Putting Complexity to Work - achieving effective Human-machine Teaming

Patrick Beautement

The abaci Partnership LLP, Tewkesbury, Gloucestershire, GL20 7TA, UK

patrick@abacipartners.co.uk

Abstract: As people 'put complexity to work' they are increasingly employing machines to help them achieve understanding, make decisions and assist with supporting tasks such as modelling, intervention-monitoring and evaluation. In effect, humans and machines are forming teams - but how effective can these teams be and what are the consequences that follow if they are not? This paper asks the question: what needs to be done to enable machines to make enough sense of complexity so that they can assist people more effectively with this work? This paper first describes a characterisation of complexity which is then used to look at the role of machines in supporting decision-making in these environments. The paper examines the issues around the formation and sustainment of so-called human-machine teams (HMTs). To this end, the theoretical ranges of collaboration and relationships between machines and people are characterised and the nature of dialogue with machines (about complex issues) is examined. From this, a list of desired machine capabilities is derived. Lastly, implications and conclusions are drawn to answer the initial question and comment on the significant issues which arise and on their potentially serious consequences. The author hopes to stimulate debate on this topic with practitioners (about the degree of contribution that it is realistic to expect from machines in the near future) and with technologists about how machine behaviour can be further augmented.

Keywords: Complexity, Decision-making, Machine Cognition, Collaboration, Teaming, Agents, Relationships, Information, Intelligence, Shared Awareness

1 Introduction

We have all experienced situations where, when we ring up a call-centre to sort out a problem, we are told by the assistant that "The computer won't let me do that". This situation is a symptom of one thing - that the computers' model of the world is too simple and has become out of step with the complexities of reality. Often, the only way to correct this situation is to jump over strange hoops to "keep the computer happy", by following the irrelevant sequence of steps that the interface offers, or by 'tricking' the system and finding a work-around.

These experiences are usually frustrating, and sometimes even comical, but they hint at a critical failing. Machines currently don't 'understand' complexity - and

this has important consequences for all of us. As people 'put complexity to work' they are increasingly employing machines to help them achieve understanding, make decisions and assist with supporting tasks such as modelling, intervention-monitoring and evaluation. In effect, humans and machines are forming teams - but how effective can this 'human-machine teaming' be if machines can't be made to appreciate complexity - and when is it appropriate to use machines anyway?
The aim of this paper is to address the question: "**What needs to be done to enable machines to make enough sense of complexity so that they could assist people more effectively with this work**?". The paper will do this as follows.
In Section 2, the paper first provides a Complexity Framework which is used to inform the nature of decision-making in complex environments and to identify the information and capabilities that machines require if they are to help.
In Section 3, the paper then examines the issues surrounding the formation and sustainment of so-called human-machine teaming (HMT). A key metaphor is: "The machine as team player", i.e. a 'purposeful entity' capable of contributing effectively. To this end, the theoretical ranges of collaboration and teaming and the grades of relationships between machines and people are characterised (Section 4). In Section 5, dialogue with machines (about complex issues) is examined and the types of information required are identified. From this, in Section 6, a list of desired machine capabilities is proposed for HMT. Lastly, implications are derived and conclusions drawn to answer the initial question and comment on the significant issues which arise. The author hopes to stimulate debate on this topic with practitioners (about the degree of contribution that it is realistic to expect from machines in the near future) and with technologists about how machine behaviour can be further augmented.

2 Characterising Complexity

To enable the issues around Human-machine Teaming to be examined, the paper proposes that the following Complexity Framework be used to characterise the complexity that we shall be discussing. This Framework is relevant when considering of any set of complex issues, as it identifies aspects common to all.
The emergent phenomena which arise from complex interactions are (mostly in system engineering [1]) often portrayed as unwelcome, chaotic and destructive, but nothing could be further from the truth. In reality, emergent phenomena

become interconnected, interdependent and creative and the whole world depends upon it (as in ecosystems and living creatures for example). These Complex Multi-modal, Multi-level Influence (CMLI) networks [2] display persistent, emergent patterns which are adaptive over time. It is well understood that these complex behaviours cannot be understood by breaking them down in a reductionist, linear manner - instead approaches such as those from complexity science can be used. It is outside the scope of this paper to provide a tutorial on complexity science as there are many excellent resources available [3, 4, 5, 6].

In complexity science, the persistent entities which come into being as a result of complex interactions have been called Complex Adaptive Systems (CAS). In the social and human context these CAS also have additional elements such as purposeful sensing, learning, problem-solving, prediction and acting. The term Complex Adaptive Reflexive Systems (CARS) [7] has been used for these purposeful entities - examples of which extend from the brain to social groupings.

The Complexity Framework

Human-machine communities, such as those which are the subject of this paper, are also CARS and we assert that the Complexity Framework we describe here, based on the understanding from complexity science, can be used to capture, visualise and understand their behaviours and expose issues and implications.

It is important to note that, if we only focussed on the dynamic activities of these communities, we would miss many critical enabling factors (such as the role of the environment and the properties of the interacting elements). The Complexity Framework is designed to address these potential shortfalls by considering the phenomena from four 'Aspects'. Each of these Aspects is described briefly below, illustrated in Figure 1 and developed further in the rest of the paper:

Aspect 1: Precursors and Enablers (The Givens / Environment).

Emergent phenomena could not arise, nor complex systems exist without the persistent existence of certain things. These 'Givens' include, for example, enduring 'precursors' of emergence (such as geology - which shapes the environment and hence human activity) which are relatively fixed and those things which are 'enablers' of influence and which people can affect (such as social class structures and economic models). In the paper, this Aspect focuses on the environmental invariants (in both the real world and in cyberspace) which enable 'components', such as people and software agents, to come into being.

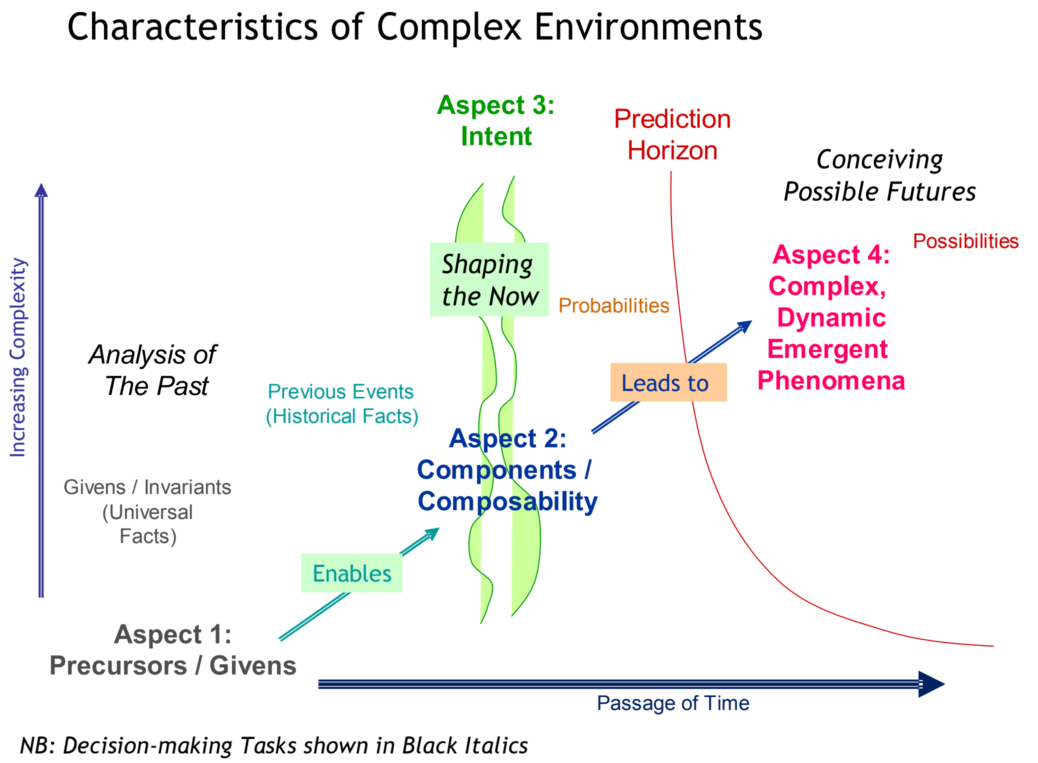

Figure 1 The Complexity Framework in the Decision-Making Context

Aspect 2: Components and Composability ('Assemble-Time').

Complex emergent phenomena arise from the multi-modal / multi-level interactions between these 'components' and their environments. The nature of the interactions which are possible, their interdependencies and influences are largely dictated by the properties of those components - their abilities to sense, perceive, communicate, interact and shape themselves and their environment - these repertoires of behaviour determining the ways in which the entities can group themselves. In the paper, this Aspect is concerned with how changes in the properties of these components (eg, people and agents) and of their 'composability' (eg, sociability) determines the kind of complex behaviours which arise - as understanding this makes dialogue about these matters easier.

Aspect 3: Purpose and Intent (Design-time).

Given the purposeful nature of social activity and decision and policy-making, there has to be some expression of intent, some formulation of a design and / or of goals - however transitory. These expressions shape much of the direction in which future endeavours go. A key aspect of this is the way in which the 'design' is either captured, 'formalised' and disseminated through artefacts or employed through on-the-fly interventions to communicate intent. For example, leaders and policy makers issue formal laws and guidelines. In contrast, people may take to the streets and indicate intent through protest. The paper uses this Aspect to identify the 'design-time' factors that drive the purposeful behaviour of CARS and to consider how things such as intent and 'values' might be shared with, and affect our developing relationships with, machines.

Aspect 4: Dynamic Change (Behaviour and Perception - 'Run-time').

During real life, dynamic adjustment and adaptation are the norm. In the 'process-driven' world of machines this is often not so - yet it is key to the effective use of machines that 'on-the-fly' dynamic changes are possible (to deal with the unexpected etc). This Aspect of the Complexity Framework tells us that the interplay between the following three elements is key to influencing emergence, adaptation and evolution: 1, top-down intervention from the 'leaders'; 2, the self-aware and self-regulatory behaviour of the people, their community and their machines (especially where the perceptions, involvement, reactions and influences of culture and society [8, 9, 10] can have profound effects) and 3, the bottom-up effects which arise from individual personalities, motivations and competencies. In addition, the environment in which the purposeful entities are active will have its own influences - though it should not be forgotten that the environment itself can be purposefully 'interfered with' to change outcomes. Given the importance (in any situation) of the environmental effects of co-evolution, the mappings between cyberspace and 'real space' are profound drivers. The paper uses this Aspect to focus on the kinds of interventions which affect the dynamic aspects of emergence and the patterns of Cause and Influence (C&IN) that propagate through the CMLI networks that we wish to discuss with machines.

Employing the Complexity Framework - Required Mindset.

The Aspects above are not independent - they are interconnected (not in the most obvious ways) and continuously active ('always-on'). It is easy to forget this and to make limiting and over-simplified assumptions about the context under discussion. Hence, it is essential to adopt a suitable multi-aspect mindset when putting complexity to work. This paper will not describe in detail the best ways to employ 'complexity thinking', but two examples will give an idea. On Figure 1, a 'prediction horizon' is shown - this indicates that the limits of prediction relate to the Aspect being considered. One can generally be confident about predicting the future of Aspect 1's Givens (eg, how landscapes evolve), but must be realistic and accept that it is impossible to predict the precise future of Aspect 4 phenomena. Another example would be that people often draw 'arbitrary' boundaries around what they perceive as the limits of the 'system of interest' - but these boundaries are often naive. For example, the notion of a community meeting sounds like something with a clear start, middle and end - but this is not so. The way that a meeting will turn out starts long before the event and continues well after it - transport to the meeting may cause delays or people may win the lottery; relationships or animosities are formed, and consequences follow - all these so-called 'external factors' impacting on the behaviours. In reality, these aspects are not 'external' and may even be purposefully manipulated to either to disrupt or coerce - something of which our machines need to be aware.

3 Dimensions of Teaming and Collaboration

Given the complex challenges described in the Introduction, a relationship between user and machine must be appropriate to the issues being dealt with (Ashby's 'Law of requisite variety'). A key metaphor for complex decision-making is that: "Any machine used for dealing with Aspects 3 and 4 must be a team player", i.e. a 'purposeful entity' capable of contributing effectively to the tasks.

Characterising Human-machine Teams.

For the purposes of this paper we characterise HMT as: "The purposeful involvement of machines in decision-making about complex issues". These human-machines teams usually consist of the following elements:
<u>Communities of Interest (CoIs)</u>: CoIs are sets of teams which can range from being relatively fixed (structured and institutionalised based on known scenarios) to being very ad-hoc (with their structure, relationships and activities being very fluid and event-driven). There are not tight, clearly-defined boundaries between CoIs, there is a great deal of nesting and overlap and many, necessarily different,

ways of working and complex multiple memberships (of humans and machines) such that, by definition, CoIs must extend into cyberspace.

Actors: are taken to cover people carrying out tasks, whether mundane (such as inputting and filtering information) or significant (such as when reasoning from other actors' perspectives or considering competing hypotheses about possible futures) - where they have decided to include a machine 'in the team'.

Machine: is taken to include all synthetic and computational entities that can display collaborative behaviours such that they can contribute effectively to human problem-solving and decision-making.

Artefacts: are taken to be physical or virtual objects manipulated by the teams or other actors or teams as part of their collaboration. These include things such as the information about the teams, their memberships, status and their activities.

Environment: is the world (real and in cyberspace) where activity tasks place.

All these elements could be part of a dialogue with machines about 'complex issues' and hence they will need to be able to 'understand' them (see Section 5).

Characterising Collaboration

Collaborative teaming requires that the machine can engage in dialogue and display ‘helpfulness’ apparent to users as follows:

- Machines acting as team-players, through the discovery of expertise and the declaration of capability. The way that the machines respond (‘adjustable autonomy’) can be affected by dynamically changing authority, obligation and control though policy frameworks (ie, organisational ‘control’);
- Machines reducing human workload - through 'mixed initiative' working. This requires the machines to respond appropriately and dynamically as they collaborate and as the situation changes and / or as the particular user directs (ie, individual ‘control’ - required because of changing contexts, perceptions, perspectives and concerns, and necessarily different human ways of working).
- Providing advice and indications from previous patterns of analysis and identifying trends and sequences of events over time (eg, as warnings). This requires machines to be capable of inference and reasoning - where 'appropriate' behaviour is determined by circumstances, not by following procedures.
- Machines able to recommend (limited) options to the human. This requires machines able to interpret information in the context of the human tasks.

• Machines engaging in subjective, dialectic analysis (e.g. arguing). This requires machines able to reflect different ‘worldviews’ and perceptions, through holding multiple, possibly contradictory interpretations of the same information.
• Machines 'understanding' complex information relationships. This requires machines to be aware of the human context and so capable of tracking multiple, dynamically linked and changing actor-to-information relationships.

4 Dimensions of Human-machine Relationships

When people use machines to support their endeavours a ‘relationship’ is inevitably created which, like all relationships, cannot be sustained without active involvement of the participants. Up to now, human-machine team ‘relationships’ have been relatively simple ones (mostly due to the lack of ‘common ground’ in the dialogue between humans and machines). Yet, for *any* type of collaboration, ‘dialogue’ *is* essential, even though it can cover a vast range of sophistication.
• Consider a situation where the topic, and the resulting dialogue between the user and machine is ‘basic’. Human: “How many people have malaria in village ‘V’?”. Machine: “There are no reports of malaria.”. Human: “Is that because we have not asked, or because we have asked and have not had any reports back?”. Machine: “We asked an hour ago and have had no reports yet”.
• However, if the topic under consideration is more complex then the resulting dialogue will be more sophisticated (current machines would be quickly be 'out of their depth'). For example. Human: “Why has the allegiance of person ‘Y’ changed?”. Machine: “Changed in which way?”. Human: “Such that we can no longer rely on their support”. Machine: “How do you want me to solve that?”.
The question is then: for effective (and appropriate) HMT, what is the range of relationships and levels of dialogue that machines should be able to sustain? In addition we would then need to ask: How are relationships and context sustained over time? What data, information, structures and mechanisms need to be in place 'behind the scenes' to enable machines to do this? (See Section 5 below).

Forms of Human-Machine Relationship

For the purpose of this paper we propose four degrees of dialogue sophistication and therefore types of relationships with machines that these dialogies can sustain (described in the ‘pen pictures’ in Table 1 below):

Sophistication:	Table 1: 'Pen-picture' Description:
Basic **(Type 1 relationship)** **Recent Past**	(E.g. a stand-alone DOS-based computer). This machine has no model of the user or context, and no memory of previous tasks, acting as if it has never met the operator before. It infers nothing from previous interactions, so the human must laboriously re-enter all inputs and preferences, even if they have been unchanged for hundreds of interactions. The interfaces provided are fixed, limited, and chosen for ease of machine implementation, at the cost of human convenience. Limited, fact-based information types (Level 1 - described in Section 5 below) can be handled via a basic local store-and-retrieve paradigm. Richer user-interfaces can improve such systems slightly, but a point of diminishing returns is rapidly reached since there is no underlying flexibility or memory.
Machine-centric **(Type 2 relationship)** **Current Technology**	(E.g. a networked XP-based computer). This machine has basic models of users, and can store some predefined preferences. It makes rudimentary use of context, offering actions to the user that are likely to be relevant, and hiding options which are clearly impossible. It may use a small degree of 'learning' to tailor the interactions with the user. It has a richer interface, and may attempt to guide the user through predetermined tasks. A variety of information types can be handled (up to Level 2), and a variety of networked information access paradigms may be provided, with basic virtualisation capabilities to present heterogeneous sources in a common manner via a browser. The machine has no real autonomy, though it may take the initiative in fixed, pre-scripted situations, though these may often be inappropriate (e.g. the Microsoft Office ‘Paper Clip’ help agent).
Human-Driven **(Type 3 relationship)** **Personal communication devices in 2020**	(Such machines only exist in research laboratories). These machines have rich user models (which draw on users’ expressed or inferred preferences and goals) which can be substantially reconfigured by the user to suit their needs. The use of context and machine learning pervades the relationship, drawing on external information (space, time, location, background knowledge) and internal information (history of interactions, ongoing tasks). The machine maintains, for each user / CoI, overlapping models of their different perceptions, perspectives, context and current concerns. The machine can access a rich set of underlying services; these provide novel possibilities for the user to discover, explore and exploit the many types of information that they provide. The machine has a limited degree of autonomy, governed by learnt preferences or explicit policy, and can carry out a range of background tasks without constant supervision. A flexible set of behaviours enables the machine to interact in both passive and active ways, and to elicit requirements via stereotyped multi-step dialogues.
Social Team-Player **(Type 4 relationship)** **Aspirational**	(Currently science fiction). This self-aware ‘cyberspace’ and its ‘social agents’ understand the preferences, interests, specialities, goals, and ways of working of teams and of individual users. It is aware of its own capabilities and potential contributions, and may proactively initiate tasks or make suggestions to humans, within the bounds of governance policy and team norms. Conversely, it may shy away from tasks that it has learned require human input (due to subtle political sensitivities, for example). Its configuration, behaviours and degree of autonomy are shaped by the wider context and its deep background knowledge, and are responsive to the current task(s) and the workload of the humans. These machines can 'argue' and may take the initiative by asking clarifying questions, reporting the status of its projects when progress is slow, or requesting enhancements to itself (for example, to enable access to authorities / information which appears relevant to the task). Conversely, it may be aware of looming deadlines or critical tasks, and avoid interrupting with minor issues. The machine community is trusted, within limits, by the human team members, who choose to invest time in forming and maintaining relationships, because of the eventual pay-off in terms of improved individual and team capability. Type 4 machines are unlikely to be achieved before 2050.

At the time of writing, technology can already sustain the ‘machine-centric’ relationship up to Type 2, whilst the Type 4 ‘social team player’, as stated, is unfeasible for the foreseeable future. Hence, to support decision-making in complex environments, the issue is to ascertain the level of ‘human-driven’ relationship required in relation to the level of ‘machine-centric’ relationship which will suffice and then identify what underlying capabilities are required to support the necessary behaviours (discussed in Section 6).

5 Characterising Dialogue about Complex Issues

Decision-making in the real world is characterised by the uncertainty and abstract nature of many of the issues. This paper offers (see Figure 2) a Four-level 'Information' Landscape (which has been successfully used in previous work) to characterise the range of issues which HMTs might be called upon to discuss. Current machines are good when dealing with facts (eg, it’s a truck, it’s here), but poor when dealing with so-called 'abstract-information' (eg, it may be a humanitarian convoy, but we are not yet sure). These two situations can be likened to the difference between doing a puzzle and solving a mystery:

• Puzzles: (eg, as in managing logistics) have the following features. They can be a solved by following a procedure or process. The puzzler knows in advance what the puzzle is and so can bound the problem. When something is missing, it is easy to classify the missing item(s), describe it in 'fact-like' terms (a red piece with a face on it) and then search for or collect the missing item(s). Once candidate pieces have been found, it is possible to compare them (because prior models exist) and match a candidate item as being the missing one. It is then possible to fit the new fact into the puzzle (as the puzzle is 'static') and confirm it is the 'right' piece. Levels 1 and 2 in Figure 2 support Puzzling:

- Level 1: Supports simple analysis of familiar environments which can be categorically analysed giving definitive factual results which can easily be stored in a conventional 'flat file database' in a Type 1 (see Table 1) Machine.

- Level 2: Supports purposeful, directed analysis about a situation or about previous events, ie, dealing with probabilistic, fact-like, results which could be stored in a relational database and explored and manipulated using a Type 2 machine.

Mysteries: (eg, People's reactions to so-called 'Swine Flu') cannot be solved by following a 'process' - mysteries involve trial-and-error and risk. The thinker starts

with little or no knowledge of the nature or extent of the problem and starts by building theories / hypotheses or 'fantasies' of what might be going on and projects these 'models' into the future. The competing hypotheses are represented as hyperlinked nets of evidence / patterns of indicators which can be compared to each other. The thinker considers potential indicators / weights of evidence that might exist or be required to support or refute hypotheses and looks to see if they exist in the perceived world. This 'looking' is purposefully directed sensing, not looking and then 'sensemaking' afterwards as is often suggested [11].
The thinker may 'shake the tree' first to 'flush out' items which may prove to be significant indicators. There is no 'final, correct' answer, instead, judgement, assessment, probability and continued purposeful probing are required with an 'open eyes, open mind and open to change' mindset.

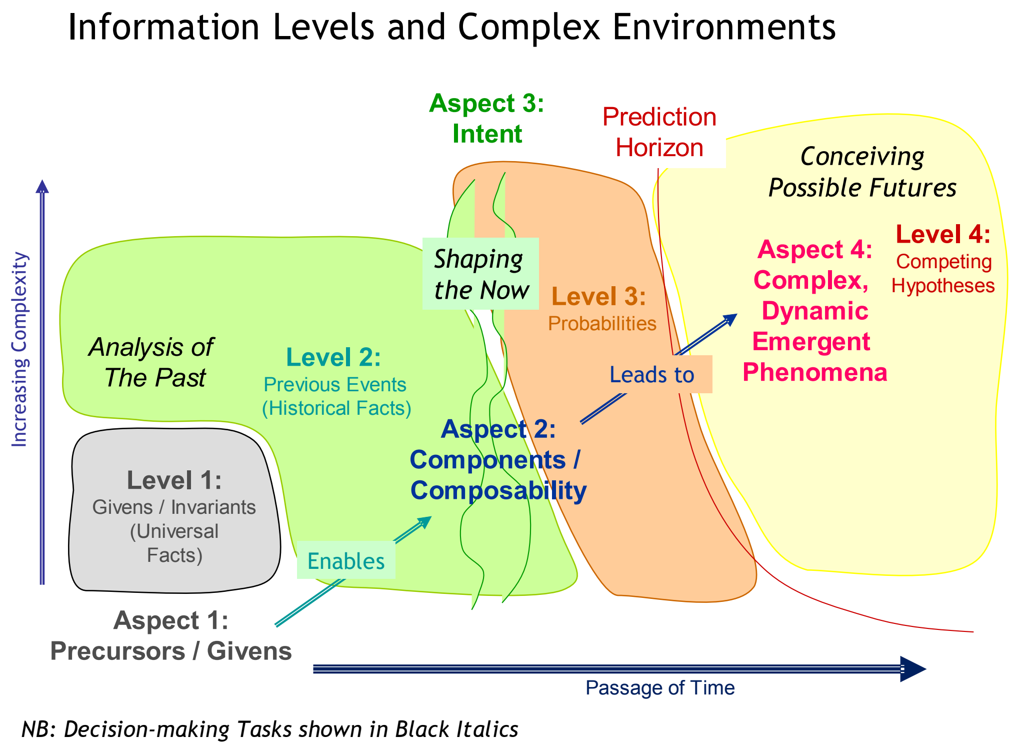

Figure 2 Information Landscape Levels mapped to Aspects of Complexity

Levels 3 and 4 in Figure 2 support the exploration of Mysteries:

- Level 3: Supports user-led exploration, problem-framing and hypothesis testing where the analysis requires creativity and insight, as the levels of significance are low and hard to determine. The Type 3 Machine will have to engage in dialogue about 'abstract information' - which is defined as consisting of hypotheses, their inter-relationships, and related perspectives and context (in the real world and cyberspace). Linkages are essential between the evidence, indicators, assessments and judgements (e.g. about significance, provenance, ‘normality’ and confidence). These are not static facts, but dynamically adapting networks of relationships between items of significance (some of which may be historical, predictive, contradictory, hypothetical or unknown).

- Level 4: Supports user-led intuition and 'brain storming' about possible futures (by 'competing' alternative hypotheses). Here, it may not even be possible to adequately determine what needs to be stored, nor how it can be acquired or visualised. Most of this abstract-information will be in peoples minds and not able to be stored elsewhere (eg, concerning the ‘mind-games’ being played by the other actors). This includes considering the information to support so-called 'red teaming' and to capture other perspectives and perceptions (including information to cope with deception and / or malicious or dysfunctional behaviour). For the foreseeable future, machines will not be able to participate in these dialogues - and the consequences of this must be fully acknowledged.

The implications for HMT are that although information technology is generally well developed to support dialogue about Levels 1 and 2, practitioners are often dealing with issues being played out at Levels 3 and 4 which are poorly supported by current machines, as the next Section will highlight.

6 Human-Machine Teaming - Capabilities Required

At the outset, it has to be said that limitations in human understanding of complexity are an impediment to HMT. If it is not possible to formulate, capture and communicate problems relating to complex environments that it is hard to communicate them - whether to human or machine - let alone solve them. Yet this is a situation which is improving as practitioners extend their trans-disciplinary working to include complexity scientists. The rest of this section will touch on a few of the issues which need to be addressed if effective HMT is to be possible. Concerning the contribution that machines might make, previous work [12] has thoroughly examined issues involved in 'terra-forming cyberspace' and in identifying what needs to be done to make 'collaborating machines' (capable of arguing constructively) acceptable to people [13]. Indeed, there are a number of major projects underway [14], though it has to be said that most of these programmes assume Information Levels 1 and 2 and are not yet capable of dealing with Levels 3 and 4 of human decision-making in complex environments. Nevertheless, they have mapped out many of the capabilities required. These include architectural and engineering principles which indicate that a so-called 'service metaphor' (not to be confused with service-orientated architecture) is an appropriate technical approach to employ for federated HMT. The service metaphor has modular software and hardware services which can be composed in a variety of ways at ‘assemble time’ (deployment) and dynamically ‘threaded’ together at run-time to generate a rich variety of behaviours.

The effective performance of joint human-machine (agent) actions requires dynamic adjustment to styles of interaction and appropriate levels of mutual supervision and control (so-called 'mixed-initiative' team working) including:

- Responsibility for approval and decision making;
- Ability to propose or initiate new actions or policy changes;
- Frequency, modality, and salience of notifications and updates;
- Acknowledgements, clarification, feedback, interruptability.

Machines must store the information necessary to support mixed initiative working with the actors (humans, agents and devices) with which it collaborates. So-called 'adjustable autonomy' is intended to lead to measurably better performance of human-agent teams by enabling the machine adaptation necessary to develop enduring Type 3 relationships. Examples of potential adjustments (made by the humans or by the machine itself) to autonomy include:

- Forbidding machines from attempting the impossible or the unachievable;
- Preventing obligations from exceeding the scope of permissions and possibilities or reducing current obligations in order to take on new tasks;
- Increasing permissions to accommodate new capabilities or changing the decision-making context to increase the range of what is possible.

Issues of trust, provenance and confidence are key and so machines should be able to clarify provenance (which actions have been directly observed and who the ‘seer’ was) especially where items may have been indirectly inferred from other observed actions. Sensitivity, especially in relation to things such as 'freedom of information' issues, is a key challenge for machines where access in any form might, by inference etc, reveal important connections that should remain 'hidden'. Maintaining the current context of a human-machine dialogue is important and this involves linking together related items to form dynamic threads. The evolution of links / development of patterns of links over time can be very significant as it indicates ‘critical pathways’ - vectors, sequences of events leading up to certain actions. Machines must be able to support the historical, current and ‘predictive’ aspects of analysis required for Type 3 dialogue including:

- Mutable and emerging categories (eg of actors, artefacts and environment) and representing irregular behaviours - even without precedence;

- Data structures to represent the variety of sources, subjects and topics of interest along with the relationships between the data items – for example, how the relationships between alternative hypotheses are represented;
- Data structures to represent ‘time’, sequencing and synchronisation;
- The need to represent capture and represent uncertain, equivocal and varying or subjective ‘context’. This requires abstract information structures that enable analysts to handle these uncertainties - even if due to source errors and ambiguities and adversary deceptions and own equivocalities or self-deceptions.

Critically, HMT must support the ability to index, search, annotate, link, share, manipulate and explore information sets and abstract-information across the diverse ways of working of the humans and communities involved.

The development of appropriate visualisation is also an important challenge, as are the means to help users comprehend, explore and manipulate new types of information which may be created as problems are explored.

In addition to the primary behaviours that are apparent to users, machines rely on a range of behind-the-scenes, so-called 'autonomic' behaviours to maintain their own readiness and health. These include: self-(re)configuration / adaptation; self-optimization / tuning; self-healing and self-protection. To provide this autonomic behaviour, the service model must be extended to include proactive ‘agents’ which make autonomous decisions within the constraints of human policy. In terms of risk and maturity, the service metaphor presented in the paper has been widely researched and at least two major series of experiments in the USA [15] have shown its value and utility. For suppliers and implementers the service metaphor provides tools and techniques that can be pragmatically employed.

The basic capabilities of current machines which seriously limits their ability to participate in HMT (or to model complex issues) is being highlighted strenuously both in the world of academic research [16] and in the real world [17]. However, these insights, though authoritative, are not yet being fully acknowledged and their consequences are not being fully debated. It is the intention of this paper to assist in raising these important issues and in stimulating that debate.

7 Implications and Conclusions

Although not considered so far, the role of machines themselves in adding to the complexity of the modern world should not be overlooked. The machine 'value-space' has its own drivers which, innately, are not the same as those in human-space. For agents to reason on our behalf, it is necessary to perform some value mapping. The question is, what are the meaningful equivalents in terms of what can be sensed / perceived / reasoned about / effected etc?

A key insight from experiences in the field is that databases of facts, however ‘complete’, is not enough. Complex decision-making requires the storage and manipulation of a subtle web of inter-linked hypotheses and inferences. Answers (leading to what may be a potentially decisive course of action) could be in the mind of an expert (conceptual domain), locked in a box in a cabinet somewhere (physical domain) or hidden in a number of apparently unrelated electronic reports (‘cyberspace’ domain). The consequence of this is that HMT will only work when these three domains can effectively interact, otherwise we cannot say with confidence that either ‘we don't know’ or ‘the answer is Z’.

For users engaged in HMT, the detailed workings of the machines will largely be hidden. Whatever their task, they will perceive that the machines provide them with awareness of, and access to, diverse types of information and abstractions, and enable them to explore and manipulate this information in collaboration with others. This access will be through logical points of interaction available to them via a variety of devices in a location-independent manner. They will notice that their preferences are remembered and that their concerns, perspectives and ways of working are supported in a flexible manner whatever the circumstances.

A key feature of the HMT will be its dynamic aspects. Users will notice learning, tracking of significant change over time, evolution of machine behaviour, extensibility of ontologies and support for human innovation, insight and counter-intuitive behaviour. In short, the electronic aspects of HMT should augment human analytical and problem-solving capabilities, not inhibit them.

However, the technology to support HMT is not mature. The service metaphor significantly extends Service Oriented Architecture (SOA), However, although SOA concepts are reasonably well developed, experience in implementing them is not. In addition, the use of a service metaphor will require some standardisation and governance to provide common service descriptions across agencies.

Hence, it must be fully acknowledged that, if it is decided to store everything concerning a complex issue on machines then, given that humans can only have a limited ‘dialogue’ with the machines, the human's ability to generate effective decisions will be inhibited. The primacy of human ability must be celebrated.

The author [18] hopes to stimulate debate on this topic with practitioners (about the degree of contribution that it is realistic to expect from machines) and with technologists about how machine behaviour can be further augmented.

Conclusion. This paper asked the question: "What needs to be done to enable machines to make enough sense of complexity so that they could assist people more effectively with this work?". The author hopes that debate has been stimulated on this topic and that practitioners will provide contributions concerning the degree of contribution that it is realistic to expect from machines in the near future and that technologists will consider how machine behaviour can be further augmented to support decision-making in complex situations.

Acknowledgements:

The author acknowledges the contribution of Dr Jeff Bradshaw (of the Florida Institute of Human and Machine Cognition, IHMC) in informing this work.

References

1. Brook, P. Arnold, Jackson and Stevens. *Systems Engineering - Coping with Complexity*. Prentice Hall ISBN 0-13-095085-8.
2. Beautement, P. Broenner C. *Complex Multi-modal Multi-level Influence Networks - Affordable Housing Case Study*. At Complex '09, Shanghai, China. Feb 2009.
3. Morowitz, H. J.: *The Emergence of Everything: How the World became Complex*. Oxford University Press (2002).
4. Santa Fe Institute Library at http://www.santafe.edu/research/topics.php
5. Lewin R.: *Complexity - Life at the Edge of Chaos*. Phoenix (1993).
6. Exystence - European Union Complexity Science Centre: http://once-cs.csregistry.org/tiki-index.php?page=Exystence+focus+documents
7. Mathieson G.: Complexity and Managing to Survive it. DSTL (2005). Instititute for the Study of Coherence and Emergence (I.S.C.E.): http://isce.edu/ISCE_Group_Site/web-content/ISCE_Events/Cork_2005/Papers/Mathieson.pdf
8. Smuts J. C.: Holism and Evolution. pp 17-18. Macmillan. (1926).
9. Maslow A. H. Theory of Human Motivation. Psychological Review 50 pp 370-96 (1943).
10. Ritchey T. Wicked Problems - Structuring Social Messes with Morphological Analysis. Swedish Morphological Society. (2005). http://www.swemorph.com
11. The oft mis-quoted 'OODA Loop' - which starts with 'Observe', then orientates before acting.
12. Bradshaw, J. M., Suri, N., Breedy, M. R., Canas, A., Davis, R., Ford, K. M., Hoffman, R., Jeffers, R., Kulkarni, S., Lott, J., Reichherzer, T., & Uszok, A. (2002). Terraforming cyberspace. IEEE Intelligent Systems, July 2001, pp. 49-56.
13. Bradshaw et al. Making Agents Acceptable to People. In Intelligent Technologies for Information Analysis. N. Zhong and J. Liu (Eds.). Springer 2004.
14. DARPA's Cognitive Assistant Learns Organizes (CALO).: http://caloproject.sri.com/darpa/
15. Coalition Agents Experiment Web site: http://www.aiai.ed.ac.uk/project/coax.
16. Beautement, P. Allsopp, D. Greaves, M. Goldsmith, S. Spires, S. Thompson, S. Janicke, H. 'Autonomous Agents and Multi-agent Systems (AAMAS) - Issues and Challenges' published in “Lecture Notes in Computer Science”. From: http://dx.doi.org/10.1007/11683704_1. Applications of Multi-Agent Systems: International Workshop, DAMAS 2005, Utrecht.
17. http://www.futurefastforward.com/component/content/article/902-military--intelligence/608-usjfcom-commanders-guidance-for-effects-based-operations-by-james-n-mattis-latest-update-251108?tmpl=component&print=1&page=
18. Supporting material available from: http://www.abaci.net/html/our_publications.html