Capability, Cognition and Autonomy

Robert M Taylor
Dstl Human Sciences
Air Systems
Ively Gate, Ively Road
Farnborough
Hants GU14 0LX
UNITED KINGDOM

SUMMARY

The aim is to consider the role of humans in intelligent and automated systems, with enhanced capability from ubiquitous information, computer and communication technologies. The key human factors (HF) issues are generally believed to be what are appropriate levels of automation for military decision functions, and how to mitigate the associated risks for human understanding, prediction and control of system functioning. The aim of this paper is to attempt to synthesise these issues into a view of human roles through frameworks of military capability, cognition and autonomy.

1.0 INTRODUCTION

Technology can provide both technical artefacts (tools and means for human problem-solving) and technical agent-based systems (software programs that mimic human properties with behaviours, goals and intentionality). Agents (human and machine), artefacts (technical), and the environment can be considered as world domains interacting in a system of systems. It is the interaction between agent, artefacts and the environmental domains that produces the changes in human roles of interest in the present context. The challenge with new technologies is to understand and predict the influences of these interactions, and the changes that they produce, for future human roles. This understanding is needed so that people remain in control of systems, and for the system functioning to be human-centric i.e. serving human functional purposes. To support this, the utility of options for technical systems (artefacts and agents) needs to be judged against principles and criteria for optimisation of human effectiveness. Human effectiveness issues should govern the design and development, and guide the exploitation and use, of the intelligent and automated systems of present concern. These are not new issues (e.g. Asimov's laws of robotics). Understanding the implications for human roles in future systems is a continuing challenge. Recent developments in intelligent agent systems and uninhabited vehicles have added special impetus to the debate. It is helpful to understand how these issues have developed. In particular, we need to the consider what success there has been in developing the HF ideas identified previously e.g. collaborative computer support, human-computer teamwork, dynamic function allocation, adaptive automation, cognitive automation, cognitive engineering (NATO RTO, 1998; Reising, Taylor, and Onken, 1999; NATO RTO, 2002).

1.1 Function Allocation

Automation is continually improving in capability with associated changes in perceptions of appropriate human roles and the suitability of functions for human and/or machine performance. Traditional engineering mostly used the *"left over"* principle for allocation of function, where the technical system was designed to do as much as is feasible from an efficiency point of view, and the rest was left for the

operator. HF engineering introduced the *compensatory* principle, where human and machine capabilities are compared on salient criteria and the function allocation is made so that the respective capabilities are used optimally. In 1951, Paul Fitts suggested some simple criteria for allocating functions between people and machines to predict roles in future air navigation and air traffic control systems (Fitts, 1951; Republished in Beevis, Essens and Schuffel, 1996). These are summarised below from Sheridan (1996):

Table 1: Fitts' List

What people are better at	What machines are better at
Detecting small amounts of visual, auditory, or chemical energy	Responding quickly to control signals
Perceiving patterns of light or sound	Applying great force smoothly and precisely
Improvising and using flexible procedures	Storing information briefly, erasing it completely
Storing information for long periods of time, and recalling appropriate parts	Reasoning deductively
Reasoning deductively	Doing many complex operations at once
Exercising judgement	

Asking what roles can the human be assigned in future systems, Fitts distinguished between four kinds of control systems, namely: (1) Fully automatic control; (2) Automatic control with human monitoring; (3) Semi-automatic control supplemented by human performance of critical functions; (4) Primary control by human operators. In the latter, the human operators would be assisted "by effective data analysis, data transmission and data display equipment", thus anticipating the role of computers. After analysis of issues of alertness and overload and breakdown under stress, and human fallibility, Fitts proposed that it is an important working principle that checking, verifying and monitoring equipment should be devised in ways that make it impossible for a human to violate basic safety rules. A recent controversial implementation of this idea has been the introduction into the civil aviation flight deck of control "limiters" that prevent the aircraft pilot from exceeding extremes of designed flight envelopes e.g. unusual positions in roll, pitch. As a general rule, Fitts proposed that machines should monitor humans, especially in matters of safety, and prevent them making serious mistakes. There has been relatively slow take-up of this machine-monitoring-man concept - a prototype pilot "cognition monitor" system has been recently developed capable of detecting and initiating counter-measures for spatial disorientation (Taylor, Brown, and Dickson, 2002). On the question of who should make decisions, Fitts suggested that other things being equal, the person who is informed is obviously the best person to make decisions. Also, he advised that decisions should be made near the point where basic information is derived. In military aviation, the importance of timely and correct decision information – coupled with knowledge and experience – is well ingrained in approaches to human systems design. This is reflected in concerns over maintaining situational awareness (SA) in highly automated cockpits and the need for combat identification (ID) capability, and in the enduring primary system design rule of keeping the pilot in charge (and resisting cognition monitor take-over until provably correct). This is illustrated by the first rule of flight control "if you don't know what to do, don't do anything". Also, it is embodied in the pilot adage of "being ahead of the airplane" or "ahead of the power curve" – having the experience to anticipate what could happen rather than just reacting to what was happening at that moment in time (Krantz, 2000). This contrasts with the old Soviet Air Force policy, governed by political priorities, of ground control of air defence aircraft combat decisions. Conflicting guidance information from a cockpit collision avoidance system and an air traffic controller has been cited as contributing to pilot decision uncertainty and delay leading up to the recent European civilian aircraft mid-air collision.

Analysing, identifying and allocating functions is commonly regarded as an essential early step in the HF engineering process. Methods of early function analysis are described in NATO Military Standard

STANAG 3994, The Application of Human Engineering to Advanced Aircrew Systems (NATO NSA/AI, 2000). Standardisation is limited to principles and procedures for function analysis, rather than defining function allocations. Beevis (1996) reported a set of often overlooked human resource functions identified from study of multi-crew operations:

- Co-ordination

- Consultation

- Resolution of ambiguity

- Crew performance monitoring,

- Maintenance of awareness of the system state

- Training

- Reversionary mode operation

- Maintenance of alertness

Functions are abstract concepts that change and evolve with the development of systems thinking. Their description depends on the level and scope of system definition. New functions and human roles are likely to emerge as the capability of advanced system develops. Including humans in the system definition and description is important for a full understanding of the human functional issues.

The continuing uncertainty over the role of humans in advanced systems was identified as one of the major research issues at 1994 NATO RSG workshop on function allocation (Beevis, Essens, and Schuffel 1996). In considering how should humans and machines work collaboratively, the questions posed by Fitts and his colleagues were considered still to lack general answers.

- Should the human monitor the (technical) system given that humans are poor monitors?

- Should the (technical) system monitor the human?

- If so what roles should the human play and what are their responsibilities?

- Are humans included in systems just to deal with those functions that engineers can not automate?

Options on decision making were noted to range from the principle that the human should make all decisions, because humans are responsible for systems, to the principle that there are some decisions that humans should never be permitted to make. Categories of programmed ultimate authority were proposed as a function of level of abnormality i.e. critical, caution, routine (Sheridan, 1996).

Advanced computing, information and decision systems seem increasingly capable of performing aspects of high-level cognitive decision functions. The term *cognitive automation* has been introduced to describe this intelligent machine capability for high level cognitive functioning. Technology is available that can recognise complex patterns, work with unstructured information and knowledge, reason with context and uncertainty, and operate with human agent-like properties (e.g. beliefs, desires and intentions). Thus, in theory, technology can be used either to support or replace human functioning. But there is uncertainty about if, and when, computers will perform these functions better than a highly trained and experienced human. There is concern about whether computers can, and should be, trusted to make critical decisions affecting the use of military force (e.g. targeting). This sits with concern about the degrading effects of battle stress on human decision skills and cognitive capability.

1.2　Human-Computer Collaboration

In the 1980's, with increasingly capable intelligent computing, ideas of human-computer teamwork, cognitive engineering, cognitive automation and joint cognitive systems (Hollnagel and Woods, 1983),

began to emerge. According to the *complementarity* principle, function allocation serves to support and sustain human ability to perform efficiently. Here, the focus shifts from human-machine interaction to human-computer co-operation, and from the internal functions and structure of the human and machine to the external functions and establishing the system boundaries (Hollnagel, 1997).

So, automation gets better all the time. But this comes with risks. In 1987, increasing automation capability was observed as having the following consequences (Kantowitz and Sorkin, 1987):

- The human must become a monitor of automation – however it is known that the human is a poor monitor, unless aided in certain ways.

- Increased automation means increased training requirements.

- Newly automated systems have bugs.

- Failure of automation leads to a loss of credibility and trust.

- Designers tend not to anticipate new problems that automation brings with it (e.g. mode errors and feelings of alienation).

The implications of increasing levels of automation for maintaining pilot's situational awareness (SA), skill, trust and dependency has been extensively documented and discussed. Under the complementarity approach, an estimate of the risks for human effectiveness of human, automation and shared/joint control of functions is attempted in Table 2 below:

Table 2: Human Effectiveness Risk List

Risks of human control	Risks of shared/joint control	Risks of automation control
Cognitive bias – conservatism expectancy, set	Authority, intent ambiguity	Automation bias - dependency
Complexity, overload, error	Model, monitor, infer ambiguity	Out-of-the-loop human performance
Pre-occupation, fixation	Incorrect aiding, knowledge	Poor situation and mode awareness, take-over
Failure to evaluate options	Interaction, comms complexity	Monitoring, vigilance, boredom
Forgetting rules and procedures	Interface over-simplicity	Surprise, unpredictability unexpected action
Time pressure	Opaqueness of dynamics	Mistrust, under-use
Endurance, fatigue, inattention	Separated desynchronised command, SA, planning	Unresponsiveness to change (ROE, context)
Breakdown of skill	Slowed decision tempo	Skill fade
	Balancing pro-active/reactive feedback/forward	
	Inadequate SA for anticipation, with low demands	
	Hand-over, take-back unreadiness	

1.3 Unmanned Systems

Recent advances in sensors, communications, computing and information technology have provided humans with remote control roles for a variety of uninhabited vehicles and platforms. This gives the

human operator improved accommodation, particularly for unmanned air vehicle (UAV) control, and the ability to perform roles in relative comfort and safety. Increased remote control capability is sought to reduce manning and training costs. A future systems design goal is for one operator to control many platforms (vice the current many controlling one). Enabling technologies include virtual teams and virtual human-machine interaction, alternative control/display technologies, cognitive automation and decision-support, as well as intelligent automated systems. The term "unmanned systems" has been coined by systems engineers to describe autonomous, uninhabited vehicles. This is somewhat misleading, and it is almost certainly inappropriate from a HF perspective. It suggests an absence of human involvement, and a lack of human system issues. This is unfortunate since the opposite is closer to the truth. The need for human involvement and input to system design and operation is paramount and needing to be in the sharpest focus. At the present time, advances in autonomous vehicle technologies are worthless without an effective operator and remote control/display interface. Currently, the human operator provides the flexibility to adapt to constraints on functioning arising from system design, on-line tolerance of variability in the external environment, and adaptation to changing dynamic mission-context. In the future, humans may have only indirect control of operations, through the initial system specification and design, and then through the mission and tasking instructions for operations. Progress in technology has made it possible to envisage future uninhabited vehicles with on-board cognitive automation operating with relatively high levels of autonomy or decision authority. With future cognitive "unmanned systems", such as Uninhabited Combat Air Vehicles (UCAV), there may no longer a "Government Furnished" human always on-line, and in-the-loop with the authority, responsibility and SA to deliver the intended effects and safe system performance. The aspiration is to achieve the requisite cognitive agility, precision, reliability and safety of operations with intelligent systems, and not necessarily man-in-the-loop. Future systems are envisaged with "intelligent" computer software agents designed to control, regulate, direct and adapt system behaviour in uncertain, novel, and unpredictable situations. For this is to be achieved there will need to be a good understanding of cognitive functions and a successful transfer from the human domain to systems domain. The challenges are in understanding the cognitive domain.

Thus, it seems that even "unmanned systems" will continue to have important human roles. The focus of HF interest has merely shifted away from habitability, operability and usability issues, towards the providing clear definition, proof and validation of purpose during system specification and design, and in system command and control. It is therefore worrying that poor specification of system requirements is the most frequently cited source of difficulty in providing good HF integration in defence equipment acquisition programs. The effective specification of high level human performance requirements, such as the correct balance of feedforward and feedback control, and balance of information gathering and decision tempo, presents significant challenges for the systems engineering community.

1.4   Key Human Role Issues

With the increasing ubiquity and capability of the computers, information, communication and decision systems that provide cognitive automation, a key question is how do we ensure the optimum contribution of humans and automation to military system effectiveness? The issue of optimisation can be decomposed into asking do we know:

- Why, and when only humans should contribute to military system effectiveness?

- Whether, and if so, when automation should contribute to military system effectiveness?

For cognitive automation, optimisation needs to be considered with reference to high level of decision functions, by asking do we know:

- Why, and when only humans should make decisions about the use of military force?

- Whether, and if so, when and what cognitive automation should be used to make decisions about the use of military force?

It is generally believed that technology should used be used to automate routine predictable tasks, and used to assist and support, rather than replace, human performance of high-level decision functions. How this assistance and support is to be provided most efficiently and effectively requires understanding about the current and future military environment and it's changing requirements (i.e. context, capability), about ideas for increasing human effectiveness (i.e. augmenting cognition, control), and of developments in advanced automation capability (i.e. intelligent systems, autonomy).

2.0 CAPABILITY

A simplistic view of human role development, compared with traditional operator controlled "master-slave" man-machine systems, sees the role of humans in intelligent and automated systems as likely to be diminished and mostly supervisory in nature. This "press and play" automation system viewpoint, mostly comes from the process control experience of automation systems (industrial production processes, air traffic control). It is an appropriate issue for relatively stable, static and predictable "closed-loop" systems and processes, with relatively closed environments and bounded problem contexts. Here, abnormalities, accidents and emergencies requiring human intervention are intended to be relatively rare, in contrast to the military battle-space environment.

2.1 Digitisation

Digitisation of future battle-space, specifically the command, control and information infrastructure (CCII) is currently proposed in support of information superiority initiatives. In contrast to industrial production and process control, the military battle-space is characterised by change, uncertainty, open-loop systems, complex environments, and the need for a highly agile, freedom of manoeuvre approach to control of military power. The need for human involvement in command and control processes of automated intelligent battle-space systems is seen as a central and critical issue. Again, this requires a strong user input during system specification and design to optimise the command and control architecture, and the fullest human involvement in delivery of command and control of operations. Here, the focus of human input is in the formulation, setting and monitoring of system function purpose, in providing reactive and pro-active strategic goals and targeting objectives, and in planning and execution of operations for the intended strategic surgical effects. So, rather than passive supervision and monitoring, the human role is to provide cognition of the system, in the sense of thinking, conceiving and reasoning. It is from appreciating the situation and command intent, and in providing the understanding, sense-making and prediction needed to plan and execute operations that the strategically required effects are produced.

2.2 Military Environment

Human roles in military systems can change for a variety of reasons other than technology insertions. Changes can be to due to a broad set of internal and external influences (political, legal, social/cultural, economic), but need to be understood with reference to the effects of the prevailing military environment. In 2002, the "fog of war" seems increasingly unpredictable, opaque and uncertain. Alongside traditional attritional warfare providing progressive degrading of adversary capability, new threats come from international terrorism and asymmetric warfare, involving adversaries that are highly wilful, if not always highly capable, and thus primary targets for influence operations. Urban operations and use of less destructive capabilities may combine with complex, peace-making and peace-keeping influence operations. Adversaries and threats must be countered by democratically approved and legally justifiable means, with operations conducted under the ever-watchful, collateral-sensitive, "CNN" public news media eye. Operations increasingly involve collaborating forces and joint operations, needing complex co-ordination, synchronisation and clear commonality of purpose.

2.3   Defence Capability Framework

Responding to these challenges, UK MOD strategy for R&D and equipment procurement programs has moved from a platform-centric approach to a capability-based framework. This capability-based approach uses a vocabulary more readily mapped onto human role effectiveness issues and dimensions than platform-centric language. The UK MOD Defence Capability Framework (DFK) has seven concentric dimensions along three axes, namely (1) to command and inform, (2) to sustain and protect, and (3) to prepare, project and operate. The structure of this capability framework is illustrated in Figure 1.
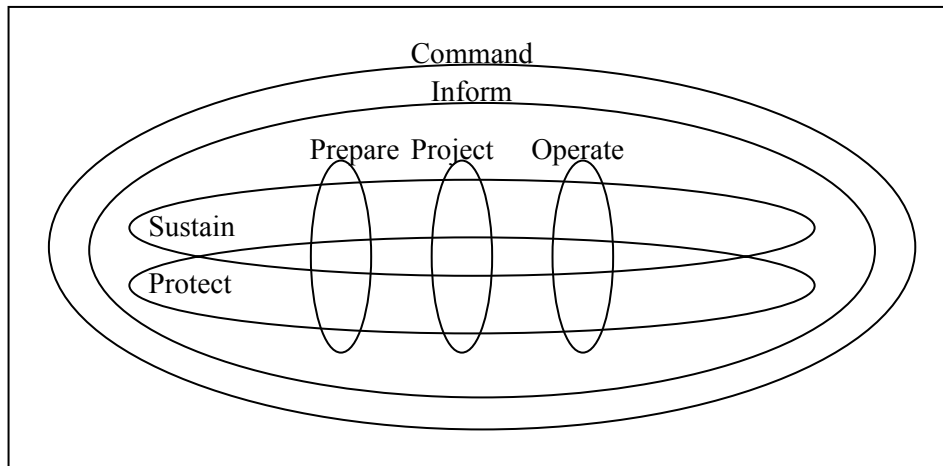

Figure 1: UK MOD Defence Capability Framework.

The DCF framework represents military understanding of a number of important inter-relationships and dynamics. This provides important clues that indicate military priorities for human roles with network enabled capability.

- A fusion and coherency of command and inform functions, with network enabled shared situation awareness adapting the degree of control to political sensitivities. (Currently, three broad categories of command and control are used – Mission Specific, Objective Specific, Order Specific – in order of increasing centralisation and specificity of directives).

- Detailed preparation using shared knowledge, with rapid transition, for timely projection of highly flexible and adaptable forces.

- Shared awareness during projection to allow flexibility and adaptation to ongoing developments.

- Decision agility at all levels of command through shared awareness of the unfolding situation.

- Early integrated logistics for the sustained endurance to deliver the right effect at the right time.

- Shared planning and awareness to balance offensive and defensive action and implicit force protection tasks.

2.4   Effects-Based Operations

Emerging military doctrine is described in a recent paper on Effects Based Operations (EBO) from the UK MOD Joint Doctrine and Concepts Centre (JDCC, 2002). This focuses on the need to identify the effects that lead to campaign success. Whilst focussing on the commander's problems, this JDCC understanding is also relevant to operator's requirements, since the approach aims to provide command intent and decision superiority at all levels. To broaden consideration of it's applicability to decision making at all levels, the term decision-maker is substituted for commander in this brief summary of the JDCC position.

EBO are enabled by knowledge superiority, battle-space exploitation and effects-based planning. The JDCC EBO paper calls for a manoeuvrist approach, providing freedom to manoeuvre with decision and force agility. JDCC military understanding of agility identifies the following attributes:

- Responsiveness – speed in the recognition of the need for action or change, and to seize the initiative.

- Robustness – the degree to which a force is effective across a wide range of situations and missions.

- Flexibility – the capability to identify and execute actions along multiple paths to success.

- Adaptability – the degree to which a force learns about it's operating environment and acts upon that knowledge.

Adaptability increases in importance when operating in unfamiliar situations. Military understanding considers that adaptability should be measured in three ways:

1) The development of the decision-maker's critical information requirements. This needs to be accurately defined, and then the correct balance needs to be formed between information gathering and the optimum tempo of decision making;

2) Identification of anomalies as predicted patterns fail to unfold;

3) Acting effectively in novel situations.

JDCC understanding is that it is imperative to preserve the human cognition dimension of command within command and control processes. What are emphasised in the JDCC EBO paper are the importance of the correct use of information, the sharing of planning, knowledge and SA through network enabled capabilities, and decision superiority. Decision superiority is described as achieved through a combination of the decision-maker's own information position (measure of information gleaned relative to that needed), and the degree of control over the adversary's information position. This preserves operational flexibility, enables courses of action to be developed ahead of the adversary and maintains the initiative in the battle-space.

From this human-centric concerned view-point, emerging military doctrine, and the battlefield digitisation S&T community see ubiquitous computing and information technology as clearly facilitating capability rather than as automating complex decision making and presenting barriers to human involvement. Information technology is expected to provide transparency of information and of the command and control architecture, and expose the inadequacies of current processes associated with functional and physical separation of force components. JDCC recognise that changes to mission command philosophy and operational processes towards coherent processes and effects-based operations will need to be made to fully exploit this technology. But rather than the desynchronised planning currently experienced aimed at providing deconfliction of force components, future network web-based information technology and seamless joint C4ISTAR is expected to enable capability and increase battle-space exploitation. This is to be achieved by providing SA and decision/knowledge superiority, balanced information gathering and decision tempo, collaborative planning and command intent, timely integration and synchronisation of components, and optimised tempo and minimised friction. The purpose of command intent is for subordinate commanders to be fully conversant with the nuances of directives, with level of understanding sufficient to enable rapid and pre-emptive decisions without recourse to senior commanders, as opportunities arise. Command intent and the provision of combat ID capability are seen as key products.

2.5 Human Role Issues

From this perspective on the required military capability, it is possible to identify some issues and concerns for human roles with intelligent and automated systems. Automation often reduces human

involvement in functions and replaces human roles. This reduction or replacement presents new and often unanticipated risks for command and control of military force.

Situation awareness. Good SA and context sensitivity are recognised as key pre-requisites for good military decision making. But how do we ensure the proper human appreciation of military situations when highly automated systems are used for gathering, distilling and disseminating information (i.e. digitisation)? Networked web-based technology should broaden and increase access to relevant information. But this potential wealth of information will need to be intelligently managed (filtered, integrated, classified, analysed, prioritised) to prevent information overload and decision delays.

Tempo. Military decisions are often time sensitive, and they need to be timely to maintain the correct tempo of action. Delaying action decisions while waiting for more situation information is sometimes not an option. How do we ensure that proper control of tempo and decision flow is achieved with highly automated systems?

Agility. How do we ensure that with highly automated systems, the military can provide the capability and levels of agility required for contemporary and envisioned operations, and needed for unanticipated future operations? What is the most appropriate and effective control architecture and functional, component and process structure for decision agility and adaptiveness when working with semi- and highly automated computer information and decision systems. How can intelligent systems help provide the correct balance of reactive and proactive responsivity?

Trust. Can intelligent systems be trusted always to make the correct decisions autonomously? Do we need to have rules that constrain and limit of the use of advanced automation technology in military systems? How do we ensure that humans have effective roles with highly automated systems, that the decision-maker's do not become dependent on computers for complex decisions, and can retain their human effectiveness (skill fade, automation dependency)? Do we understand how computers should be used to assist and interacted with to support humans in the performance of critical decision functions, without developing dependency, overtrust and incorrect influence?

Validity. How can we validate the proper and most effective integration of advanced automation with reference to critical military capabilities, such as acting effectively in novel situations?

3.0   COGNITION

It is generally believed that human involvement in critical decisions on the use of military force is paramount. Humans should decide the objectives, strategy, influences and intended effects of military force. Cognition provides humans with essential mental abilities for judging and deciding the appropriate and effective use of military force. This includes the ability to appreciate novel, complex and dynamic situations involving uncertainty and time pressure, to visualise complex new patterns, relationships and interactions in situations, and to interpret unstructured and incomplete information. Cognition enables humans to hypothesise and draw inferences about possible consequences of courses of action, to predict patterns, and to identify opportunities and alternatives. Cognitive skills provide the ability to anticipate consequences, to make decisions at the right tempo, and visualise how situations might evolve over time.

3.1   Alternative Frameworks

It would be useful to have a framework of cognition suitable for consideration of human roles in advanced automated systems. Frameworks for thinking about cognition are legion. Cognition by definition is a broad term used to refer to such activities as thinking, conceiving, reasoning. Cognition concerns any class of mental behaviours where the underlying characteristics are of an abstract nature and involve

symbolising, insight, expectancy, complex rule use, imagery, belief, intentionality, problem-solving etc. One class of frameworks for cognition, often used for function allocation purposes, provides structural descriptions of fundamental stages of internal mental processes, free of context, and caste as a limited capacity information processing system (sensory processes, attentional filtering, central processing, memory and retrieval, action selection). This includes the simple perceive-decide-act, and the air combat-based OODA loop (observe, orient, decide, act). But these provide a vocabulary with weak associations and limited mapping to the functional terminology of the capability framework. An alternative class of pragmatic frameworks considers cognition in context and natural situations. This approach recognises that human performance is constrained by the conditions under which it takes place, and focuses on the variety of human performance and what cognition does, rather than what cognition is and the internal mechanisms for achieving it. This approach is referred to as "situated cognition", "naturalistic decision-making", and "cognition in the wild". Hollnagel (2002) summarises this approach in the following terms:

- Cognition is distributed across multiple natural and artificial cognitive systems and not confined to an individual cognitive agent.

- It is part of a stream of activity and not confined to a short moment in response to an event.

- Sets of active cognitive systems are embedded in a social context which constrains their activities and provides resources.

- The level of activity has transitions and evolutions and is not constant.

- Most of the activity is supported by something or someone beyond the individual cognitive agent i.e. by an artefact or another agent.

3.2   Cognition and Control

Rassmussen (1986) has provided an error-based classification of behaviour with skill, rule, and knowledge (SRK) levels of performance, or generic decision layers, corresponding to decreasing levels familiarity with the task or environment, or expertise. This approach recognises that both goals and experience drive behaviour. It recognises that humans should be allowed to be flexible and variable, and that error observability and reversibility are important features for safe task and system design. The SRK framework can provide a means of considering the potential scope and contribution of levels of automation aiding (Taylor, 1997). Advanced automation can be thought of as mostly replacing skill-based behaviour (well-defined, highly structured domain), as both replacing and supporting some rule-based behaviour, and as supporting some knowledge-based behaviour (ill-defined, unstructured domain). The question of how control is passed between SRK levels, has led Hollnagel (1993) to consider the functions necessary to explain the orderliness of human action, with levels of cognitive control. The Contextual Control Model (COCOM) is intended to be applicable to a range of systems, including individuals, joint cognitive systems and complex socio-technical systems. COCOM has three main constituents.

(1)   Competence - set of possible actions, responses available to apply to a situation.

(2)   Control – a continuum of modes to characterise the orderliness of performance with different planning horizons and levels of goals (scrambled, opportunistic, tactical, strategic).

(3)   Constructs – what the system knows about the situation in which action takes place.

Scrambled is the least efficient, irrational trial and error, with little if any cognition involved. Strategic has the longest planning horizons and looks ahead at higher level goals. It is the most efficient, and optimal in terms of being able to control a situation, with multiple goals, but requires so much cognitive effort that it is difficult to sustain for extended periods of time. Humans usually function in opportunistic (responding to salient features, heuristics, with limited planning) and tactical mode (following a known procedure, limited planning) with an efficient equilibrium between feedforward and feedback, pro-active and reactive control.

In the Extended Control Model (ECOM), Hollnagel (2002) recognises that performance can take place at several levels simultaneously, or as multiple concurrent control loops; some are closed-loop and reactive, some are open-loop and pro-active. The idea of layered control protocols has been used previously to model the management of interlocutor speech dialogue using principles of perceptual control theory (Taylor, 1992). Four levels of activity or control loops are currently distinguished, namely controlling, regulating, monitoring, and targeting or goal-setting. Figure 2 provides an illustration derived from Hollnagel (2002) with acknowledgement. These levels correspond to the layers in the SRK generic decision model. The difference with SRK is that the ECOM layers concern different aspects of system control, and human-machine system development.
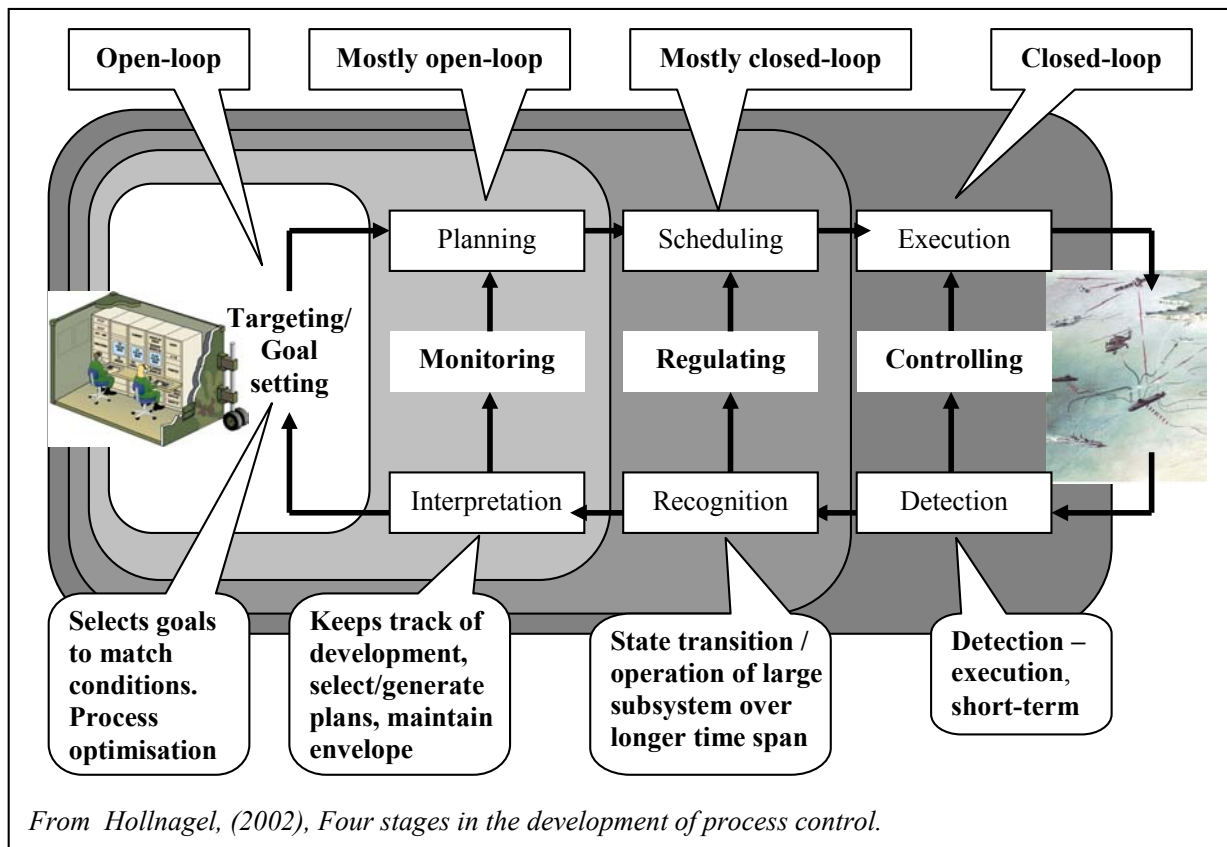

From Hollnagel, (2002), Four stages in the development of process control.

Figure 2: Illustration of Eric Hollnagel's Extended Cognitive Control Model.

The COCOM/ECOM framework is intended to support understanding of human roles and collaborative functioning with intelligent automation and joint cognitive systems. Automation take-over can be seen to be applicable to most controlling activities. These would be activities that for the skilled operator are performed automatically, without much effort, but which may breakdown under stress or need regulating when the situation changes. Problems can arise when the automation makes it difficult to maintain the appreciation of the situation for other levels of control. Automation of regulation is common, mostly of the closed-loop variety. Monitoring is mostly open-loop performance and automation is less common. Automation has had little impact on targeting. Regulating and controlling can be regarded as working with relatively well-defined, organised and ordered information and problem domains, more amenable to simple automation take-over. In contrast, monitoring and particularly targeting concern ill-defined, unstructured information and problem domains and require advanced automation and support, following the "complementarity" principle. The problem is how to support targeting and monitoring or knowledge-based behaviour with automated decision support, without unduly biasing the decision-making, when the

situation is novel and involves uncertainty, when the correct decision is difficult to determine *a priori.* The balancing of rules and knowledge based behaviour, coupled with preparation and planning, is well illustrated in the mission control activities that averted disaster on the Apollo 13 mission (Krantz, 2000). Decision systems can help by providing structuring of the information and knowledge-base for decision-making (e.g. CommonKADS), and by fast automating of complex reasoning heuristics (Shadbolt et al, 2000). Current approaches for the design of decision systems seeking to avoid automation decision bias, favour critic cognitive interfaces. Here the decision is arrived at through a dynamic human-computer cognitive interaction dialogue, involving hypothesising and negotiation, where prioritised options are posed with measured saliency, with estimated confidence, and justified rationally i.e. as provably correct (Silverman, 1992; Taylor and Dru-Drury, 2001).

3.3   Cognition and Capability

There are several conjunctions between the concepts of COCOM/ECOM theory of cognitive control and the UK MOD Defence Capability Framework, reported earlier. The coherency of the DFK command and inform dimensions, and the command intent initiatives are relevant for strategic and tactical cognitive control modes, and for targeting and monitoring, being open-loop, feed-forward and proactive. Shared situation appreciation, collaborative planning and command intent and the enabling of decision agility and combat ID in operations are important for a balance of tactical and opportunistic cognitive control modes. This enables cognition and control at the levels of monitoring, regulating and controlling with an appropriate and efficient mix of feed-forward and feedback, open-loop and mostly closed-loop cognitive control. The interaction between the DFK command and inform axis and the sustain and protect dimensions is understandable in terms of strategic and tactical cognitive control, with proactive, planful, feed-forward performance. The coupling of detailed preparation and planning coupled with high situation awareness and command intent, are the foundation of decision agility and superiority in novel situations, and this makes sense in terms of cognitive control requirements. Preparation and planning enable predictable decisions to be made ahead of time and frees cognitive resources to control the situation both tactically and efficiently. An illustration of the ECOM model operating to provide decision agility through strong CCII, command intent, planning and situation awareness is given in Figure 3.
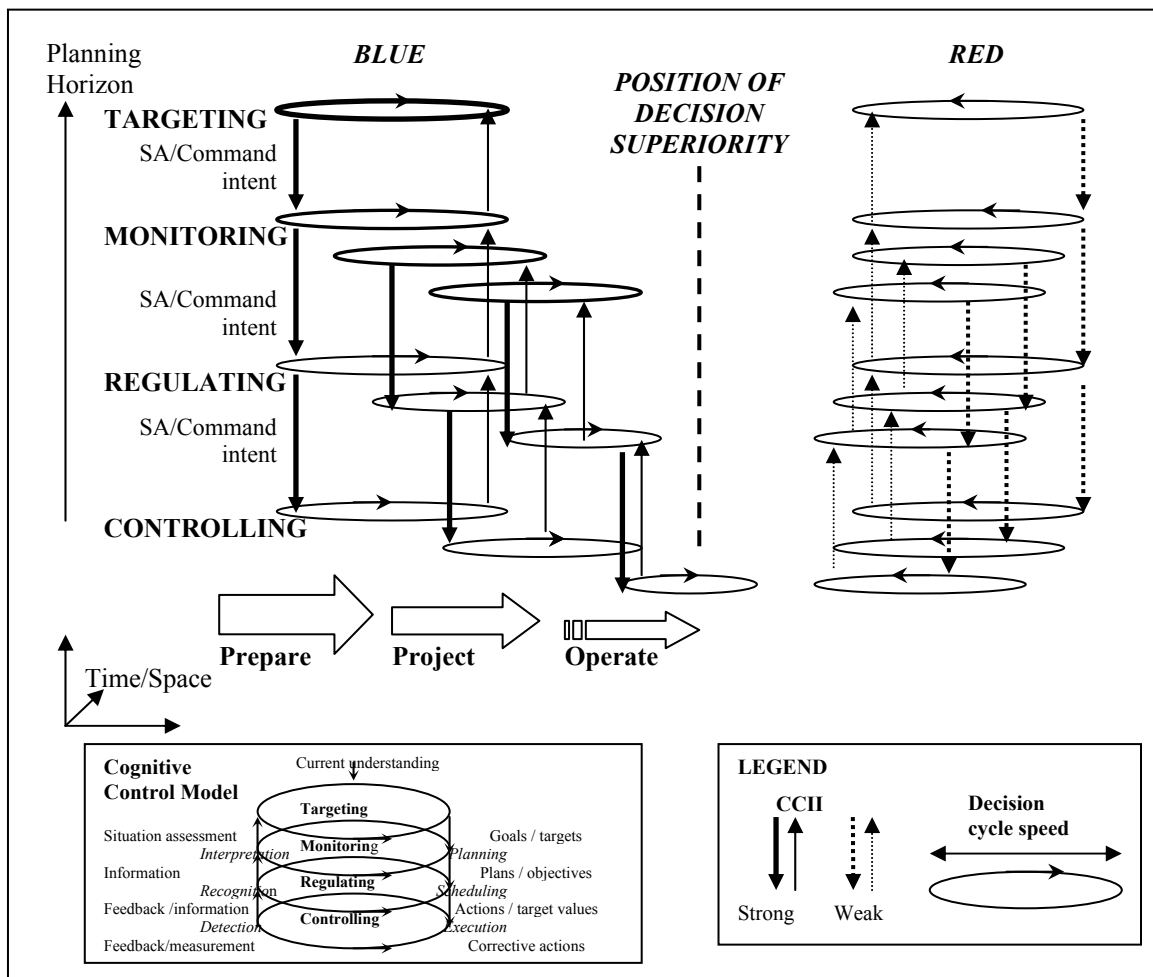
Figure 3: Decision Agility through Cognitive Command and Control.

4.0   AUTONOMY

4.1   Intelligent Systems

The term autonomy has been introduced to describe the bounding of functioning and decision authority of advanced automation and intelligent decision systems. Autonomy can be defined simply as the capability to make decisions. Thus, levels of autonomy can be considered in terms of the freedom to make decisions, considering constraints on decision-making (limitations, boundaries, rules, regulations), decision-making abilities (authority, responsibility, competency), and the capability to make different kinds of decisions (classes, functions, levels). Defining intelligent systems is probably done best with reference to functional rather than technological criteria. The traditional way of defining intelligent automation is in terms of the criteria of the Turing test i.e. systems capable of behaviour indistinguishable from humans. To pass the Turing test intelligent automation probably needs the following attributes (Geddes and Shalin, 1997):

- Active collector of goals
- Goal driven
- Reasons at multiple levels
- Context sensitive communicator
- Learns from experience

In addition, the capability for collaborative working with other agents, including humans, is a goal for intelligent automation. Taylor and Reising (1998) noted that in order to work collaboratively with humans, intelligent automation probably requires a functional architecture with the following attributes:

- A model of human decision making and control abilities,

- The ability to monitor human performance and workload through behavioural and physiological indices,

- The ability to predict human expectations and intentions with reference to embedded knowledge of mission plans and goals.

The building of trust between the operator and the computer automation system has been identified as a key issue in enabling the capability of cognitive automation. Trust is built when consistency and correctness is observed in the computer system's decisions and actions. Two important guidelines for building trust have arisen (Reising, 1995):

Define the Prime Directives. These are overall governing rules which bound the behaviour of the aiding system, and yet provide a logical structure for aiding system to act in a rational and reliable manner, avoiding arbitrary behaviour, so that the human does not experience any surprises e.g. Asimov's Laws of Robotics.

Specify the Levels of Autonomy. These also bound the behaviour of the aiding system by limiting its decision authority for the performance of specific sub-functions to a set of system configurations specified and set by the operator.

4.2   Levels of Automation

Sheridan and Verplanck (1978) first proposed 10 possible levels of allocation of decision-making tasks, or levels of autonomy, between humans and computers. More recently, Parasuraman, Sheridan and Wickens (2000) have considered the application of automation to a four-stage model of independent information processing functions (information acquisition, analysis, and decision selection and action implementation). In doing so, they have sought to apply a revised set of levels of autonomy. Both the original and the revised redefined levels of automation are listed for comparison in Table 3.

Table 3: Sheridan's Original and Revised Levels of Automation

Levels of Automation of Decision and Action	
1978 Original Set *Sheridan and Verplanck (1978)*	**2000 Revised Set** *Parasuraman, Sheridan & Wickens (2000)*
10. Computer does the whole job if it decides it should be done, and if so tells human, if it decides human should be told.	10. The computer decides everything and acts autonomously, ignoring the human.
9. Computer does the whole job and tells human what it did. The computer decides whether or not human should be told.	9. The computer informs the human only if it, the computer, decides to.
8. Computer does the whole job and tells human what it did only if human explicitly asks.	8. The compute informs the human only if asked.

Levels of Automation of Decision and Action	
1978 Original Set *Sheridan and Verplanck (1978)*	**2000 Revised Set** *Parasuraman, Sheridan & Wickens (2000)*
7. Computer does the whole job and tells human what it did.	7. The computer executes automatically, then necessarily informs the human.
6. Computer selects action, informs human in plenty of time to stop it.	6. The computer allows the human a restricted time before automatic execution.
5. Computer selects action and implements it, if human approves.	5. The computer executes the suggestion if the human approves.
4. Computer selects action and human may or may not do it.	4. The computer suggests an alternative.
3. Computer helps determine the options and suggests one, which human may or may not follow.	3. The computer narrows the selection down to a few.
2. Computer helps by determining the options.	2. The computer offers a complete set of decision alternatives.
1. Human does the whole job up to the point of turning it over to the computer to implement.	1. The computer offers no assistance. The human must make all the decisions and actions.

4.3   Pilot's Associate LOA

In the 1980's, the DARPA/USAF Pilot's Associate (PA) program provided a practical implementation of intelligent pilot aiding based on prime directives and levels of autonomy (see Taylor and Reising, 1998, for review). A summary of the PA design approach underpinning the levels of autonomy is shown in Table 4. PA design was guided by a top-level operational philosophy based on the pilot being in charge. The goal of the PA was to provide consistently correct information, and to aid the pilot's decision making by helping to manage workload, reduce confusion, and simplify tasks. This led to the philosophy of the PA as an intelligent subordinate to the pilot, with specific capabilities for decisions and actions. These top level requirements led to specific operational relationships (ORs) for discrete PA sub-functions interactions, with increasing degrees of automation and autonomy. From these ORs, pilot selectable levels of autonomy (LOA) were obtained for groups of functions governed by the required pilot operational relationship and interaction (Krobusek et al 1989). Four discrete LOA modes were proposed, namely: Inactive, Standby, Advisor, Assistant, Associate. Each LOA mode was associated with tailorable functional clusterings for flexible responding to avoid too rigid automation imposed by design. These modes were aimed to provide a bounded, communicable structure for delegated levels of authority, minimising mode confusion, and building trust and confidence. HF research indicates that the required control structure should be cognitively simple, and not complex. Pilots tend to view computer autonomy simply as either automatic, with or without status feedback; semi-automatic, telling what will happen and asking permission to proceed; or advisory, providing information only.

Table 4: Pilot's Associate Design Approach for Levels of Autonomy

Pilot's Associate Design			
Operational Philosophy	**PA Capabilities**	**Operational Relationships**	**Modes for Levels of Autonomy**
The pilot is in charge - i.e. the pilot shall always have the capability to act according to his desires. PA's plans may be: Approved or rejected explicitly with little effort Approved or rejected pre-mission Approved or rejected implicitly by pilot action, or Ignored with predictable results The PA must operate in a predictable manner. The PA is required to monitor the pilot, not the other way around. The PA must notify the pilot of key mission events (as defined and set by the pilot). The effort required of the pilot to control the PA must be less than the effort saved by the PA. PA shall save more effort for the pilot than it creates - it shall be responsive to the pilot and not demanding of his resources.	PA could not act on its own. PA could make recommendations. PA could take actions based on pilot discretion. PA could fly the aircraft tactically on autopilot. PA could take action based on interpreting pilot intent. PA could diagnose malfunctions, identify mis-communications, & determine correct response. PA could deal with ambiguities in human speech in the context of the mission.	OR2. The activity is performed automatically by the PA OR7. PA may perform an action only if various conditions are met. OR6. PA has been given authority to perform, but with pilot consent. OR5. PA may prompt the pilot. OR4. PA may remind the pilot. OR3. PA may remind the pilot, if the pilot asks, or has authorised such. OR1. The pilot must perform the activity	**Associate.** In Associate mode, under full dynamic function allocation (DFA), the proposed system maintains advisory functions and accepts pilot allocated tasks, but also takes over tasks as the context demands. **Assistant.** In Assistant mode, the PA would maintain advisory functions and also assume responsibility for tasks explicitly allocated to it by the pilot. **Advisor** **Standby** **Inactive**

4.4 Cognitive Cockpit Pact

More recently, the UK DERA *Cognitive Cockpit* project on technology proof-of-concept, has identified a limited set of four automation assistance, or variable autonomy, levels for integrating knowledge-based decision support with adaptive automation (Taylor et al, 2001; Taylor, 2001a; Taylor, Brown and Dickson 2002). This policy for (pilot) authorisation and control of tasks, or PACT framework, is used in conjunction with concepts for a *tasking interface* manager whereby mission functions or tasks are assigned for computer automation or computer support (Bonner, Taylor and Miller, 2000). The PACT framework is summarised in Table 5, and illustrated in Figure 3.

Table 5: Bonner-Taylor Pact System

| Primary Modes	Levels	Operational Relationship	Computer Autonomy	Pilot Authority	Adaptation	Information on Performance
AUTOMATIC		Automatic	Full	Interrupt	Computer monitored by pilot	On/off Failure warnings Performance only if required.
ASSISTED	4	Direct Support	Advised action unless revoked	Revoking action	Computer backed up by pilot	Feedback on action. Alerts and warnings on failure of action.
	3	In Support	Advice, and if authorised, action	Acceptance of advice and authorising action	Pilot backed up by the computer	Feed-forward advice and feedback on action. Alerts and warnings on failure of authorised action.
	2	Advisory	Advice	Acceptance of advice	Pilot assisted by computer	Feed-forward advice
	1	At Call	Advice only if requested.	Full	Pilot, assisted by computer only when requested.	Feedforward advice, only on request
COMMANDED		Under Command	None	Full	Pilot	None performance is transparent.
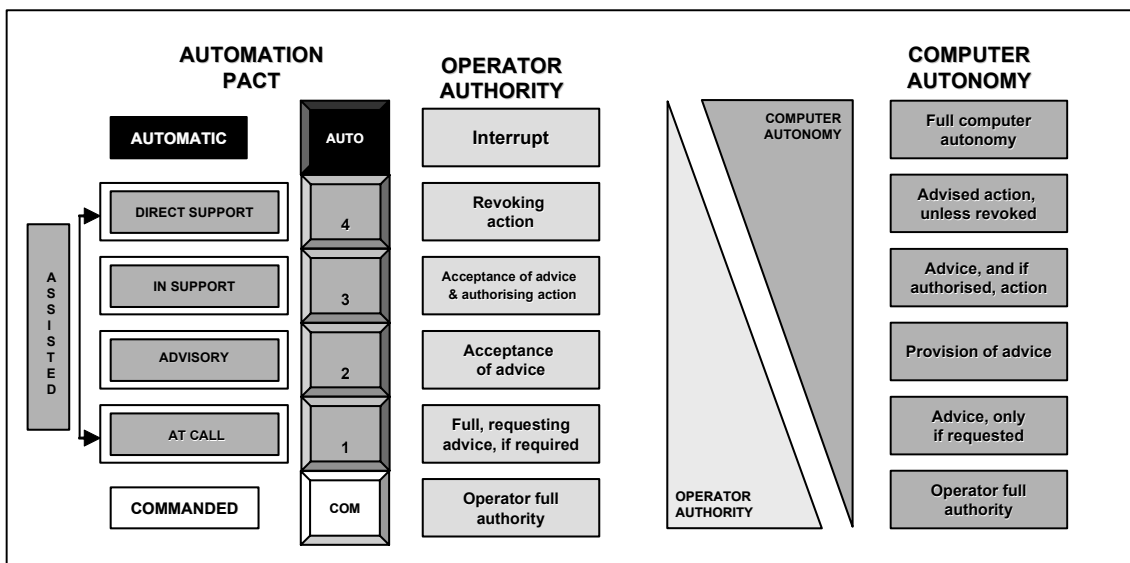
Figure 3: PACT Progression of Operator Authority & Computer Autonomy.

The PACT system succeeds in reducing the number of automation or autonomy modes required to three – namely, fully automatic, assisted or pilot commanded - with a further four secondary levels nested within the semi-automatic, assisted mode, which can be changed adaptively or by operator/pilot command. Military terminology is used derived from categories of support in Army land forces military operations (At Call, Advisory, In Support, Direct Support). This is to afford usability and compatibility with military user cognitive constructs, schemata and models. It provides realistic operational relationships for a logical,

practical set of levels of automation, with progressive operator authority and computer autonomy supporting situation assessment, decision making and action. Mission functions and tasks, at different levels of abstraction allocated individually or grouped in related scripts or plays, can be set to these levels in a number of ways:

- Pre-set operator preferred defaults.

- Operator selection during pre-flight planning.

- Changed by the operator during in-flight re-planning, probably using Direct Voice Input commands.

- Automatically changed according to operator agreed, context-sensitive adaptive rules.

The setting of functions and tasks to PACT levels is described as the creation of personal binding *contracts* between the operator and the computer. This is to provide the operator with implicit if not explicit control, and to engender trust through understanding of automation functioning. Fighter pilots develop similar inter-personal contracts in planning control of multi-aircraft manoeuvres in co-operative air defence. PACT autonomy contracts govern and constrain the behaviour of the computer according to rules of operation (context, resources). Figure 4 illustrates a set of mission functions and tasks with PACT contractual autonomy levels arranged along a timeline in a hypothetical task network.
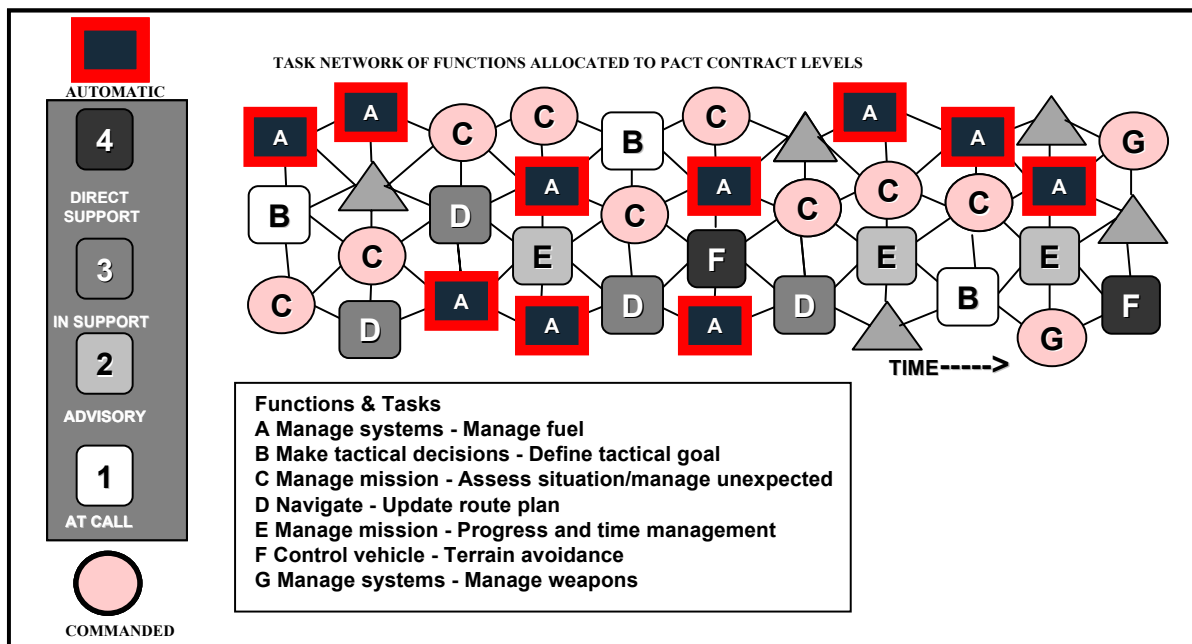
Figure 4: Task Network of Functions and Tasks Set to PACT Contract Levels.

The PACT system is designed to support the pilot's cognitive work. The support ranges from providing advice to providing action. The cognitive work required can be represented in terms of a SRK perception-assessment-targeting-execution decision ladder using state flow transition diagrams. Control task analysis (Vicente 1999) can been used to identify the structure of the cognitive work performed by the pilot and by automation at each PACT level. Figure 5 and 6 provide examples of the control task analysis for PACT Level 2 Assisted-Advisory and PACT Level 3 Assisted – In Support, expressed in decision ladder terms (Taylor, 2001b). This enables estimates of the resultant or residual pilot cognitive load for different degrees of pilot critical involvement. The two analyses contrast different degrees of expert pilot involvement ranging from immediate acceptance (PACT 3), to critical acceptance (PACT 2), as might occur under time pressure. In both instances, full independent high level analysis of the implications for

situation status and goal evaluation is not represented as taking place. A fuller human decision analysis might occur with novel situations or with novice operators during training, and before the development of trust in the computer assistance.

Figure 5: Control Task Analysis for PACT Level 2 Assisted Advisory.

Figure 6: Control Task Analysis for PACT Level 3 Assisted in Support.

4.5   UAV/UCAV Autonomy

Initially developed for individual pilot cockpit decision support, subsequently the PACT system has been applied in research on the management of multiple UAVs from manned cockpits, to help reduce pilot cognitive workload, and it is seen as equally applicable to control of UAVs from ground stations (White, 2002). Recent DARPA sponsored work on air vehicles (AV) indicates similar autonomy solutions. The DARPA UCAV Advanced Cognition Aids Integration project for target engagement and multiple AV identifies four levels of autonomy, namely automate, exception (informs immediate action, OK or revoke), consent (authorisation required), manual (Leahy, 2001). The DARPA ICAV Intelligent Control of Unmanned AV project on mixed initiative distributed intelligence architecture for UAV operations identifies four levels of authorisation, namely autonomous, veto (proposal implicitly accepted after time out), permissive (proposal implicitly rejected after time out), manual (Elmore, 2001). For future envisioned UCAV operations, involving real-time, multiple (group) collaborating autonomous vehicles in joint operations with manned platforms, it seems likely that autonomous control levels will need extending beyond human command and computer support, to cover classes of autonomous complimentary, co-ordinated and co-operative planning and interactions.

4.6   Multi-Agent Adjustable Autonomy

Autonomy issues and implementation solutions have been addressed in work on multi-agent intelligent systems for problem solving in complex dynamic environments (Barber et al 2000). Mixed-initiative systems, dynamic adaptive autonomy and adjustable autonomy have been proposed to enable multi-agent systems to perform effectively with adaptability and flexibility. In the context of single-agent to human-user interaction, autonomy has generally been viewed as freedom from human influence. But for multi-agent systems, where the human user may be remote from operations, autonomy becomes a matter of the agents self-direction and goals, and the capability to dynamically form, modify or dissolve the agent organisation into goal-oriented, problem-solving groups. The degree of autonomy is considered to be implicit or explicitly linked to individual goals, and focuses on the decision making process used to determine how a goal is pursued free from intervention, oversight, or control by another agent (technical or human). Autonomy with respect to goals can be considered to be on a variable scale:

- Consensus or distributed control through consensus (working as a team member, sharing decision-making control equally with all other decision-making agents, all with equal authority).

- Master control (makes decisions alone, may communicate or give orders to other agents with authority).

- Locally autonomous (makes decisions alone, only agent with authority).

- Command-driven or centralised control (makes no decisions about how to pursue goals, has authority, but must obey orders given by another agent).

Table 6 adds these agent autonomy levels to the PACT levels with a summary of the responsibilities in cognitive control model terms of advising and performing targeting, monitoring, regulating and controlling. A representation of autonomy levels in terms of the COCOM/ECOM framework is shown in Figure 7. This enables consideration of the flow and transitioning of control in functional context rather than in terms of internal decision-making processes. Further exploitation of the PACT framework can be suggested as follows:

- Assign functions to multi-agent resources in CCII. Use PACT levels to define operational relationships.

- Assign a broad range of inactive reserve functions and operational relationships to PACT Level 1 Assisted at Call i.e. pre-set at PACT Levels 2, 3, 4.

- Use PACT to define multi-agent support inter-relationships at the Master Control autonomy level.

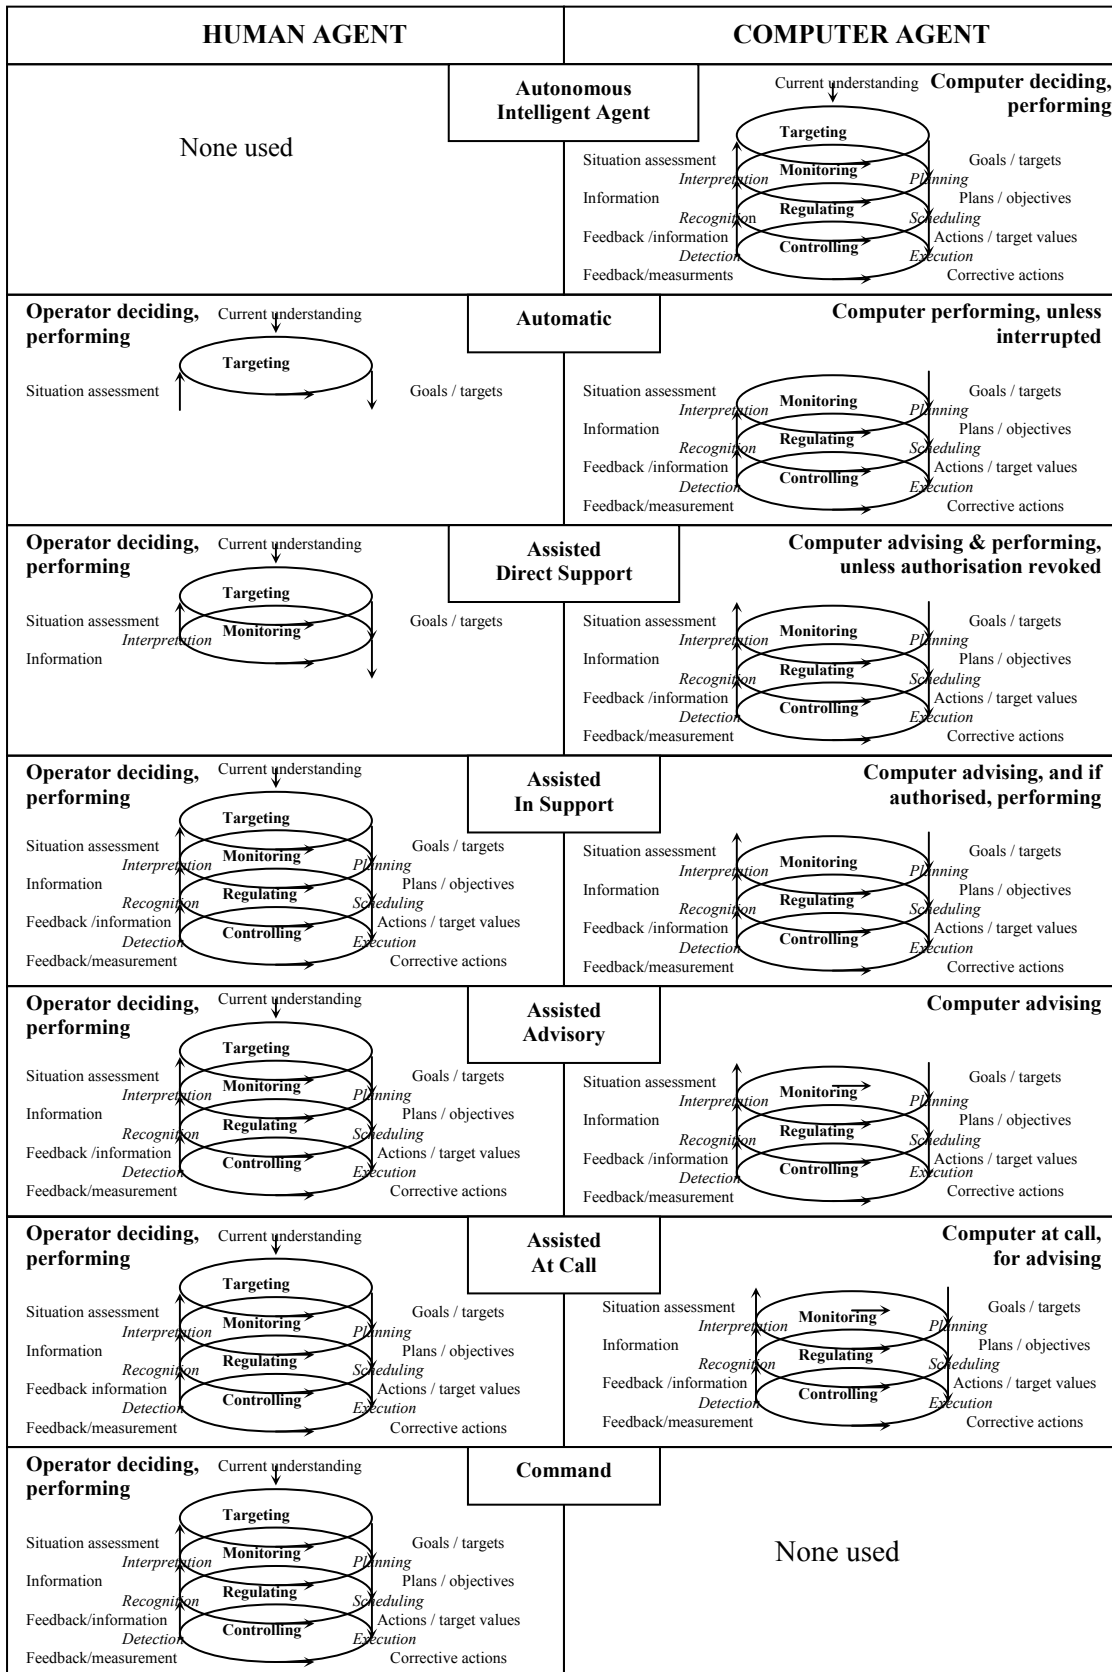- Use PACT agents to organise and filter prioritised information in CCII for command intent and SA.

Figure 7: Levels of Autonomy within the Extended Control Model Framework.

Adjustable autonomy gives the agent architecture the ability to adapt their problem-solving to situations particularly in domains with unreliable communications and the possibility of agent failure, high degrees of uncertainty and resource contention needing distribution of tasks and co-ordinated planning to resolve conflicts. Distributed problem solving structures are generally thought to perform faster for complex tasks, when operating under uncertainty and changes in the environment, when few resources are shared, and when communication is unreliable. Centralised structures perform faster for simple tasks, when many resources are shared, when communication is reliable, and when there is no requirement to negotiate. Autonomy level agreements and communication protocols, joint intentions, and employing conventions for explicit commitment to specific interaction styles are considered necessary to establish reliability and trust. A central problem in adjustable autonomy is the determination of whether and when transfers of control to the operator/user should occur (Scerri et al, 2001). The transfer of control from agent to human is believed to require a balancing of the costs of interrupting a human user with the benefits for highest quality decision making when the human has superior decision-making expertise. One technique proposes that transfer should occur when the expected utility of transfer is greater than that of retaining the decision-making. Another forces the agent to relinquish and transfer control if the uncertainty is high. Others transfer if any incorrectness in the agents decision can cause significant harm, if the agent lacks decision-making capability, or on the basis of thresholds of learnt rules. In multi-agent applications, cognitive strategies are needed for reasoning with adjustable autonomy in the operating context (situated autonomy) to provide the correct co-ordination, reordering and scheduling and to balance the costs, benefits, uncertainty and implications within the multi-agent group (Hexmoor, 2000).

There is considerable potential for read-across for control architectures from cognition and joint cognitive systems for the control of distributed multi-agent systems. They use decision resources efficiency and enable the decision agility and adaptiveness needed for the manouevrist approach to military problem-solving. The use of cognitive control models will increase the transparency of control architectures and control authority for human user appreciation of the planning and interaction situation during collaborative problem-solving.

Table 6: Adjustable Autonomy Levels for Computer Assistance and Intelligent Multi-Agent Systems

AUTONOMY	TARGETING	MONITORING	REGULATING	CONTROLLING
Consensus Autonomy	Multiple intelligent computer agents	Multiple intelligent computer agents	Multiple intelligent computer agents	Multiple intelligent computer agents
Master Autonomy	Intelligent computer agent	Intelligent computer agent	Intelligent computer agent + Authorised support agents	Intelligent computer agent + Authorised support agents
Local Autonomy	Intelligent computer agent	Intelligent computer agent	Intelligent computer agent	Intelligent computer agent
Automatic/ Commanded Autonomy	Operator	Computer agent performing some interpretation & planning + Operator interrupt	Computer agent performing recognition & scheduling + Operator interrupt	Computer/intelligent agent performing detection & execution agent + Operator interrupt
Assisted Direct Support	Operator	Operator authorising + Computer agent performing some interpretation & planning	Operator authorising + Computer agent performing recognition, & scheduling	Operator authorising + Computer agent performing detection & execution
Assisted In Support	Operator	Operator performing + Optional computer agent advising & performing some interpretation & planning	Operator performing + Optional computer agent advising & performing recognition & scheduling	Operator performing + Optional computer agent advising & performing detection & execution

AUTONOMY	TARGETING	MONITORING	REGULATING	CONTROLLING
Assisted Advisory	Operator	Operator performing + Computer agent advising interpretation & planning	Operator performing + Computer agent advising recognition & scheduling	Operator performing + Computer agent advising detection & execution
Assisted At Call	Operator	Operator + Optional computer agent	Operator + Optional computer agent	Operator + Optional computer agent
Command	Operator	Operator	Operator	Operator

5.0   CONCLUSIONS

A number of fundamental questions and key issues are identified concerning the role of humans in advanced automated and intelligent systems. In particular, there is uncertainty over how to optimise the use of human and computer decision resources, while preserving a human-centric system. These matters need to be understood in the context of the changing capability requirement responding to new military problem-solving challenges. Important changes are being made in the way in which military force is to be used in the future with the introduction of effects-based approach to the planning and conduct of joint operations. This will be enabled by network CCII, and will provide shared planning and situation appreciation, command intent and Combat ID. The prime reason for human involvement in military decision-making seems unquestioned – human knowledge and experience provides unique capability to analyse and think ahead in uncertain and novel situations. The challenge is to provide information and decision systems that protect and preserve the human user's key role, and that augment and enhance the user's cognition rather than replaces the user in complex decision making. Recent developments in theory of cognition provide pragmatic approaches that are likely to improve understanding of the human factors issues, problems and solutions of human-computer collaboration. In addition, new approaches to the use of automation propose adjustable levels of computer autonomy with a strong socio-technical and cognition basis. These seem likely to provide sensible architectures for distributed, multi-agent intelligent systems that can be more readily appreciated by human users than traditional automation approaches. New applications of PACT to multi-agent CCII are proposed.

6.0   REFERENCES

[1]   NATO RTO 1998. Collaborative Crew Performance in Complex Operational Systems. Papers presented at the RTO Human Factors and Medicine Panel (HFM) Symposium held in Edinburgh, United Kingdom, 20-22 April 1998. NATO RTO-MP-4, AC/323(HFM)TP/2, ISBN 92-837-1008-8. NATO Research and Technology Organisation, Neuilly-sur-Seine Cedex. December 1998.

[2]   Reising, J., Taylor, R.M. and Onken, R. 1999. The Human-Electronic Crew: The Right Stuff? AFRL-HE-WP-TR-1999-0235. US Air Force Research Laboratory, Human Effectiveness Directorate, Wright Patterson AFB, Dayton, Ohio.

[3]   NATO RTO 2002. Human Factors in the 21st Century. Proceedings of the RTO Human Factors and Medicine Panel (HFM) Specialists' Meeting held in Paris, France, 11-13 June 2001. RTO-MP-077 AC/323(HFM-062)TP/38, ISBN 92-837-0021-X. NATO Research and Technology Organisation, Neuilly-sur-Seine Cedex. May 2002.

[4]   Fitts, P.M. 1951. Human Engineering for an Effective Air Navigation and Traffic Control System. Washington DC. National Research Council.

[5]   Beevis, D., Essens, P., and Schuffel, H. 1996. (Eds) Improving Function Allocation for Integrated Systems Design. CSERIAC SOAR 96-01. Crew Systems Ergonomics and Analysis Centre, Wright Patterson Air Force Base, Ohio.

[6]   Sheridan, T.B. 1996. Allocating Functions Among Humans and Machines. In Beevis, D., Essens, P., and Schuffel, H. (Eds) Improving Function Allocation for Integrated Systems Design. CSERIAC SOAR 96-01 pp. 7-27. Crew Systems Ergonomics and Analysis Centre, Wright Patterson Air Force Base, Ohio.

[7]   Taylor, R.M., Brown, L., and Dickson, B. 2002. From Safety Net to Augmented Cognition: Using Flexible Autonomy Levels for On-Line Cognitive Assistance and Automation. In, Proceedings of NATO RTO Human Factors and Medicine Panel Symposium on Spatial Disorientation in Military Vehicles: Causes, Consequences and Cures, La Coruna, Spain, 15-17 April 2002. Paper No. 27. NATO Research and Technology Organisation, Neuilly-sur-Seine Cedex (In Press).

[8]   Krantz, G. 2000. Failure is Not an Option. New York, Berkley.

[9]   NATO NSA/AI, 2000. Application of Human Engineering to Advanced Aircrew Systems. STANAG 3994 AI, Edition 2, NATO Standardisation Agency, Air Board Aircrew Integration Working Group, NATO HQ, Brussels.

[10]   Beevis, D. 1996. Human Functions and System Functions. In Beevis, D., Essens, P., and Schuffel, H. (Eds) Improving Function Allocation for Integrated Systems Design. CSERIAC SOAR 96-01, pp. 63-79. Crew Systems Ergonomics and Analysis Centre, Wright Patterson Air Force Base, Ohio.

[11]   Hollnagel, E. and Woods, D.D. 1983. Cognitive Systems Engineering: New Wine in Old Bottles. International Journal of Man-Machine Studies, Vol. 18, pp. 583-600.

[12]   Hollnagel, E. 1997. Control versus Dependence: Striking the Balance in Function Allocation. In, M.J. Smith, G. Salvendy, and R.J. Koubek (Eds), Design of Computing Systems. Advances in Human Factors/Ergonomics, 21B, pp. 243-246. Amsterdam: Elselvier.

[13]   Kantowitz, B.H. and Sorkin, R.D. 1987. Allocation of Functions. In G. Salvendy (Ed). Handbook of Human Factors, pp. 365-369. New York: Wiley.

[14]   JDCC, 2002. A Concept for Effects Based Operations. JDCC/7/11/7/1, UK Ministry of Defence Joint Doctrine and Concepts Centre, Shrivenham. 20 May 2002.

[15]   Hollnagel, E. 2002. Cognition as Control: A Pragmatic Approach to the Modelling of Joint Cognitive Systems. Special Issue of IEEE Transactions on Systems, Man and Cybernetics A: Systems and Humans – Model-Based Cognitive Engineering in Complex Systems (In Press).

[16]   Rassmussen, J. 1986. Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering. New York: North Holland.

[17]   Taylor, R.M. 1997. Human-Electronic Crew Teamwork: Cognitive Requirements for Compatibility and Control with Dynamic Function Allocation. In, M.J. Smith, G. Salvendy, and R.J. Koubek (Eds), Design of Computing Systems. Advances in Human Factors/Ergonomics, 21B, pp. 247-250. Amsterdam: Elselvier.

[18]   Hollnagel, E. 1993. Human Reliability Analysis: Context and Control. London: Academic Press.

[19]   Taylor, M.M. 1992. Principles for Integrating Voice I/O in a Complex Interface. In AGARD-CP-521. Paper 21, pp. 1-14. NATO Advisory Group for Aerospace Research and Development, Neuilly-sur-Seine.

[20] Shadbolt, N.R., Tennison, J., Milton, N. and Howells, H. 2000. Situation Assessor Support System: A Knowledge-Based Systems Approach to Pilot Aiding. In P.T. McCabe, M.A. Hanson and S.A. Robertson (Eds.), Contemporary Ergonomics 2000, pp. 60-64. London: Taylor and Francis.

[21] Silverman, B.G. 1992. Critiquing Human Error: A Knowledge Based Human-Computer Collaboration Approach. San Diego: Academic Press.

[22] Taylor, R.M. and Dru-Drury, R. 2001. Cognitive Cockpit Systems: The Effects of Reliability and Saliency of Aiding Information on Map Situation Assessment. In D. Harris (Ed), Engineering Psychology and Cognitive Ergonomics, Aerospace and Transport Systems, Vol. 5, Chapter 11, pp. 89-97. Aldershot: Ashgate.

[23] Geddes, N.D. and Shalin, V.L. 1997. Intelligent Decision Aiding for Aviation. HFA Report. Swedish Centre for Human Factors in Aviation, Linkoping, Sweden. 8 September 1997. ISSN 1402-7585. ISBN 91-7219-125-2.

[24] Taylor, R.M. and Reising, J.M. 1998. The Human-Electronic Crew: Human Computer Collaborative Teamworking. In Collaborative Crew Performance in Complex Operational Systems. Papers presented at the RTO Human Factors and Medicine Panel (HFM) Symposium held in Edinburgh, United Kingdom, 20-22 April 1998. NATO RTO-MP-4, AC/323(HFM)TP/2, ISBN 92-837-1008-8. Paper 22, pp. 1-17. NATO Research and Technology Organisation, Neuilly-sur-Seine Cedex. December 1998.

[25] Reising, J.M. 1995. Must the Human-Electronic Crew Pass the Turing Test? WL-TR-96-3039, pp. 103-108. Wright Patterson AFB, Ohio. December 1995.

[26] Sheridan, T.B. and Verplank, W.L., 1978. Human and Computer Control of Undersea Teleoperators. Technical Report. MIT Man-machine Systems Laboratory, Cambridge, MA.

[27] 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, Vol. 30, No. 3, pp. 286-297. May 2000.

[28] Krobusek, R.D., Boys, R.M., and Palko, K.D. 1989. Levels of Autonomy in a Tactical Electronic Crewmember. WRDC-TR-89-7008, Wright Patterson AFB, Ohio.

[29] Taylor, R.M., Abdi, S., Dru-Drury, R., and Bonner, M.C. 2001. Cognitive Cockpit Systems: Information Requirements Analysis for Pilot Control of Cockpit Automation. In D. Harris (Ed), Engineering Psychology and Cognitive Ergonomics, Aerospace and Transport Systems. Vol. 5, Chapter 10, pp. 81-88. Aldershot: Ashgate.

[30] Taylor, R.M. 2001a. Cognitive Cockpit Systems Engineering: Pilot Authorisation and Control of Tasks. In R. Onken (Ed), CSAPC'01. Proceedings of the 8th Conferences on Cognitive Sciences Approaches to Process Control, Neubiberg, Germany, September 2001. University of the German Armed Forces, Neubiberg, Germany.

[31] Bonner, M., Taylor, R.M. and Miller, C. 2000. Tasking Interface Manager: Affording Pilot Control of Adaptive Automation and Aiding. In P.T. McCabe, M.A. Hanson, and S.A. Robertson (Eds.), Contemporary Ergonomics 2000, pp. 70-74. London: Taylor and Francis.

[32] Vicente, K.J. 1999. Cognitive Work Analysis: Towards Safe, Productive, and Healthy Computer-Based Work. New Jersey: Lawrence Erlbaum.

[33] Taylor, R.M. 2001b. Cognitive Cockpit Control Task Analysis. DERA Technical Memorandum, DERA/CHS3/6.3/14/7, dated 9 March 2001.

[34] White, A.D. 2002. The Human-Machine Partnership in UCAV Operations. Proceedings of the 17th UAV Bristol Conference, 10 April 2002.

[35] Leahy, M. 2001. DARPA UCAV Advanced Cognition Aids Integration. Paper presented to Unmanned Systems 2001. The 28th Annual Technical Symposium of the Association of Unmanned Vehicles International, Baltimore, July 2001.

[36] Elmore, W. 2001. DARPA ICAV Intelligent Control of Unmanned Air Vehicles. Paper presented to Unmanned Systems 2001. The 28th Annual Technical Symposium of the Association of Unmanned Vehicles International, Baltimore, July 2001.

[37] Barber, K., Goel, A. and Martin, C. 2000. Dynamic Adaptive Autonomy for Multi-Agent Systems. Journal of Experimental and Theoretical Artificial Intelligence. Vol. 12. Part 2, pp. 129-148.

[38] Scerri, P., Pynadath, D.V. and Tambe, M. 2001. Adjustable Autonomy in Real World Multi-Agent Environments. In Agents '01, Proceedings of the 5th International Conference on Autonomous Agents.

[39] Hexmoor, H. 2000. A Cognitive Model of Situated Autonomy. In Proceedings of PRICAI-2000. Workshop on Teams with Adjustable Autonomy, pp. 11-20.