Emerging challenges in cognitive ergonomics: managing swarms of self-organizing agent-based automation

JOHN D. LEE*

University of Iowa, Iowa City, IA 52242, USA

Keywords: Emergent behaviour; Information architecture; function allocation.

As technology becomes more sophisticated, the problems of appropriate function allocation, mode errors, and misuse of automation will continue to challenge system safety and efficiency. Addressing these problems will require the field of cognitive ergonomics to consider three important challenges. First, to understand the human implications of self-organizing, multi-agent automation may involve recognizing the unique monitoring and control requirements. While current research has studied how people control a small number (2–10) of agents, the future will likely introduce the challenge of supervising hundreds of agents. Multi-agent automation that consists of hundreds of loosely connected intelligent agents may exhibit powerful new adaptive behaviours that may be difficult for people to understand and manage. Secondly, to understand human interaction with increasing complex automation may require more comprehensive analysis and modelling techniques. Current analysis techniques such as analysis of variance, tend to rely upon static representations of the human–system interaction when dynamic representations are needed. Thirdly, understanding human interaction with this increasingly complex automation may benefit from reconsidering new constructs to explain behaviour. The constructs of the information processing approach may not be sufficient to explain reliance on multi-agent automation. Addressing the challenge of this new technology will require a theoretical understanding of human behaviour that goes beyond a task-based description of well-defined scenarios. Cognitive ergonomics must develop an understanding of the basic cognitive demands associated with managing multi-agent automation, tools that consider the dynamics of the interaction, and constructs that address the dynamic decision making that governs reliance.

1. Introduction

Until recently, automation has supported a few well-defined functions that do not interact with many other system components. As automation continues to evolve, it will generate important challenges for cognitive ergonomics. For example, technological advancement may enable automation to evolve into a swarm of interacting agents that may become increasingly powerful and autonomous. This multi-agent automation offers increased robustness, flexibility, and adaptability; however, understanding how to support the human supervisory control of swarms of agents remains an unresolved issue. Many researchers in computer science and robotics have realized the important capabilities of multi-agent automation, and further development of this technology seems likely (Beni and Wang 1993, Brooks and Flynn 1993, Fukuda *et al.* 1998). This paper addresses three emerging cognitive ergonomics issues associated with increasingly complex automation. First, cognitive ergonomics must

* e-mail: jdlee@engineering.uiowa.edu

Theoretical Issues in Ergonomics Science ISSN 1463–922X print/ISSN 1464–536X online © 2001 Taylor & Francis Ltd
http://www.tandf.co.uk/journals
DOI 10.1080/1463922011010492 5

identify the particular cognitive demands associated with the emergent properties of multi-agent automation that may make it difficult to understand. Secondly, to understand these cognitive demands, cognitive ergonomics must develop more comprehensive analysis and modelling techniques. Thirdly, cognitive ergonomics might benefit by rethinking the constructs underlying the dynamic decision making associated with reliance on multi-agent automation. The objective of this paper is to outline considerations and initial directions associated with each of these emerging issues.

2. Agent-based automation and new cognitive demands for supervisory control

Early analysis of human interaction with automation considered technology that replaced human operators in performing well-defined functions. Functions were statically allocated to humans or automation based on the inherent capabilities of each (Fitts 1951). With *static function allocation*, the division of labour between human and automation is fixed by the designer, and functions once performed by the human are now performed by the automation. Subsequently, more subtle forms of automation have evolved and researchers have described multiple levels of automation (Sheridan 1987), different types of automation (Lee and Sanquist 2000), combinations of types and levels (Parasuraman *et al.* 2000), and the dynamic allocation of function. With *dynamic function allocation*, the division of labour between human and automation depends on the moment-to-moment inclination of the human or the automation, each intervening or delegating functions (Hancock and Scallen 1996, Sarter and Woods 1997). Issues of authority and independence that have emerged with dynamic function allocation have become increasingly important (Sarter and Woods 1994). Agent-based automation requires a new description of the interaction between automation and the human supervisor because swarms of agents adapt to the environment in unpredictable ways. This may lead to an *emergent function allocation*, where the division of labour between the human and automation depends on capabilities that emerge as the human and automation adapt to a dynamic environment. Unlike current automation, future automation may not have a static capability that operators dynamically allocate to fulfil specific system functions. Instead, its capability may evolve as it interacts and adapts to the environment. Automation that consists of many agents may require fundamentally different types of control, such as the dynamic manipulation of agent autonomy and authority, as well as indirectly guiding the emergent behaviour of swarms of agents. The emergent behaviour of the swarm may extend human capabilities far beyond more conventional automation, but it may introduce important demands on the operator. Figure 1 summarizes the change in display and control interactions as technology moves toward agent-based automation.

2.1. *A few functionally discrete agents*

This represents the situation in many current systems, where operators must supervise and dynamically allocate function with several separate elements of automation that have not been assembled into an integrated system. Automation fills a relatively simple role of augmenting the humans' perception and control. The operator requests a specific behaviour and the automation responds in a deterministic manner. In this situation, significant human performance issues emerge as the operator is forced to bridge the gaps between functionally isolated automation (Bainbridge 1983, Lee and Sanquist 1996).

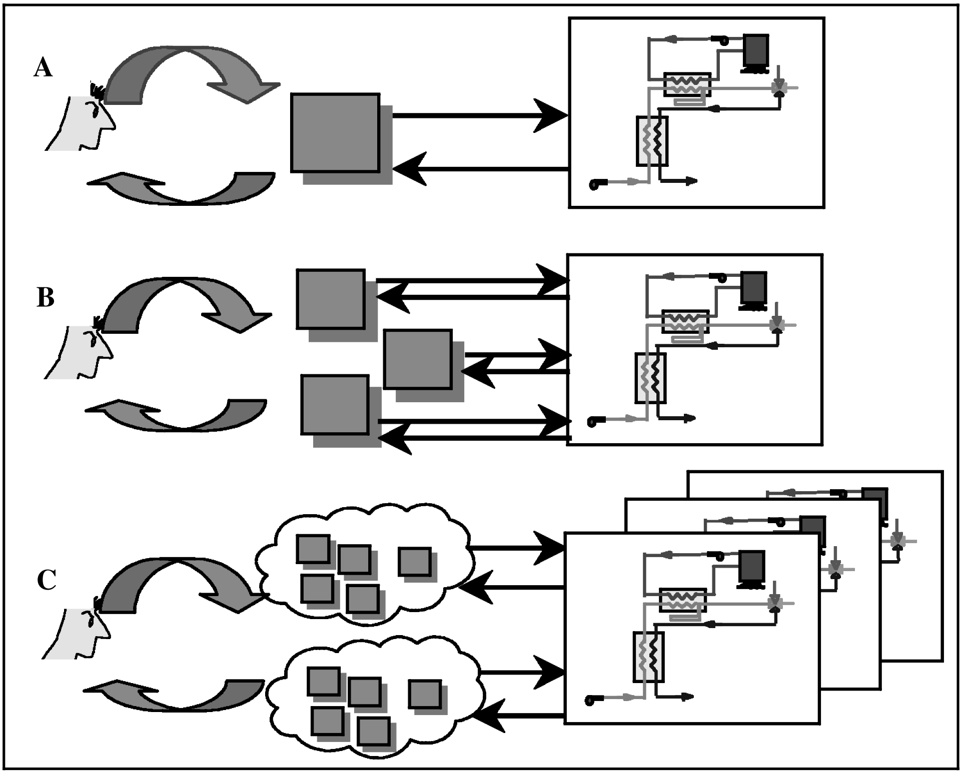

Figure 1. Possible supervisory control situations associated with increasingly complex automation.

2.2. *Several uncoordinated and loosely coupled agents*
This situation shows a more sophisticated sort of automation, which changes modes in response to the human operator, to other elements of automation, and to the environment. In this situation, operators must monitor multiple elements of the automation to ensure that their joint behaviour is productive. This sort of automation can greatly extend operator capabilities; however, mode errors illustrate the potential for inducing errors and degrading system safety and performance (Degani *et al.* 1995, Sarter and Woods 1995).

2.3. *Swarms of self-organizing agents*
This situation shows a qualitative change in supervisory control associated with an increasing number of agents. The many interacting agents may induce macro-level behaviour that cannot be easily predicted by the behaviour of individual agents. Although this added complexity can undermine the operator's monitoring efficiency, it supports a much more powerful and adaptive system. The operator is now responsible for managing the macro-level, emergent behaviour of a swarm of agents that interact in a complex manner, while also monitoring the micro-level behaviour of individual agents.

Multi-agent automation is an alternate design paradigm that may make it possible to respond to environmental variability while reducing the chance of system failure. These capabilities have important applications in a wide range of domains including planetary exploration, landmine neutralization, or even data exploration, where hundreds of simple agents might be more effective than a single complex agent. Biology-inspired roboticists provide a specific example of agent-based automation. Instead of the traditional approach of relying on one or two larger robots, they employ swarms of insect robots as an alternative (Brooks *et al.* 1990, Johnson and Bay 1995). The swarm robot concept assumes that small machines with simple reactive behaviours can perform important functions more reliably and with lower power and mass requirements than can larger robots (Beni and Wang 1993, Brooks and Flynn 1993, Fukuda *et al.* 1998). Typically, the simple programs running on the insect robot are designed to elicit desirable emergent behaviours in the insect swarm (Sugihara and Suzuki 1990, Min and Yin 1998). For example, a large group of small robots might be programmed to search for concentrations of particular mineral deposits by building upon the foraging algorithms of honeybees or ants.

In addition to physical examples of multi-agent automation, agent-based automation has a huge potential in searching large complex data sets for useful information. For example, the pervasive issue of data overload and the difficulties associated with effective information retrieval suggest a particularly useful application of multi-agent automation. Current approaches to searching large complex data sources, such as the Internet, are ineffective. People are likely to miss important documents, disregard data that is a significant departure from initial assumptions, misinterpret data that corroborates or conflicts with an emerging understanding, and disregard more recent data that could revise interpretation (Patterson 1999). These issues can be summarized as the need to *broaden searches* to enhance opportunity to discover highly relevant information, *promote recognition of unexpected information* to avoid premature fixation on a particular viewpoint or hypothesis, and *manage data uncertainty* to avoid misinterpretation of inaccurate or obsolete data (Woods *et al.* 1999). These represent important challenges that may require innovative design concepts and significant departures from current tools (Patterson 1999).

Current tools rely on two general approaches for making complex systems comprehensible. The first is a bottom-up process where algorithms search, group, and transform data according to syntactical patterns in the data. An example of this approach is a search engine that identifies relevant documents based on syntactic or lexical relationship between keywords. The second approach is a top-down organization of the data based on a normative model of the system under consideration. An example of this approach is Yahoo's hierarchical ontology used to organize Internet content. The bottom-up approach fails because linking meaning to statistical relationships between keywords is problematic and the top-down approach fails because the ontology used to organize the data reflects a single perspective, which is not necessarily that of the user. Developing a generic ontology may be a fundamentally intractable problem.

The inherent limits of these approaches require an innovative approach. Agent-based automation could combine the strengths of the top-down and the bottom-up approaches. Swarms of agents could forage for information and collaborate with the user to create an emergent ontology that organizes the information according to the user's goals. Although this innovative design concept offers substantial benefits, it is not clear how designers should support the cognitive demands associated with managing swarms of agents.

Organization and control of swarm behaviour presents several unique monitoring and control challenges compared to traditional systems. The study of biological systems, where many simple individuals produce complex emergent behaviour, may provide some useful insights into how multi-agent automation might be managed (Bonabeau *et al*. 1997). Three important factors govern emergent behaviour: swarm behaviour emerges from parallel interaction between many agents, positive feedback accentuates certain activities, and random variation generates new activities and encourages adaptation (Resnick 1991). For example, swarms of bees dynamically adjust their foraging behaviour to the environment in a way that does not depend on the performance of any individual. A colony of honeybees functions as a large, diffuse, amoebic entity that can extend over great distances and simultaneously tap a vast array of food sources (Seeley 1997). Positive feedback and random variation interact to support robust and effective swarm behaviour (Stickland *et al*. 1995). A specific mechanism that underlies the self-organizing behaviour of swarms is stimergy communication. Stimergy communication is based on a dynamically evolving structure, and is a powerful alternative to a static set of instructions that specify a sequence of activity. Through stimergy, social insects communicate directly through the products of their work. This sort of communication promotes the swarm to evolve into a self-organizing system that can generate many forms of collective behaviour. In this way, the interaction among many simple individuals produces complex behaviour for the group (Bonabeau *et al*. 1997). A specific example of stimergy is the self-organizing foraging behaviour of ants. Stimergy control of foraging behaviour involves a trade-off of speed of trail establishment and search thoroughness; a trail that is more quickly established will sacrifice the thoroughness of the search. Parameters that govern this trade-off include the strength of the positive feedback and the amount of random variation (Stickland *et al*. 1995). The effectiveness of this behaviour can easily be generalized to foraging of robot swarms as they explore a planet's surface or software agents as they explore a complex data set.

Control mechanisms, such as stimergy, offer great potential in extending human capabilities, but only if a thorough empirical and analytic investigation identifies the display requirements, viable control mechanisms, and range of swarm dynamics that can be comprehended and controlled by humans. For example, as a self-organizing system, the swarm could dynamically adjust its information or material foraging behaviour to a dynamic environment to accomplish its goals effectively. This characteristic of multi-agent automation, in contrast to conventional automation, has important implications for how people monitor and control the control of individual robots and the overall swarm. Interacting with agent-based automation requires people to consider swarm dynamics independently of the individual agents. Control of the swarm may involve manipulating global parameters of positive feedback and random variation, which are not traditional ways of controlling automation. Importantly, the functionality of the automation may emerge over time as the swarm of agents adapts to the environment. Because of this, multi-agent automation introduces the need to support emergent function allocation. Cognitive ergonomics should take a leading role in defining the nature of multi-agent automation so that the behaviour and interaction mechanisms are designed to be compatible with human goals and cognitive limits.

3. Modelling approaches to understand human interaction with multi-agent automation

Describing how people interact with multi-agent automation and the associated emergent allocation of function may require analysis and modelling tools that go beyond the traditional techniques. Operator interaction with the emergent functionality and the non-linear dynamics of the multi-agent automation cannot be fully understood with traditional statistics. Understanding human interaction with such system will likely require approaches that depart from standard experimental design and analysis of variance. Figure 2 summarizes a range of analytic approaches to consider the dynamic interaction of operators and multi-agent automation.

The top of figure 2 illustrates a traditional approach, in which observations are averaged over experimental conditions and compared with bar charts and analysis of variance. The focus of most theories and analytic techniques in the area of decision-making and social systems follows this static approach that ignores interactions over

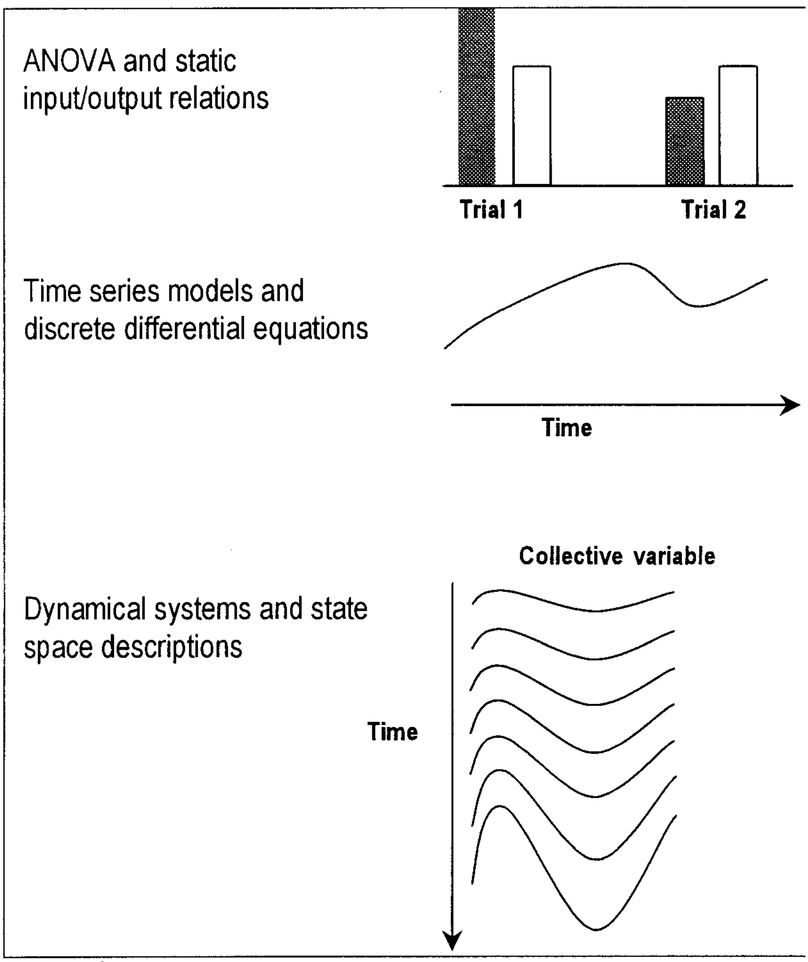

Figure 2. Analytic techniques that provide a progressively comprehensive view of system dynamics that is needed to understand the interaction between operators and the emergent behaviour or multi-agent automation.

time. A static approach can be defined as one in which cognition is defined as symbol manipulation and the passage of time is not considered. For example, traditional decision making theories describe the outcome of the decision process, but not the vacillation that accompanies most decisions (Busemeyer and Townsend 1993). Analysis of variance techniques follow this logic and focus on static aggregations of data without regard to how it varies across time. For example, analysis of variance typically combines data into categories according to experimental conditions and variations over time are simply treated as experimental error.

The middle of figure 2 illustrates a time series approach, where observations are used to estimate parameters of discrete differential equations and plotted as a trajectory over time. As an example, time series analysis techniques produced a discrete differential equation that predicted reliance as a function of trust and self-confidence. This equation accounted for between 60–86% of the variance in operators' reliance on automation (Lee and Moray 1994). Beyond accurately predicting reliance, this analysis approach shows that trust and reliance have an inertia that is critical to understanding the factors that affect reliance. Specifically, these techniques can help define new measures of the calibration of trust and how that calibration changes over time. The static approaches that have been used to identify miscalibration of trust (Lee and Moray 1994) and self-confidence (Lichtenstein and Phillips 1982, Tomassini *et al.* 1982, Wright *et al.* 1994) only consider a snapshot in time and do not consider changes over time. The importance of understanding how a system arrives at a particular state suggests a need for an approach that can provide a process model of system dynamics. As an example, Busemeyer and Townsend (1993) provide such a model and illustrate the power of a dynamical model that describes the behaviour of systems that vary over time. By capturing the time-dependent nature, dynamical systems theory provides a powerful approach to consider how behaviour changes over time. Time series analysis provides a first step in this direction by describing the characteristics of a typical trajectory of system variables over time, a perspective lacking in current analysis of variance approaches.

The bottom of figure 2 shows a dynamical systems approach where the focus shifts from describing a particular trajectory of behaviour to a description of the potential field that guides behaviour. Specifically, time progresses from top to bottom in the figure and the collective variable is depicted from left to right. The collective variable describes the state space of the joint human–automation system. Each horizontal line depicts the probability that the system will be in any of the states described by the collective variable. Attractor states are represented by troughs in the lines. The deeper and steeper the trough the more stable the state. Figure 2 shows that the stability of the various attractor states changes over time. This representation of system provides a rich description of how behaviour evolves over time, providing a field description of the probabilities of various states rather than a description of a particular trajectory. These differences from traditional data analysis and time series techniques may make a description based on attractor states and collective variables particularly well-suited to the analysis of situations where behaviour depends on the confluence of several factors that evolve over time.

Going beyond data analysis techniques, the differences between traditional analysis of variance techniques and those associated with a dynamical approach can lead to differences in the underlying representation of human behaviour. Figure 3 shows the difference between a computational approach, which tends to rely upon traditional analysis of variance techniques, and a dynamical approach to

Figure 3. A comparison of computational and dynamical approaches to modelling social systems showing that the dynamical approach is more appropriate for continuous phenomena like trust in automation (Port and van Gelder 1995).

describing cognition. A dynamical approach considers continuous change of systems over time, whereas the computational approach considers discrete static states. The dynamical approach capitalizes on the power of differential equations to represent continuous interaction among system components over time. The dynamical systems approach, based on attractor states and collective variables, also considers important phenomena such as bifurcation, stability points, and phase transitions. This approach has met with considerable success in describing motor control, but it has only begun to be applied to cognitive and social behaviour (Shaw and Kinsella-Shaw 1988, Kelso 1995, Beer 2000, Thelen *et al.* 2001).

4. Constructs for describing human interaction with complex, counterintuitive agents

Agent-based automation confronts humans with challenges not seen with current forms of automation. Mode errors, misuse and disuse of automation could dramatically increase. The factors that induce mode errors with current automation include indirect mode changes, inadequate feedback, and inconsistent behaviour (Sarter and Woods 1992). The emergent behaviour of agent-based automation may exacerbate all of these factors. The fundamental challenge is that agents may interact to produce emergent behaviour that is not an intuitive extension of individual agents' behaviour. This emergent behaviour may be a very useful characteristic if it is properly designed and managed; however, the challenge of anticipating emergent behaviour may promote new types of errors. For example, people may focus on micro-level behaviour of individual agents and devote inadequate attention to macro-level behaviour.

Attention must be distributed between the macro (swarm behaviour) and micro (individual agent behaviour) levels. An improper focus on one level may allow problems to go undetected on the other. Understanding the factors that guide people to focus attention at a particular level and what factors influence their interventions in controlling the automation may require constructs beyond those traditionally used to describe decision making.

Operators' attitudes toward the automation may play a major role in guiding the users focus of attention and control strategies (Sheridan 1975, Lee and Moray 1992, 1994). In particular, trust is a sociological construct that has been used to describe relationships between people, and much research has shown that trust is an important attitude that mediates how people rely on each other (Rempel *et al*. 1985, Ross and LaCroix 1996). Just as trust influences the relationships between humans, trust may also mediate the relationship between people and the swarm of agents. A series of experiments beginning with Lee and Moray (1992, 1994) and Muir and Moray (1996) have shown that trust is an attitude toward automation that affects reliance and can be reliably measured. Field studies of operators who are confronted with newly installed automation have consistently identified trust as a critical factor guiding reliance on new automation (Zuboff 1988). Highly trusted automation may be used frequently, whereas operators may choose to control the system manually rather than engage automation they distrust (Lee and Moray 1994). A series of recent studies have shown that trust mediates reliance on many types of automation in many different domains. For example, trust was an important explanatory variable in understanding how people react to imperfect traffic congestion information in a navigation system for a car (Kantowitz *et al*. 1997). Trust has also helped to explain reliance on augmented vision systems for target identification (Conejo and Wickens 1998), pilots' perception of cockpit automation (Tenney *et al*. 1998), and control of a tele-operated robot (Dassonville *et al*. 1996). These results show the wide applicability of trust in explaining people's use of automation.

Just as trust plays an important role in mediating the relationship with more traditional automation, trust may guide the control strategies of those supervising agent-based automation. If people trust automation too much they will tend to misuse the automation and rely upon it when it is not appropriate. If they trust the automation too little then the automation may fall into disuse, preventing them from taking advantage of its capabilities. Because of this, calibration of trust is critical to ensure the effective use of automation. Accurately calibrating trust promotes appropriate reliance on automation. Recent neurological evidence provides converging evidence that emotions, such as trust, play an important role in decision making. This research is particularly important because it suggests important measures of the calibration of trust. Damasio *et al*. (1996) shows that people with brain lesions in the ventromedial sector of the prefrontal cortices retain reasoning and other cognitive abilities, but that their emotions and decision-making ability is critically impaired. A series of studies have carefully isolated the decision-making deficit and have demonstrated that it stems from a deficit of affect and not from deficits of working memory, declarative knowledge, or overt reasoning as might be expected (Bechara *et al*. 1997, 1998). The somatic marker hypothesis explains this effect by suggesting that marker signals from the physiological response to the emotional aspects of decision situations influence the processing of information and subsequent response to similar decision situations. Through the bioregulatory process of the mind, emotions and feelings play a key role in guiding people away from situations in which negative outcomes are likely (Damasio 1996). In a gambling decision-making task, those with prefrontal lesions performed much worse than a control group, responding to immediate prospects and failing to accommodate the long-term consequences (Bechara *et al.* 1994). Interestingly, in a subsequent study, a physiological measure of affect, galvanic skin response (GSR), showed a substantial response of normals to the large loss, but not in those with prefrontal cortical lesions (Bechara *et al.* 1997). Interestingly, normals also demonstrated an anticipatory GSR whenever they considered a risky alternative, even before they explicitly recognized the alternative as being risky. These results support a strong argument that emotions play an important role in decision-making, providing neurological evidence that suggests that attitudes such as trust might play an important role in deciding to rely on automation. In addition, these studies show that physiological measures, such as GSR, might be a useful tool in evaluating the calibration of trust. Highly calibrated people will generate an anticipatory GSR when it becomes risky to rely on the automation, whereas poorly calibrated people will not.

Converging evidence from analysis of operator reliance on automation and from neurological investigations into decision-making show that emotions may play an important role in how people manage agent-based automation. This suggests that new constructs that reflect the role of emotions, such as trust and self-confidence, need to be developed. These constructs will complement the more traditional constructs associated with the information-processing paradigm. Identifying, describing and integrating these constructs into the analysis and design of multi-agent automation is a substantial challenge for cognitive ergonomics.

5. Conclusions

Function allocation and effective use of automation have long been important research and design issues. Technological developments in the coming years will challenge the cognitive ergonomics community to address increasingly complex issues associated with these enduring problems. Specifically, the emergence of multi-agent automation poses challenging problems that go beyond those experienced with conventional automation. Like swarms of insects, agent-based automation may generate emergent behaviour that may also produce new error tendencies and require additional display and control considerations. People will be faced with challenges of emergent allocation of function and particularly important considerations include:

- requirements to support monitoring and control of multi-agent automation at the individual and swarm level; and
- requirements associated with stimergy control that may involve manipulations of the agents or the environments the agents inhabit.

Understanding and predicting human interaction with multi-agent automation will require new modelling and analysis approaches. Traditional analysis of variance approaches fail to describe how systems evolve over time, a critical consideration in describing the adaptive capabilities of agent-based automation. In particular, dynamical systems theory may be needed to capture the emergent behaviour that evolves over time and that would otherwise be poorly described. Important contributions of a dynamical systems perspective include:

- time series analysis techniques that describe trajectories of system behaviour in terms of discrete differential equations; and
- collective variables that define attractor states that evolve over time and define a field of potential behaviours.

Just as cognitive ergonomics must consider new modelling and analysis techniques, it must also consider new constructs that might complement an information processing approach to the decision making that governs reliance on multi-agent automation. Particularly important considerations include:

- how operators' trust in individual and collective behaviour of the agents can be calibrated; and
- how emerging findings of neurology can be used to develop novel measures of trust calibration.

This paper demonstrates that addressing the challenge of agent-based automation will require a theoretical understanding of human behaviour that goes beyond a task-based description of well-defined scenarios. Cognitive ergonomics must develop an understanding of the basic cognitive demands associated with managing multi-agent automation, tools that consider the dynamics of the interaction, and constructs that address the dynamic decision making that governs reliance on the automation.

Acknowledgements

This paper is a product of discussions with Geb Thomas, Emily Pollack and Jerry Steele of the University of Iowa.

References

Bainbridge, L. 1983, Ironies of automation, *Automatica*, **19,** 775–779.

Bechara, A., Damasio, A. R., Damasio, H. and Anderson, S. W. 1994, Insensitivity to future consequences following damage to human prefrontal cortex, *Cognition*, **50,** 7–15.

Bechara, A., Damasio, H., Tranel, D. and Anderson, S. W. 1998, Dissociation of working memory from decision making within the human prefrontal cortex, *Journal of Neuroscience*, **18,** 428–437.

Bechara, A., Damasio, H., Tranel, D. and Damasio, A. R. 1997, Deciding advantageously before knowing the advantageous strategy, *Science*, **275,** 1293–1295.

Beer, R. D. 2000, Dynamical approaches in cognitive science, *Trends in Cognitive Sciences*, **4,** 91–99.

Beni, G. and Wang, J. 1993, Swarm intelligence in cellular robotic systems, in *Robots and Biological Systems: Towards a New Bionics* (Berlin: Springer-Verlag).

Bonabeau, E., Theraulaz, G., Deneubourg, J. L., Aron, S. and Camazine, S. 1997, Self-organisation in social insects, *Trends in Ecology & Evolution*, **12,** 188–193.

Brooks, R. A. and Flynn, A. M. 1993, *A Robot Being, Robots and Biological Systems: Towards a New Bionics* (Berlin: Springer-Verlag).

Brooks, R. A., Maes, P., Mataric, M. J. and More, G. 1990, Lunar base construction robots, *Proceedings of the 1990 International Workshop on Intelligent Robots and Systems*, 389–392.

Busemeyer, J. R. and Townsend, J. T. 1993, Decision field theory: a dynamic–cognitive approach to decision making in an uncertain environment, *Psychological Review*, **100,** 432–459.

Conejo, R. A. and Wickens, C. D. 1998, The effects of highlighting validity and feature type on air-to-ground target acquisition performance. *Paper presented at the Proceedings of the 16th Symposium on Applied Behavioral Sciences Symposium*, Colorado Springs, Colorado.

DAMASIO, A. R. 1996, The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, **351,** 1413–1420.

DAMASIO, A. R., TRANEL, D. and DAMASIO, H. 1990, Individuals with sociopathic behaviour caused by frontal damage fail to respond autonomically to social-stimuli, *Behavioural Brain Research*, **41,** 81–94.

DASSONVILLE, I., JOLLY, D. and DESODT, A. M. 1996, Trust between man and machine in a teleoperation system, *Reliability Engineering & System Safety*, **53,** 319–325.

DEGANI, A., SHAFTO, M. and KIRLIK, A. 1995, Modes in automated cockpits: Problems, data analysis, and a modelling framework. *Paper presented at the Proceedings of the 36th Israel Annual Conference on Aerospace Science*, Haifa, Israel.

FITTS, P. M. 1951, *Human engineering for an effective air navigation and traffic control system* (Washington, DC: National Research Council).

FUKUDA, T., FUNATO, D., SEKIYAMA, K. and ARAI, F. 1998, Evaluation on flexibility of swarm intelligent system, *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, 3210–3215.

HANCOCK, P. A. and SCALLEN, S. F. 1996, The future of function allocation, *Ergonomics in Design*, October, **4,** 24–29.

JOHNSON, P. J. and BAY, J. S. 1995, Distributed control of simulated autonomous mobile robot collectives in payload transportation, *Autonomous Robots*, **2**, 43–63.

KANTOWITZ, B. H., HANOWSKI, R. J. and KANTOWITZ, S. C. 1997, Driver acceptance of unreliable traffic information in familiar and unfamiliar settings, *Human Factors*, **39,** 164–176.

KELSO, J. A. S. 1995, *Dynamic patterns: Self-organisation of brain and behaviour* (Cambridge, MA: MIT Press).

LEE, J. and MORAY, N. 1992, Trust, control strategies and allocation of function in human-machine systems, *Ergonomics*, **35,** 1243–1270.

LEE, J. D. and MORAY, N. 1994, Trust, self-confidence, and operators' adaptation to automation, *International Journal of Human-Computer Studies*, **40,** 153–184.

LEE, J. D. and SANQUIST, T. F. 1996, Maritime automation, in R. Parasuraman and M. Mouloua (eds), *Automation and Human Performance* (Mahwah, NJ: Erlbaum), pp. 365–384.

LEE, J. D. and SANQUIST, T. F. 2000, Augmenting the operator function model with cognitive operations: Assessing the cognitive demands of technological innovation in ship navigation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, **30,** 273–285.

LICHTENSTEIN, S. B. F. and PHILLIPS, L. D. 1982, Calibration of probabilities: the state of the art to 1980, in D. Kahneman, P. Slovic and A. Tversky (eds) Judgement Under Uncertainty: Heuristics and Biases (New York: Cambridge University Press), pp. 306–334.

MIN, T. W. and YIN, H. K. 1998, A decentralized approach for cooperative sweeping by multiple mobile robots. *Proceedings of the* 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems.

MUIR, B. M. and MORAY, N. 1996, Trust in automation 2: experimental studies of trust and human intervention in a process control simulation, *Ergonomics*, **39,** 429–460.

PARASURAMAN, R., SHERIDAN, T. B. and WICKENS, C. D. 2000, A model for types and levels of human interaction with automation, *IEEE Transactions on Systems, Man, and Cybernetics- Part A: Systems and Humans*, **30,** 286–297.

PATTERSON, E. S. 1999, A simulation study of computer-supported inferential analysis under data overload, *Proceedings of the Human Factors and Ergonomics 43rd Annual Meeting*, **1**, 363–368.

PORT, R. F. and VAN GELDER, T. 1995, *Mind as Motion* (Cambridge, MA: MIT Press).

REMPEL, J. K., HOLMES, J. G. and ZANNA, M. P. 1985, Trust in close relationships, *Journal of Personality and Social Psychology*, **49,** 95–112.

RESNICK, M. 1991, *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds* (Cambridge. MA: The MIT Press).

ROSS, W. and LACROIX, J. 1996, Multiple meanings of trust in negotiation theory and research: a literature review and integrative model, *International Journal of Conflict Management*, **7**, 314–360.

SARTER, N. B. and WOODS, D. D. 1992, Pilot interaction with cockpit automation I: operational experiences with the flight management system, *International Journal of Aviation Psychology*, **2**, 303–321.

SARTER, N. B. and WOODS, D. D. 1994, Decomposing automation: autonomy, authority, observability and perceived animacy, in M. Mouloua and R. Parasuraman (eds), *Human Performance in Automated Systems: Current Research and Trends* (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 22–27.

SARTER, N. B. and WOODS, D. D. 1995, How in the world did we ever get in that mode? Mode error and awareness in supervisory control, *Human Factors*, **37,** 5–19.

SARTER, N. B. and WOODS, D. D. 1997, Team play with a powerful and independent agent: operational experiences and automation surprises on the Airbus A-320, *Human Factors*, **39,** 553–569.

SEELEY, T. D. 1997, Honey bee colonies are group-level adaptive units, *American Naturalist*, **150,** S22–S41.

SHAW, R. and KINSELLA-SHAW, J. 1988, Ecological mechanics: a physical geometry for intentional constraints, *Human Movement Science*, **7,** 155–200.

SHERIDAN, T. B. 1975, Considerations in modelling the human supervisory controller. *Paper presented at the Proceedings of the IFAC 6th World Congress*, Boston, MA.

SHERIDAN, T. B. 1987, Supervisory control, in G. Salvendy (ed.), *Handbook of Human Factors* (New York: Wiley), pp. 1243–1268.

STICKLAND, T. R., BRITTON, N. F. and FRANKS, N. R. 1995, Complex trails and simple algorithms in ant foraging, *Proceedings of the Royal Society of London Series B-Biological Sciences*, **260,** 53–58.

SUGIHARA, K. and SUZUKI, I. 1990, Distributed motion coordination of multiple mobile robots, *5th IEEE International Symposium on Intelligent Control*, 138–143.

TENNEY, Y. J., ROGERS, W. H. and PEW, R. W. 1998, Pilot opinions on cockpit automation issues, *International Journal of Aviation Psychology*, **8,** 103–120.

THELEN, E., SCHONER, G., SCHEIER, C. and SMITH, L. B. 2001, The dynamics of embodiment: a field theory of infant perservative reaching, *Behaviour and Brain Sciences*, **24,** 1–86.

TOMASSINI, L. A., SOLOMON, I., ROMNEY, M. B. and KROGSTAD, J. L. 1982, Calibration of auditors' probabilistic judgments: some empirical evidence, *Organizational Behaviour and Human Performance*, **30,** 391–406.

WOODS, D. D., PATTERSON, E. S., ROTH, E. M. and CHRISTOFFERSEN, K. 1999, Can we ever escape from data overload? A cognitive systems diagnosis, *Proceedings of the Human Factors and Ergonomics 43rd Annual Meeting*, **1,** 174–179.

WRIGHT, G., ROWE, G., BOLGER, F. and GAMMACK, J. 1994, Coherence, calibration, and expertise in judgmental probability forecasting, *Organizational Behaviour and Human Decision Processes*, **57,** 1–25.

ZUBOFF, S. 1988, *In the Age of Smart Machines* (New York: Basic Books Inc).

About the author

John D. Lee is associate professor of industrial engineering at the University of Iowa. He has a background in engineering and psychology, with a PhD in mechanical engineering from the University of Illinois at Urbana-Champaign. He has 10 years of research and consulting experience aimed at matching human capabilities to the demands of technologically intensive systems. His research interests include interface design, human adaptation to technological innovation, and human performance modelling. Current research projects address modelling driver behaviour, understanding the cognitive demands of in-vehicle technology and trust in automation.