chain). This process happens somewhat automatically, albeit supplemented by deliberate questioning and reflecting. The amount of understanding one has about a given problem or situation often simply comes through direct experience and observation (i.e., data collection). Occasionally, distilled high-level models (i.e., understanding and wisdom) can extrapolate from a seemingly small amount of data, which is indicative of a high Understanding-to-Data Ratio. In other words, in a high-UDR operation, one's level of understanding about a given situation is relatively high, while the data required to make immediate judgments is relatively low.

Unlike SNR's quantitative measures of both Signal and Noise, UDR as a measure is somewhat limited by one's ability to define and measure the level of one's *understanding*. Quantitative or even qualitative measures of understanding are not plentiful, since the upper levels of Ackoff's heirarchy reflect the (mostly unobservable) states, processes, and relationships of the human psyche. Lonergan describes understanding as "not the mere apprehension of any [data], not the mere memory of all, but a quite distinct activity of organizing intelligence that places the full set of clues in a unique explanatory perspective" (1992). When we write about understanding or insight, he says, "we write about a moving target, from a moving viewpoint" (1992). Hence the inherent difficulty in definition and measurement. With respect to ISR systems, though, there exist (or should exist) some qualitative ideas of what *type of understanding* a system is intended to provide – whether specific occurrence of a singular event in a specific context, or broad situational awareness in various contexts.

A key consideration of this point is that narrower goals of understanding (i.e., specific questions about a limited number of objects or events) imply a higher UDR operation, since narrower goals imply a clearer, more tightly scoped definition of the "problem" motivating the system design in the first place. A clear grasp of a simpler problem allows one to glean answers from specific data collects due to basic assumptions that allow extrapolation of the data's implications. On the other hand, a more general problem – for example, explanation of human behavior over time – prevents simple assumptions from being used (or very reliable), meaning a lower understanding of the basic problem, and thus a lower UDR operation. This seemingly small amount of distinction between types of understanding can be useful enough in distinguishing relative levels of UDR operations, as following examples demonstrate.

The divisor of the UDR measure, *data*, is relatively easily described numerically, having multiple established quantitative measures. One such measure is "richness", in which the number of types of data observed are known, but the relative abundance of each is not known (Gotelli and Colwell 2001). Another is "diversity", in which the number of types of data observed and their relative abundance are both known (Peet 1974, Jost 2006). And closely related to this latter measure is "entropy", one of the foundations of information theory (Shannon 1948). These measurements can all be used to inform and compare the UDR measurements of different systems, since systems are generally designed around a limited range of data collection types.

In high UDR operations, questions at the upper-levels of Ackoff's hierarchy can be almost completely informed by relatively specific data collects. Examples of these types of high UDR systems include automated collection and analysis systems such as those dealing with radar detection, internet traffic, and phone records. The collection systems in each of these areas expect limited signal types, whether only radio frequencies, IP addresses, or phone numbers and metadata. In addition, the analysis (i.e., processing and exploitation) is fairly straightforward, since the data collected is directly related to the analysis desired: radar signatures, internet traffic, and phone records can all be effectively described by only a handful of quantifiable parameters (i.e., relatively low *richness* of data being collected, and in some cases relatively low *entropy* as well). Equally important, each of these is collected in a narrowly-defined operating context, along with very specific goals of identification (narrowly defined *understanding*), which serve as a kind of pre-formed high-level model into which the low-level data can be extrapolated. In summary, automated systems such as these generally operate in relatively barren datascapes – ignoring or unable to collect different types of data – and with very well-defined (and pre-defined) targets in the environment of signals collected. These are some of the basic characteristics of high UDR operations, shown in Table 1 below.

Traits of High UDR Operations	**Traits of Low UDR Operations**
-Limited set of questions and targets -Clearly defined questions and targets -Many assumptions about implications of data -Designed to conclude -Limited, specific data collects -Limited, specific operating contexts	-Unbounded set of questions and targets -Questions and targets ill-defined in advance -Assumptions are questioned -Exploratory and explanatory -Big data (i.e., streaming HD video) -Variety of operating contexts

Table 1. High UDR and Low UDR operations contrasted.

Modern weapon systems also incorporate high UDR operations. In the Center for a New American Security's report "Introduction to Autonomy in Weapon Systems", Scharre and Horowitz review various weapons systems with respect to which functions require humans to be "in the loop", "on the loop", and "out of the loop" (www.cnas.org). While the majority of weapon technology deployed today requires humans to be "in the loop", some systems do operate in highly automated modes with humans simply "on the loop". But according to Scharre and Horowitz: "To date, these [latter modes] have been used for defensive situations where the reaction time required for engagement is so short that it would be physically impossible for humans to remain 'in the loop' and take positive action before each engagement and still defend effectively…In all of these cases, automation is used to defend human-occupied bases or vehicles from being overwhelmed by a rapid barrage of missiles or rockets" (www.cnas.org).

Thus, the only demonstrated military cases of humans "on the loop" today are very short-duration, narrowly-defined situations with well-defined and predictable targets, in which there are no other options (since a delay on the order of human reaction time may be an existential matter)[1]. Based on a few radar pings, these systems can effectively launch missile interceptors because the narrowly defined context of operation allows the lowest-level *data* of radar detection to greatly inform certain higher-level models of the hierarchy – even up to essential *understanding* of an existential threat. These radar-based systems with predictable targets in very limited operating contexts are prime examples of collection and analysis systems in a high UDR operation.

As an example of an ISR system in a low UDR operation, consider a full-motion video sensor mounted on a UAV meant to conduct surveillance in geographical regions around the world. The UAV collects a large amount of image data (perhaps 30 frames per second at a given resolution of pixel values), and may have other types of sensors on board as well. The amount and variety of low-level *data* being collected (i.e., pixel values and other electronic signals) immediately lowers the UDR of this operation compared to previous examples. With regard to the *understanding* desired, very diverse information and knowledge – namely, identification of almost any type of physical object or spatio-temporal event over time – are intended to be gleaned from the system in order to provide a high degree of situational awareness (i.e., broad, virtually unbounded understanding of the world at large). Due to this diversity desired from collections as well as the diversity of operating contexts, target information is not well-formed, and often cannot even be known in advance. Absent are the simple rules for extrapolating data to answer higher-level questions, unlike the case of missile interceptors or phone collection records. (In addition, all kinds of confounding technical factors can be present when attempting to infer knowledge and understanding from this system, including weather conditions that affect sensor signals and pixel patterns, sensor movement due to turbulence and platform movement, changing lines of sight, perspective geometries, partially or mostly obscured objects, and many other factors.) This inability to define (or even bound) the knowledge and understanding desired leads to an even lower UDR as compared to previous examples.

In this type of low UDR collections operation, humans are essential components of the system due to their ability to operate effectively in such environments and due to the inability of automated systems to cope with factors that ultimately remain unquantifiable and therefore un-computable. Sterman's description of mental models again explains why humans excel in this type of low UDR collection operation (Sterman 1991), while Franchi and Guzeldere, writing about state-of-the-art artificial intelligence (AI) methods, are helpful in explaining the inadequacies of automation to effectively operate this type of system: "[Artificial intelligence] is…heuristic search in a search space game-theoretically defined…explained in terms of

---

[1] With respect to these systems, Allenby and Sarewitz (2011) ask the very appropriate question, "Does the word 'robot' signify a type of artifact, a type of capability, or a certain level of computational competence?" One could ask the same question about the more general word 'automation'.

satisficing a set of rigid constraints by searching heuristically the space that those constraints define" (Franchi/Guzeldere, p. 55). If a problem cannot be completely defined in game-theoretic terms[2], if the constraints on a problem are vague, unknown, or unquantifiably uncertain, then the heuristic search becomes something altogether different than the state-of-the-art artificial intelligence methods. Humans, on the other hand, easily conduct their own natural, quite effective forms of "heuristic search" to process big data in the form of "full motion, high-definition video" (i.e., human vision) because their *situational awareness* allows them to continually classify lower-level features observed as well as project the currently perceived situation into the future, through synthesizing their basic knowledge about the world, human behavior, and many other unquantifiable factors. In low UDR operations such as vision-based surveillance, this ability means that humans' higher-level mental models can help them selectively ignore most of the data and focus on only that most relevant (i.e., critical cues) to making sense of a visual scene and *understanding* the scene's implications (Endsley 1997).

While humans have spent most of their existence dealing with causal relationships and factors that remain unquantified (many of which are arguably unquantifiable), computers and automated systems are simply unable to deal with these most basic features of the world at large. Recent advances in artificial intelligence (AI) have largely been about automating machine actions in high UDR environments (e.g., speech recognition on mobile devices, handwriting recognition in banking systems, robotic systems on factory floors), with a view toward eventually automating machine actions in low UDR environments (e.g., autonomous vehicles, household robotic systems). But efforts to automate machine actions in low UDR environments cannot be approached as a strictly technical problem that is detached from human analysis. Franchi and Guzeldere write: "'Creating intelligence' as an engineering project makes it difficult to appreciate the complicated nature of human mental life, behavior, culture, and social practices – a territory generally studied and much better understood by the humanities and social sciences" (Franchi and Guzeldere, p. 18). These "complicated" aspects of life are precisely the information and factors necessary to effective ISR operation and effective intelligence analysis, since ISR systems have always been purposed with uncovering *human* actions and intents, in varying cultures with varying social practices. Even in systems "simply" tasked with operating safely in human-intensive environments, these "complicated" human matters prove to be a major challenge (www.nytimes.com). In systems that are tasked with not simply operating, but observing human actions and making sense of them in open-ended environments, the current state of AI – and its state for the foreseeable future – cannot begin to address these problems effectively[3].

---

[2] Indeed, often in engineered systems it is the case that the perception of the problem motivating a system's creation and operation are not always well known in advance, nor do they remain constant throughout the life of the system (Rhodes & Ross 2010, Ricci et al 2014). This is almost certainly the case with any ISR system today.

[3] Many computer scientists would obviously disagree, making grandiose claims of imminent leaps in artificial intelligence within the next 10-20 years. Such claims today are, of course, substantially no different than those of Simon, Newell, and others over the past 40-50 years with respect to digital computers, "thinking machines", and

Human involvement in low UDR collection environments has additional benefits beyond increasing the local system's performance. In addition to providing broad and synthesized situational awareness in novel and ill-defined environments, humans add tremendous value by collaborating with the wider community of other ISR systems in concurrent operation. It is often the case that one sensor system perceives only part of a developing real-time situation, while other systems can help fill in the gaps of understanding to form a more complete picture. On this topic, Ganter comments, "Situational landscapes emit contradictory evidence in different ways at different times, so the work requires interactive maneuvers of different sensors and thus interactive negotiations of different tribes" (Ganter 2007). Such interactive maneuvers, and especially interactive negotiations, require human operators due to the inherent vagueness and unquantifiable uncertainties encountered in dynamically evolving situations, as well as due to the communicative and political skills required. Though somewhat messy in practice, these cooperative activities can enable ISR systems to both harness and provide more diverse information from which to make sense of an unfolding situation than a single system operating in isolation.

The Tradeoff: ***A Spectrum of Three Tightly-Correlated Dimensions***

The observations above regarding the benefits of human involvement in operationally adaptive and situationally aware ISR systems lead to the conclusion that these three dimensions of ISR systems are inextricably linked. The ability of a system to incorporate and provide situational understanding is tightly linked to its operational adaptability, as well as to its incorporation of human operators as processing and cooperating agents. These three relationships are depicted in the expanded spectrum shown below in Figure 3.

Figure 3. The proposed spectrum of ISR system operations.

Each of the three dimensions on the spectrum can be considered separately for new system acquisitions, but each of the dimensions is tightly correlated with the others. A reduction of human-in-the-loop roles simultaneously reduces the potential adaptiveness and situational awareness of an ISR system. This framework can be used to help make explicit the capability tradeoffs involved with any newly proposed system, whether it is envisioned to be a fully autonomous system, a human-operated system, or something in between. It can also be used to

"General Purpose Solvers" (Simon and Newell, 1958; Simon 1978). For a more complete treatment on the history of failed predictions about AI, see chapter 1 of (Franchi and Guzeldere 2005).

better consider the appropriateness of automation in the various operating environments envisioned for a system.

At first glance, "traversing left" (i.e., increasing autonomy) on the spectrum might seem to be reasonably accomplished by slowly incrementing autonomous behavior step-by-step into new systems, while simultaneously decrementing human involvement required. To successfully do this, however, the operating contexts and goals must also be incrementally bounded, by deconstructing broad surveillance goals and/or operating environments into specific, narrowly-scoped contexts and corresponding questions that can be answered by extrapolating from low-level data collects. While this latter activity likely sounds straightforward, it most assuredly is not. In fact, it displays many traits of so-called "wicked" problems, including "no definitive formulation", "solutions are not true-or-false, but good-or-bad", "there is no immediate or ultimate test of a solution", and "innumerable set of potential solutions", among others (Rittel and Webber 1973). As previously noted, human understanding is difficult to write about, much less formalize in an agreed-upon way (which is one reason that measuring the performance of broad ISR systems is a nontrivial task). If increasing autonomy is desired, however, the goals of operation must be well-understood and tightly bounded, keeping in mind that this may mean separate systems for separate purposes.

Another challenge in "traversing left", one that is perhaps counterintuitive, is that partial automation can lead to several new types of dangers (Inagaki 2011, citing Woods 1989, Wickens 1994, Endsley and Kiris 1995, Sarter and Woods 1995, Parasuraman and Riley 1997, and Sarter et al 1997). For systems intended to reduce human-in-the-loop activities during operation, more design effort will need to be spent on human factors considerations, specifically with regard to enabling human operators to stay aware of newly automated portions of the system. As Rechtin and Maier (1991) state, "The greatest leverage in system architecting is at the interfaces. The greatest dangers are also at the interfaces." The obviousness of this statement with respect to the interface between autonomous and human actions in a complex system is striking. On human and automated processing and their interface, Ware (2008) comments, "It is useful to think of the human and computer together as a single cognitive entity, with the computer functioning as a kind of cognitive co-processor to the brain…Each part of the system is doing what it does best. The computer can pre-process vast amounts of information. The human can do rapid pattern analysis and flexible decision making." Scoping these activities early on in the conceptual design and requirements definition phases of a new system can provide unambiguous guidelines for later interface design and lower-level implementation activities, ensuring that the systems' eventual operations will be leveraging the strengths of these fundamentally different types of processing.

Conclusion

As demand increases for new ISR solutions that require less manpower, the tradeoffs inherent in such a design decision must be made explicit in order for stakeholders to understand how capability will be affected. The need for manpower can be reduced – and in some cases even eliminated – as has been demonstrated in limited historical and present-day examples. However, with this reduction comes a direct reduction in operational adaptability as well as a reduction in the situational awareness provided by the system. The goal for ISR systems cannot always be "to automate the human out of the loop" – especially those systems with real-time tasking, diverse informational sensors, and low UDR operations. And with systems that cannot be fully automated, Scharre and Horowitz state, "The key place to focus attention is which tasks are being automated and which does the human retain" (www.cnas.org). Only by focusing on this most critical scoping issue – and on the accompanying interface between automated and human components that keeps situational awareness for human operators – can new ISR systems truly deliver value to stakeholders through effective operations and comprehensive situational awareness.

Acknowledgements

The author is grateful to Sandia National Laboratories for partially funding this research, and to John Ganter and Kerstan Cole for many thought-provoking discussions and preliminary feedback on these ideas.

References

1. http://www.gao.gov/htext/d08374.html. Accessed April 20th, 2015.
2. http://www.rand.org/content/dam/rand/pubs/technical_reports/2008/RAND_TR459.pdf. Accessed April 20th, 2015.
3. Lowenthal, M. M. (2014). Intelligence: From secrets to policy. CQ press.
4. Sterman, J. D. (1991). A skeptic's guide to computer models. Managing a nation: The microcomputer software catalog, 2, 209-229.
5. Shannon, C. E. (1948). A note on the concept of entropy. *Bell System Tech. J*, *27*, 379-423.
6. Peet, R. K. (1974). The measurement of species diversity. Annual review of ecology and systematics, 285-307.
7. Gotelli, N. J., & Colwell, R. K. (2001). Quantifying biodiversity: procedures and pitfalls in the measurement and comparison of species richness. Ecology letters, 4(4), 379-391.
8. Jost, L. (2006). Entropy and diversity. Oikos, 113(2), 363-375.
9. Lonergan, B. (1992). *Insight: A study of human understanding* (Vol. 3). University of Toronto Press.
10. http://www.cnas.org/sites/default/files/publications-pdf/Ethical%20Autonomy%20Working%20Paper_021015_v02.pdf. Accessed June 4, 2015.
11. Endsley, M. R. (1997). The role of situation awareness in naturalistic decision making.
12. Ganter, J. H. (2007, June). Naturally Aligned: A Psycholinguistic Study of Shared Situation Awareness Among Military Sensor Operations Teams. In Proceedings of the

Eighth International Naturalistic Decision Making (NDM) Conference, K. Mosier & U. Fischer, eds., Pacific Grove, CA.

13. Ackoff, R. L., "From Data to Wisdom", Journal of Applies Systems Analysis, Volume 16, 1989 p 3-9.
14. Ware, C. (2010). Visual thinking: For design. Morgan Kaufmann.
15. http://www.nytimes.com/2015/09/02/technology/personaltech/google-says-its-not-the-driverless-cars-fault-its-other-drivers.html?_r=0. Accessed September 12, 2015.
16. Rechtin, E., & Maier, M. W. (2010). The art of systems architecting. CRC Press.
17. Franchi, S., & Güzeldere, G. (2005). Mechanical bodies, computational minds: artificial intelligence from automata to cyborgs. MIT Press.
18. Ricci, N., Schaffner, M. A., Ross, A. M., Rhodes, D. H., & Fitzgerald, M. E. (2014). Exploring stakeholder value models via interactive visualization. *Procedia Computer Science*, *28*, 294-303.
19. Rhodes, D. H., & Ross, A. M. (2010, April). Five aspects of engineering complex systems emerging constructs and methods. In *Systems Conference, 2010 4th Annual IEEE* (pp. 190-195). IEEE.
20. Allenby, B. R., & Sarewitz, D. (2011). *The techno-human condition*. MIT Press.
21. Herbert, Simon A. "The science of the artificial." (1978).
22. Simon, Herbert A., and Allen Newell. "Heuristic problem solving: The next advance in operations research." *Operations research* 6.1 (1958): 1-10.
23. Rittel, H. W., & Webber, M. M. (1973). Dilemmas in a general theory of planning. Policy sciences, 4(2), 155-169.
24. Inagaki, T. (2011). Technological and legal considerations for the design of interactions between human driver and advanced driver assistance systems. In *Proc. NeTWork Workshop: Control and Accountability in Highly Automated Systems*.
25. Wickens, C.D. (1994). Designing for situation awareness and trust in automation. In: Proceedings of IFAC integrated systems engineering, pp 77-8.
26. Woods, D. (1989). The effects of automation on human’s role: Experience from non-aviation industries. In Flight deck automation: Promises and realities (NASA CR-10036, pp.61-85)
27. Endsley, M.R., Kiris, E.O. (1995). The out-of-the-loop performance problem and the level of control in automation. Human Factors, 37(2):3181-319.
28. Sarter, N.B., & Woods, D.D. (1995). How in the world did we ever get into that mode? Mode error and awareness in supervisory control. Human Factors, 37(1):5-19.
29. Parasuraman, R., & Riley, V. (1997). Humans and automation: use, misuse, disuse, abuse. Human Factors, 39(2):230-253.
30. Sarter, N.B., Woods, D.D., & Billings, C.E. (1997). Automation surprises. In: Salvendy, G. (ed). Handbook of human factors and ergonomics, 2nd edn. Wiley, New York, pp 1926-1943.