

Designing for dynamic task allocation

Kees van Dongen¹ and Peter-Paul van Maanen^{1,2}

¹ TNO Human Factors
P.O. Box 23, 3769 ZG Soesterberg, The Netherlands
Email: {kees.vandongen, peter-paul.vanmaanen}@tno.nl

² Department of Artificial Intelligence, Vrije Universiteit Amsterdam
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands

ABSTRACT

Future platforms are envisioned in which human-machine teams are able to share and trade tasks as demands in situations change. It seems that human-machine coordination has not received the attention it deserves by past and present approaches to task allocation. In this paper a simple way to make coordination requirements explicit is proposed and for dynamic task allocation a dual-route approach is suggested. Advantages of adaptable automation, in which the human adjusts the way tasks are divided and shared, are complemented with those of adaptive automation, in which the machine allocates tasks. To be able to support design for dynamic task allocation, a theory about task allocation decision making by means of modeling of trust is proposed. It is suggested that dynamic task allocation is improved when information about situational abilities of agents is provided and the cost of observing and re-directing agents is reduced.

Keywords

Dynamic task allocation, decision making, trust, coordination costs.

INTRODUCTION

The goal of this research is to develop knowledge about how activities should be divided, shared and coordinated between human and machine on future platforms. The future of the Royal Dutch Navy is characterized by operational, organizational and technological challenges. Operationally it has to be robust: effective across a range of missions, situations and conditions. Organizationally it has to be both efficient and resilient. Manning will be reduced but the organization has to be able to anticipate and adapt to potential performance degradations. To meet these demands human-machine teams have to be flexible. Future platforms are envisioned on which both human-human and human-machine teams are able to dynamically share and trade tasks as demands in the situation change. This poses the question how activities should be allocated and coordinated between human and machine.

PAST AND PRESENT APPROACHES TO TASK ALLOCATION

MABA-MABA

Traditional methods for task allocation using Fitts lists have several problems (Dekker and Woods, 2002). Three are mentioned here. First, there is no method to describe cognitive tasks, but work is and increasingly will be cognitive in nature. Second, traditional methods stimulate a rigid allocation of tasks to either human or machine. This because lists are used that describe what Man Are Better At and Machines Are Better At (MABA-MABA) and because no methods are provided to determine how tasks should be shared between human and machine. But as manning will be reduced, human-machine cooperation is anticipated to become more important. Third, traditional task allocation methods seem to assume that the abilities of the human and the machine are stable and context-independent. Trading tasks dynamically, when one of the team members is not expected to be able to perform a task successfully, is not considered. For a resilient organization, however, back-up behavior must not disappear with the manning that is reduced.

Levels of automation

The 'levels of automation' approach, as an alternative to the traditional method, provides a simple way to characterize cognitive tasks in terms of information processing activities. Cognitive tasks are described in terms of information acquisition, information analysis, decision making and action implementation. Different ways of dividing and sharing these activities between human and machine have been developed and tested (Wright and Kaber, 2005). In general this

body of research supports automation of information acquisition, information analysis, and action implementation as compared to higher order cognitive activities, such as decision making. Although an improvement to the traditional approach, several problems remain. Again three are mentioned. First, the way information processing activities are represented might lead one to conclude that these activities are sequential and independent (Parasuraman, Sheridan and Wickens, 2000). By representing cognitive tasks as a sequence of information processing activities the coordination that goes on between these activities are obscured. But as Dekker and Woods (2002) have argued, successful automation is not about ‘who has control over what or how much?’ The question for successful automation is ‘How do we get along together’. By making both feedforward and feedback processes between activities explicit, the obscured coordination requirements appear (Figure 1). To illustrate: when situation understanding is insufficient because critical information is missing, new information may need to be acquired (1); when the right action cannot be decided upon because the problem is not fully understood, the situation may need to be reanalyzed (2); when selected actions cannot be implemented because resources are not available, new actions need to be decided upon (3). The information processing activities do not need to be performed in a sequential order.

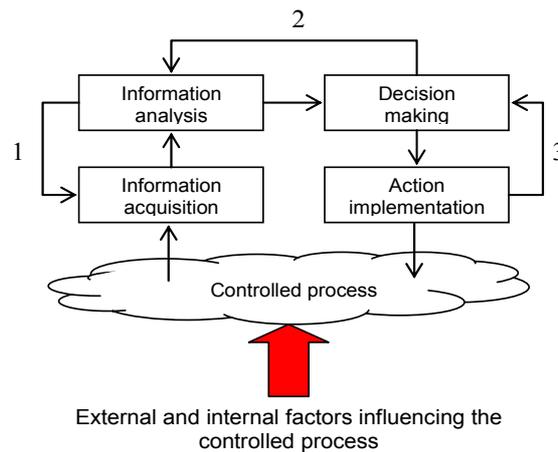


Figure 1. Interdependence between information processing activities.

To coordinate activities in a team in which actors are allocated to interdependent activities, not only the results of activities need to be communicated (feedforward) but also the processes underlying these activities need to be both observable to some degree and re-directable (feedback). Whatever kind or degree of autonomy agents have in performing information processing activities, to collaborate their autonomy needs to be adjustable (Falcone and Castelfranchi, 2001). Further, there is a risk that the benefits of dividing activities between actors do not exceed the cost associated with coordination. Coordination takes time and effort, especially when the interface and underlying functionality for interaction is not properly designed. A second problem with the ‘level of automation’ approach is that it is not clear on when and how human and machine should collaborate within information processing activities. There are many ways in which activities can be shared. How is one to choose? For information analysis, for instance, should the machine integrate and extrapolate acquired information, with the human making the final diagnosis? Should the machine suggest a procedure for diagnosis, which the human may or may not follow? Should the machine critique the procedure the human is or isn’t following or should it critique the diagnosis the human proposes (Guerlain, Smith, Obradovich, Rudmann, Strohm, Smith, Svirebely and Sachs (1999)? Finally, like the traditional approach, in looking for an optimal level of automation this approach also seems to assume that the way activities should be divided and shared is stable and independent of the context. But as was pointed out, the reliability of human and machine performance can change from one situation to the next and the way activities are divided and shared should change correspondingly. To be effective across a range of situations and conditions, human machine teams should be adaptive.

Dynamic task allocation

To overcome problems with static task allocation, recently approaches to dynamic task allocation have been proposed. According to these approaches the way activities are divided and shared between human and machine need to be adjustable after the design phase. Two concepts for dynamic task allocation have been suggested. Both give different answers to the question: who performs the task of allocating tasks? For ‘adaptive automation’ it is the machine that adjusts the way activities are divided and shared, for ‘adaptable automation’ it is the human (Opperman, 1994). The advantage of adaptive automation is that it can serve as a back-up for the human, reallocating tasks when demands, such as time pressure, are too high for the human to react in an accurate and timely manner. The disadvantage is that automation of task allocation can be unpredictable in complex and ill-defined situations and may cause automation surprises when for instance modes are shifted automatically. Adaptable automation, on the other hand, keeps the human

in the loop, foresees in the human need for control, takes advantage of the human ability to anticipate and prepare for changes and provides the flexibility that is needed in novel situations. However, humans are known to make errors in task allocation, potentially misusing and disusing automation (Parasuraman and Riley, 1997; Dzindolet, Peterson, Pomransky, Pierce and Beck, 2003). Although both approaches have different advantages and disadvantages, the choice for either adaptive or adaptable automation is misleading. This can be illustrated by also describing the task allocation task in terms of information processing activities. In this case the controlled processes are the cognitive tasks that are divided between human and machine and not the primary tasks (Figure 2).

Dual route approach to dynamic task allocation

Hinted by research on levels of automation, we suggest that information about performance on cognitive tasks and predictors and causes of potential performance degradations need to be acquired by the machine as much as possible (Figure 2). This can for instance be information about external factors influencing the reliability of human and machine performance, such as the available time for task execution. Second, concerning information analysis and decision making a dual route is proposed. This means that either the machine alone or the cooperative human-machine team determine whether actual or anticipated performance degradations are problematic and whether, when and in what way activities need to be reallocated. For this knowledge about desired and expected performance, and options for task allocation and their expected consequences, need to be taken into account (Figure 2). The machine only route (inner loop) is taken in high dynamic, relatively simple and well defined situations, that is when time for task execution is insufficient for some form of human-machine cooperation (Inagaki, 2003; Moray, Inagaki and Itoh, 2000). In these situations the machine decides to allocate control on critical tasks to automation, in order to respond in a timely manner (Scott, 1999). In this process the machine can notify the human of its diagnoses and decisions (dotted arrows) trying to take the human in the loop and allowing the human to take control. In low dynamic, relatively complex, and ill-defined situations the human-machine cooperation route is taken (outer loop). This route takes advantage of the human ability to reason about their own situational ability to perform tasks, the situational ability of others, in this case the machine, and the costs of coordination in reallocating tasks. In order to be able to reason about the situational ability of others, not only information about internal factors influencing the ability of agents, but also about external factors need to be acquired (Falcone and Castelfranchi, 2001). The dual route approach provides the flexibility that is needed in complex situations in which many factors may have to be considered or in novel and ill-defined situations. This route takes advantage of human experience, with both his own performance and that of the machine, his ability to recognize patterns in complex situations and his ability to use knowledge to reason with incomplete information. Finally, concerning action implementation, when new ways of sharing or dividing of activities have been decided upon, the implementation needs to be performed by the machine as much as possible so as to free up mental resources for concurrent activities.

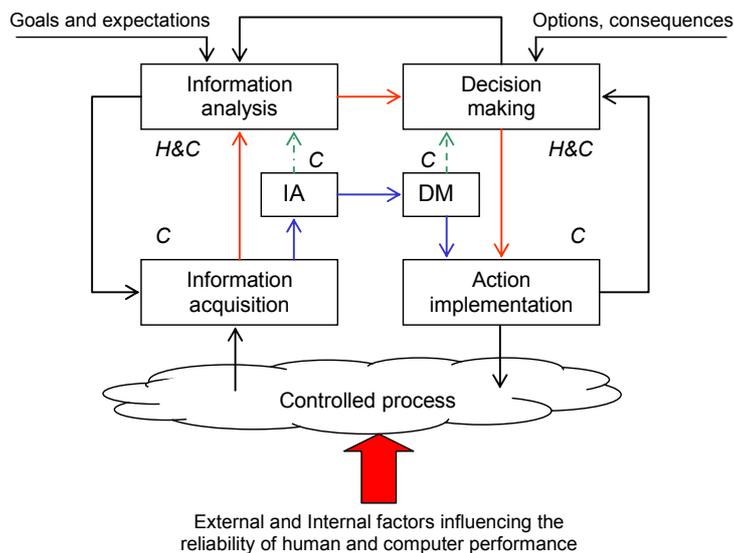


Figure 2. Task allocation: dividing, sharing and dynamic trading of activities.

MODELLING TASK ALLOCATION DECISION MAKING

By complementing the rule-based advantages of adaptive automation (inner loop) with the knowledge-based advantages of adaptable automation (outer loop), the proposed dual route approach is expected to provide the flexibility necessary for resilience and robustness. Although humans are expected to be able to reason about their own and others situational ability and the costs of coordination, they are also known to inappropriately rely on automation. To find out why and to be able to support joint task allocation decision making, we argue that an understanding of the cognitive processes underlying task allocation is required. Because knowledge about situational abilities is often inaccurate and incomplete many theories suggest that task allocation decisions are based on trust. However, Dzindolet, Pierce, Beck and Dawe (2002) point out that trust concepts and their relations are not always well defined in current theories about automation use. To understand task allocation decision making a computational model of task allocation decision making is being developed (Van Maanen and Van Dongen, 2005). How this theory can be used to support dynamic task allocation is addressed at the end of this paper. Below task allocation decision making is described from the human perspective. Key concepts and their hypothesized relations are defined.

Performance goals and trust in self and other

Trust refers to a mental state, a belief of a cognitive agent *i*, in this case the human, about the achievement of a desired goal through another agent *j*, that is the machine, or through agent *i* itself. In trusting automation, the human has, to some level, an expectation that the machines actions will achieve the goal the human desires. In Figure 3 the desired level and expectations about task performance of agents are represented. On the vertical and horizontal axis, representing trust in automation and trust in the human or self respectively, a threshold value distinguishes between unsatisfactory and satisfactory trust. Situational trust takes into account both strengths and weaknesses internal to agents as well as threats and opportunities that are external to agents. Tasks that fall into the U_h region, indicate that the human must depend on automation to meet performance goals, for instance because multitasking demands are too high. Tasks that fall into the U_a region indicate that the human must depend on him or herself to meet performance goals, for instance because a support system is does not work in that situation. Tasks in which trust in both automation and human is unsatisfactory (U_{ah}) may for instance require a renegotiation of the desired level of performance or require support by an alternative agent. The utility of performance the human expects of automation or him or herself, is represented on the axes of Figure 3. The blue diamonds indicate a situation in which the human trust in him or herself is less than that in the machine. The red triangles represent a situation in which the opposite is true.

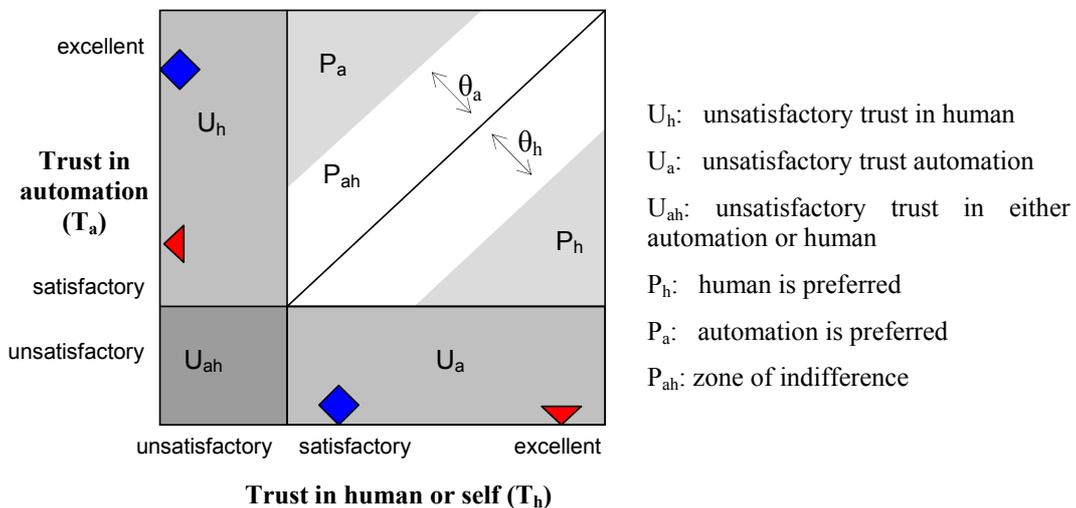


Figure 3. Task allocation decision making (inspired by Price et al. (1982)).

Trust dynamics, direct and indirect sources of knowledge and biases

Trust is not static, it changes in time as it is influenced by direct and indirect sources of knowledge. Human expectations about the utility of machine performance, as a result of their situational ability are calibrated by direct experience. Trust is dynamic, but it does not simply increase or decrease with positive or negative experiences. How trust is changed by successes and failures, for one, depends on how discrepancies between expected and actual performance are causally attributed (Falcone and Castelfranchi, 2004). Trust is indirectly influenced by reputation and gossip. This enables agents to form expectations without direct experience. Trust can also be based on analogical judgments, judgment about the

trustworthiness of a category rather than on the actual performance of one of its presumed members. Etiquette may for instance influence trust because category membership associated with adherence to a particular etiquette helps people to infer how automation will perform (Miller, 2002). Trust in own and others ability is not always veridical and can be biased (Lee and See, 2004). Humans, for instance, overestimate the number of tasks they can complete in a given period of time (Buehler, Griffin and Ross, 1994) and initial expectations about machine performance can be too high (Dzindolet, Pierce, Beck and Dawe, 2002).

Task allocation decisions, relative trust, workload, and coordination costs

The level of trust is distinguished from the decision to allocate a task to an agent or rely on an agent. The term 'task allocation decision' is used to refer to the decision to rely on an agent's actions to achieve a desired goal. Many theories in the human factors literature about reliance include a notion of relative trust, i.e. the difference between levels of trust in agents. Empirical results show that as trust in machine performance is significantly higher than trust in own performance, humans prefer to allocate tasks to the machine, and when the reverse is true, humans prefer to rely on themselves (De Vries, Midden and Bouwhuis, 2003; Moray, Inagaki and Itoh, 2000; Dzindolet, Beck and Pierce, 2000). One might expect that an agent is more likely to rely on another agent when its workload is high compared to when it is moderate or low. Parasuraman and Riley (1997), however, point out that the relation between workload and the reliance decision has not been empirically validated. Humans do not simply allocate tasks to automation so as to free up mental resources for concurrent tasks. There can be many reasons. Workload may just not be high enough to sufficiently decrease trust humans have in their own performance relative to their trust in the machine. Humans may prefer to rely on themselves to maintain their skills on that task and compensate the increases in workload with extra effort (Veltman and Jansen, 2004). They may prefer to rely on themselves because of their need for control or simply because other tasks are not very important to them. Allocating tasks just based on changes in workload seems to be too simple. The 'task allocation decision' can conceptually be distinguished from the overt actions of allocating a task or actually relying on an agent. Kirlik (1993) has pointed out that reliance decisions are not only influenced by individual differences in for instance skill or by the costs of delaying concurrent tasks, but also by expected costs of coordination, for instance the effort or time to engage and disengage automation. It is expected that in task allocation decisions the costs of coordination are taken into account, including effort and time required for deciding how to reallocate tasks. It is further expected that the effect of coordination costs will be particularly evident on task allocation decisions when the workload is already high. Again referring to Figure 3, tasks that fall into the region P_h , P_a or P_{ah} , in which both human and automation are expected to perform satisfactory offer more options for task allocation. In these regions decisions are based on preference beliefs instead of dependence beliefs. In region P_a , the human prefers to allocate the task to automation instead of to him or herself, because trust in own performance is significantly lower than trust in automation, taking the costs of coordination into account. In region P_h , the task is preferably allocated to the human. One factor that is expected to influence coordination costs is whether the preferred agent was already performing the task. Coordination costs may not be symmetric in the sense that allocation to oneself may not result in the same costs as allocating to the other. This thus results in the possibility that allocation preferences are different from what you might expect from the trust values. In region P_{ah} the difference in trust, corrected for coordination costs, is not high enough to exceed the decision threshold. In this region the decision-maker is indifferent. It is expected that the height of the threshold is influenced by uncertainty about trust and coordination costs, the importance, and urgency of task allocation decision at hand. More on the formalization of these decision processes see (Bussemeyer and Townsend, 1992; Bussemeyer and Diederich, 2002).

CONCLUSION

In this section we propose how theory on task allocation decision making can be used to guide the design for dynamic task allocation in human-machine cooperative systems. According to the presented theory, flexibility in task allocation is gained when information about situational abilities of agents is provided and the cost of coordination is reduced. In other words, dynamic task allocation is supported when the costs involved in observing and re-directing agents are reduced (Christofferson and Woods, 2002). By reducing uncertainty about situational abilities of agents, for instance by displaying their context-dependent abilities, the task allocator can make better reliance decisions, trust can be calibrated by direct experience and causal attribution of unexpected performance is supported. To reduce the costs involved in observation, displays which rely on human pattern-matching abilities may be developed. Dynamic task allocation is also expected to improve when options to re-direct agents are bounded and when the consequences of re-direction are predictable. Although many kinds and levels of autonomy can be thought of and perhaps designed, coordination costs and uncertainty about expected performance increase when many options with unpredictable consequences are provided. Given opportunities to choose modes of automated behavior, operators must spend effort in making a decision, keep track of it and act consistent with it. Too many options may result in inappropriate decisions, failure to maintain awareness of modes or lead to mode errors (Jamieson and Vicente, 2005). Operators may fail to intervene when automation takes inappropriate action or can engage a mode that is not appropriate for a given situation (Sarter and Woods, 1992). In short, to design for dynamic task allocation, keep it simple and transparent.

Theory on task allocation decision making can also be used to teach operators how to make task allocation decisions or to guide the design of training scenarios. To train for dynamic task allocation, scenarios can be developed in which operators learn when to rely on themselves or the machine. When designed for, operators need to be trained to notice and recognize cues that signal the need for task allocations. Training is expected to reduce uncertainty about estimated coordination costs and trust, thereby reducing the decision threshold. This would result in better task allocation decisions. Scenarios can also be used to decide on initial ways to effectively divide and share tasks between human and machine. It is also possible to track and display past performance of machine performance and use across situations. In this way operators are supported to decide on how to divide and share activities based on the reputation machines have in these situations.

It is also important that artificial agents are able to take the initiative in supporting humans in making task allocation decisions. For instance, in taking initiative in directing attention in the case of information acquisition, to advise or critique in the case of information analysis and decision making, or to support action implementation (see outer loop, Figure 2). In agent technology research, the subjects *delegating agents* and *adjustable autonomy* consider agents to reason about the agent's own ability to perform tasks, that of others, and the costs of coordination. Human-machine cooperation is supported by adjusting each other's autonomy by means of adjusting the style of delegation and support. For instance, in (Falcone and Castelfranchi, 2001) such styles are proposed in which trust plays a major role. In order for an artificial agent to adapt to the current situation, it is important to let it have the right situation-dependent initiative to support. This suggests that such agents should be enabled to acquire information about the specific goals of the human or whether certain support will be appreciated. Subsequently the agent should be able to decide on what goals it adopts and which tasks it should select for reaching them. An important issue here is the trade-off between the gain of agent initiative and the costs of the resulting coordination. As mentioned before, increased initiative may lead to unwanted decrease of human autonomy or lack of control.

The present research is work in progress. To validate implications of the theory a simple experimental task is being developed (Van Maanen and Van Dongen, 2005). After being confident about the replicability of previously found experimental findings, the environment will be used for further research. The following manipulations are considered: reliability of machine performance (e.g., simulated errors), reliability of human performance (e.g., changes in task difficulty, number of parallel task), level of autonomy (mixed initiative and types of delegation and support), coordination costs (e.g., effort to engage in allocation), and time pressure. The following theoretical aspects of task allocation decision making are studied: performance estimation (e.g., knowledge acquisition based on experience and on indirect sources, such as reputation, gossip, and analogical judgments), coordination costs, the allocation decision threshold (e.g., effects of uncertainty, urgency and importance). Furthermore, future research will focus on means to acquire information about trust, its integration to trust concepts, task allocation decision making based on trust, and the implementation of the allocation decision.

ACKNOWLEDGMENTS

This research was funded by the Dutch Ministry of Defense under program Human-Machine Task Integration Programme V206. Gratitude goes to Jan Maarten Schraagen for his helpful comments.

REFERENCES

1. Buehler, R., Griffin, D. and Ross, M. (1994). Exploring the "planning fallacy": Why people underestimate their task completion times. *Journal of Personality and Social Psychology*, 67, 366–381
2. Busemeyer, J.R. and Diederich, A. (2002). Survey of decision field theory. *Mathematical Social Sciences*, 43, 345–370
3. Busemeyer, J.R. and Townsend, J.T. (1992). Fundamental derivations from decision field theory. *Mathematical Social Sciences* 23 255–282
4. Christofferson, K. and Woods, D. D. (2002). How to make automated systems team players. In E. Salas (Ed.), *Advances in Human Performance and Cognitive Engineering Research, Vol. 2*. JAI Press, Elsevier.
5. De Vries, P., Midden, C. and Bouwhuis, D. (2003) The effects of errors on system trust, selfconfidence, and the allocation of control in route planning. *International Journal of Human-Machine Studies*, 58, 719–735
6. Dekker, S.W.A. and Woods, D.D. (2002). MABA-MABA or Abracadabra? Progress on human-automation coordination. *Cognition, Technology & Work*, 4, 240-244
7. Dzindolet, M.T., Beck, H.P. and Pierce, L.G. (2000). Encouraging human operators to appropriately rely on automated decision aids. In Proceedings of the 2000 Command and Control Research and Technology Symposium, Monterey, CA
8. Dzindolet, M.T., Peterson, S.A., Pomransky, R.A., Pierce, L.G., & Beck, H.P. (2003). The role of trust in automation reliance. *International Journal of Human Machine Studies*, 58, 697–718.
9. Dzindolet, M.T., Pierce, L.G., Beck, H.P. and Dawe, L.A. (2002). The perceived utility of human and automated aids in a visual detection task. *Human Factors*, 44, 79–94

10. Falcone, R. and Castelfranchi, C. (2001). The human in the loop of a delegated agent: the theory of adjustable social autonomy. *IEEE Transactions on Systems, Man, and Cybernetics*, 31, 406-418
11. Falcone, R. and Castelfranchi, C. (2004). Trust dynamics: How trust is influenced by direct experiences and by trust itself. In Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), New York, USA, 740-747
12. Guerlain, S., Smith, P. J., Obradovich, J. H., Rudmann, S., Strohm, P., Smith, J. W., Svirebely, J., and Sachs, L. (1999). Interactive critiquing as a form of decision support: An empirical evaluation. *Human Factors*, 41, 72-89
13. Inagaki, T. (2003). Automation and the cost of authority. *International Journal of Industrial Ergonomics*, 31, 169-174
14. Jamieson, G. A. and Vicente, K.J. (2005). Designing effective human-automation-plant interfaces: a control theoretic perspective. *Human Factors*, 47, 12-34.
15. Kirlik, A. (1993). Modeling strategic behavior in human-automation interaction: Why an "aid" can (and should) go unused. *Human Factors*, 35, 221-242
16. Lee, J.D. and See, K.A (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46, 50-80
17. Maanen, P.P. van, Dongen, K. van, Towards Task Allocation Decision Support by means of Cognitive Modeling of Trust, In: Proceedings of the Eighth International Workshop on Trust in Agent Societies (Trust 2005), To appear, 2005.
18. Miller, C.A. (2002). Definitions and dimensions of etiquette. Technical Report FS-02-02, American Association for Artificial Intelligence, Menlo Park, CA
19. Moray, N., Inagaki, T. and Itoh, M. (2000) Adaptive automation, trust, and self-confidence in fault management of time-critical tasks. *Journal of Experimental Psychology: Applied*, 6, 44-58
20. Opperman, R. (1994). Adaptive user support. Hillsdale, NJ; Erlbaum.
21. Parasuraman, R. and Riley, V.A. (1997) Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39, 230-253
22. Parasuraman, R., Sheridan, T.B., & Wickens, C.D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics*, 30, 286-297
23. Price, H. E., Maisano, R. E. and VanCott, H. P. (1982). *The allocation of function in man-machine systems: A perspective and literature review* (NureG-CR-2623). Oak Ridge, TN, Oak Ridge National Laboratory.
24. Sarter, N.B. and Woods, D.D. (1992). Mode error in supervisory control of automated systems. In Proceedings of the Human Factors Society 36th Annual Meeting (pp. 26-29). Santa Monica, CA: Human Factors and Ergonomics Society.
25. Scott, W. B. (1999). Automatic GCAS: 'You can't fly any lower'. *Aviation Week & Space Technology*, February, 76-79.
26. Veltman, H.J.A. and Jansen, C. (2004). The adaptive operator. In D. A. Vincenzi, M. Mouloua, P. A. Hancock. Proceedings of the Second Human Performance, Situation Awareness and Automation Conference (HPSAA II), Daytona Beach, FL.
27. Wright, M.C. and Kaber, D.B. (2005). Effects of automation of information-processing functions on teamwork. *Human Factors*, 47, 50-66.