Dimensions of Teleautonomy in Mobile Agents

John Anderson[1] **and Alfred Wurr**
Autonomous Agents Laboratory
Department of Computer Science
University of Manitoba
Winnipeg, Manitoba, Canada R3T2N2

andersj@cs.umanitoba.ca *awurr@cs.umanitoba.ca*
http://www.cs.umanitoba.ca/~aalab

Abstract

Teleautonomy, the remote issuing of instructions for control of an agent that is intended to function semi-autonomously while employing those instructions, is an important issue in intelligent agency. In most cases where intelligent agents are employed, we neither wish to directly control every detail, nor trust them to deal completely autonomously with every important task. Similarly, in multi-agent settings, agents must similarly instruct and be instructed by other agents. This paper breaks the concept of Teleautonomy into a spectrum along several dimensions for research purposes, emphasizing the demands each dimension places upon an intelligent agent. We also describe two current approaches to teleautonomy in mobile agents that together cover a breadth of this range.

1. Teleautonomy and Intelligent Agents

Creating autonomous agents, operating intelligences that can accomplish goals in complex environments without human intervention, has always been one of the major goals of artificial intelligence. Indeed, the dependence on human intervention and human encoding of domain knowledge (among other elements) has been a major criticism of AI in general in the past [Dreyfus].

In the past ten years, the view of autonomy as a flexible concept has taken root (e.g. [Evans et al., 1992; Anderson and Evans, 1993]), wherein an agent may indeed be intended to function autonomously, but allow the sacrifice of some of that autonomy in return for some greater overall good. These developments are in part a recognition of the nature of human autonomy: all of us are biologically able to function as lone entities, but enter into social contracts causing us to take on responsibilities to others in order to lead a more enriched existence than we might otherwise. Similarly, there has been increased focus in the AI community on multi-agent settings, where the dynamics of such settings make similar social contracts advantageous.

Beyond the sacrifice of autonomy to function well as a member of a larger population, there has also been the realization that in many cases we do not want systems to function completely autonomously. Rather, we want them to function within bounds (sometimes flexible, sometimes very rigid) that we set upon tasks and goals. Even in a relatively simple setting such as mail filtering, for example, there are situations where one would be uncomfortable having items deleted without a direct interaction. We deal similarly with other humans providing contracted services for us: we provide general instructions and expect these instructions to be carried out within written, verbal, or understood (by shared context and culture) bounds. We neither expect to be consulted on every detail nor have exceptional elements dealt with without a consultation.

Dealing with this breadth of autonomy is one of the most significant challenges to the acceptance of intelligent agents today. Kitano [1], for example, proposes a layered view of autonomy to deal with such situations, and argues that viewing autonomy in this manner is a necessary step in the evolution of acceptance of agent-based systems (issues of context-dependent autonomy and agent acceptance are also noted by many others (e.g. [2,3]).

While many agents can be thought of as operating in a desktop environment, the ubiquity of the internet means that these will often not be operating in a strictly local sense. Many other agent types will be operating in remote geographic locations because their nature suits them for such purposes (e.g. small robotic units for exploration in disaster settings [Kitano]) or because their task environments make their work dangerous to human agents (e.g. space exploration). In such settings we have the same need for a degree of direct agent control, and also the same need for agents to deal with one another in a similar fashion. An understanding of these needs has led to the current focus on teleautonomy as a major issue in robotic control. Teleautonononomic control involves the issuing of instructions, recommendations, or direct internal control remotely, by humans or other agents. Like autonomy in

[1] This work is supported by an operating grant from the National Science and Engineering Research Council (NSERC)

general, there is a wide range of teleautonomic situations, and dealing with this breadth is part and parcel of dealing with a wide range of autonomy.

While direct agent control is the most obvious application of teleautonomy, there are exciting overlaps to intelligent user interfaces as well: telling a robotic agent to perform some general task and leaving the details to the agent itself is strongly analogous to instructing a piece of software via it's user interface to perform a task for us and leaving the details open for the interface's interpretation. Thus the results of current work in teleautonomy should have some applicability to this area as well.

The remainder of this paper introduces a categorization of teleautonomy, and then describes two approaches that deal with a significant component of this breadth.

2. Dimensions of Teleautonomy

Given the range of situations in which an agent can accept input from another and function to some degree within that input, or indeed even selectively ignore that input in appropriate situations, the breadth of the range of potential teleautonomic situations is clearly significant. For the purposes of actually designing agents that can be controlled in a teleautonomous manner, however, a more precise breakdown is more useful.

For the purposes our research in teleautonomy in intelligent agents, we split the range of teleautonomous situations across three dimensions as shown in Figure 1. In each dimension we move from that which places the least demands on the agent itself in terms of additional facilities to support teleautonomy, (the origin in the Figure) to that which places the most.

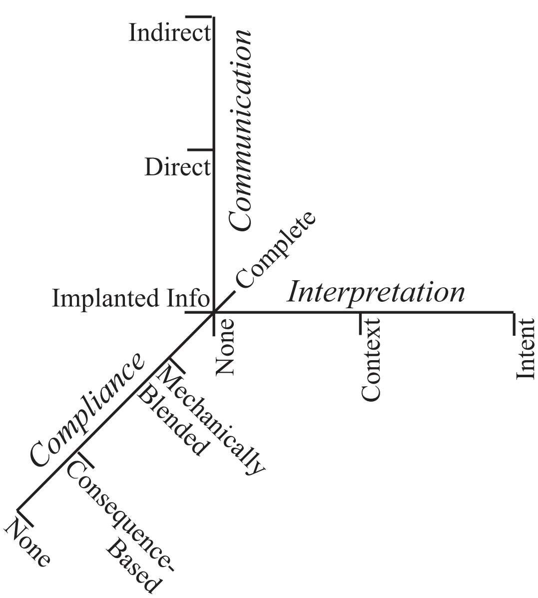

Figure 1. Dimensions of Teleautonomy

In a teleautonomous agent, we care about the extent to which an agent can put its own *interpretation* on input given by another agent teleautonomously. At its simplest, this disallows any freedom of interpretation whatsoever: a teleautonomous controller provides specific instructions that are meant to be taken literally. In a robotic control situation, stating to move left, or to bring some precise object to a precise location, for example. While this may suffice for some situations, many will require at least some form of contextual interpretation. The controller may specify an instruction in a non-specific context ("get me one like this"), for example. Like the other dimensions, moving from the origin on an axis requires additional components in an agent (here, contextual interpretation abilities to varying degrees), and moving from the origin can also be viewed as adding both freedom and responsibility to an agent. In this case, an agent is responsible for performing some interpretation, but looking at the same ability from the agent's perspective, it is *able* to take an interpretation of the controller's wishes if it feels the need to as well. At a more extreme level, the controller can communicate abstract intent to the agent rather than a specific detailed command. The latter is taking context abstraction to an extreme, and is really the ideal in terms of replacing a traditional software user interface. Instead of specific commands, if the controller's intent can be communicated my abstract intent (e.g. "I want these in order"), an agent can maximally work within its abilities to achieve this without bothering me for details. Communicating intent may even be non-verbal.

Analogous to the interpretability of commands, we also have the degree to which the agent is forced to comply directly with a teleautonomous controller's instructions. Where the autonomy in teleautonomy is lacking, an agent may have no choice but to follow instructions to the letter (indeed, the extreme would be the complete lack of autonomy in a remote-controlled car). Adding complexity, an agent can blend the influence of an instruction (i.e. the direct response that instruction should provoke) along with others it may consider to produce an overall alteration. This would allow the agent to blend an instruction such as "move left" along with the direction it is currently desiring, and average the two for example, allowing compliance with both. Beyond this, an agent may be given the ability to reason about the consequences of compliance and non-compliance and choose to comply or not depending on higher-level goals. For example, a Mars rover told to move left may know there is a steep drop off the teleoperator cannot see, and refuse to comply. An agent may similarly have an existing goals that are threatened by the teleoperator's instructions, and reason about the importance and consequences of each. At the far extreme an agent is completely autonomous and has no requirement for compliance of any sort. While this may seem counterproductive in a human teleoperational setting, it useful with both human teleoperators and other agents to receive advice that can be followed or ignored as appropriate in addition to explicit instructions.

The third dimension in a teleautonomous agent involves the nature of communication in that agent. At its simplest, the input of the teleoperator can be placed in an internal format directly into the agent's computational components, much as a joystick signal interfaced directly with the motors of a remote controlled car. This is communication in that information is obtained externally, but it's implementation would be the human equivalent of a direct neural signal, bypassing auditory or visual communication abilities and integrating directly into the decision-making components themselves. It certainly bears no resemblance to anything we would traditionally call communication. Moving to greater agent complexity, we have traditional direct communication, where an agent receives textual information and parses it, or obtains an auditory or analogous signal and performs some understanding of that signal to constitute meaning. This would follow with the integration and processing of that instruction according to the limits or choices made on the other dimensions. Moving beyond this we move to the realm of indirect communication through the environment, or *stigmergy*. Stigmergic communication involves modification of the environment to convey information indirectly to others that are receptive to those modifications, and is most commonly associated with insects. Ants, for example, have evolved to secrete pheromones when moving food from a location where it is found, and have similarly evolved an attraction to these same pheromones. As more ants detect pheromone trails and follow them, they lay down additional pheromones, making the trail more attractive to others [Holldobler and Wilson, 1990]. As food runs out, more ants wander away from the food source and fewer ants on the trail lead to less pheromone being deposited. This similarly lessens the attraction of the trail. Coordination of a primitive sort is thus achieved not by directly communicating ant to ant, but by modifying the environment in such a way that others can use this information [Resnick, 1997]. This manner of communication is common in complex human tasks as well: on a highway for example, specific symbols are used to remove guesswork and information processing load (e.g. marking curves for night driving). While the latter is certainly more complicated than ant stigmergy, from the point of view of teleautonomy the same principles apply.

It should be noted that indirect communication has been placed on the communication scale as requiring additional supports beyond direct communication. This may at first seem contradictory given the known complexity of interpreting communication. This scale has been organized in terms of how obvious the communication is and how much the agent must do to respond appropriately to the communication. In directly implanted information, the agent has no processing requirements in interpreting the communication whatsoever – the agent does not need to recognize that communication has occurred, and the information is present and already integrated in an internal form. In direct communication, the agent must extract information and integrate it internally, but the complexity of this has to do with the complexity of the communication itself – the agent generally knows that communication has occurred and this itself is a piece of useful knowledge. In indirect communication, the agent still has the requirements of parsing communication (while stigmergic communication is usually simple in nature, there is nothing in stigmergy that precludes complexity; it is a method rather than a restriction on content). However in stigmergy, the agent must also keep watch on the environment to see that communication has in fact occurred at all (i.e. whether any alterations are in fact something it needs to know about, or naturally occurring phenomena). In adversarial situations, where alterations in the environment may also be interpreted by an an opponent, making these subtle is important [Wurr, 2002]. From the point of view of teleautonomy and its effect on agent design, indirect communication can be seen as adding complexity to an agent.

3. Supporting Teleautonomy

As can be seen from the previous Section, the further we move from the origin of Figure 1, the more facilities we require. Teleautonomous situations close to the origin have been shown to require only a minimal addition to agent components. The most notable of these efforts is that of Ali and Arkin [7], who illustrate a method of including user-input into directly into a behaviour-based autonomous robot. This input forms one of a number of behaviours that are active at any one point in time. The agents response to its environment is computed through a mathematical combination of behaviour activation levels, and thus the user's input is considered in exactly the same fashion as the agent's own motivations. From the agent's internal standpoint, a mysterious urge to move in the direction the teleoperator specifies is experienced. Ali and Arkin show that this can be done with no significant changes to an overall behaviour based architecture. In figure 1, this approach is at the origin (level-0) on communication and interpretation (information is received directly into the agent's internals and forms one of its own motivations – the agent does not have the ability to consider any aspect of it). On the third dimension, compliance, this approach moves to level-1: remote instructions are mechanically mixed with internal motivations to produce a single output. Chapman[5] provides a similar approach, where a symbolic instruction such as "get it" is hard-coded as a connection to a network for decision-making – the information is once again received directly as an internal connection, and while it may seem there is context here, the context is pre-determined by the connection location within this network (that is, the instruction is directly received as referring to a specific object in the agent's knowledge).

We are interested in developing agent designs that cover a broader range of these dimensions and allow for more significant elements of teleautonomy as depicted in Figure 1. We are currently working with two approaches, stigmergic teleautonomy and intention-based teleautonomy. The remainder of this paper examines these approaches in light of the dimensions of teleautonomy we have already presented. Space allows only a brief examination, and references are made to longer works on these.

3.1. Stigmergic Teleautonomy

As stated in Section 2, Stigmergic teleautonomy involves communication based on modifying the environment as opposed to implanted or direct information communication. We are working with an approach to stigmergic teleautonomy in robotic agents that while keeping the architecture of the agent as simple as possible, supports instruction and information receipt via stigmergic methods. This approach is designed to perform in complex environments with timely response requirements. The intent of this work is to illustrate the power of stigmergy in agents that are still relatively simple (behaviour- or schema-based agents incorporating state, as opposed to agents that do considerable planning and world modeling).

Our stigmergic agents are motor schema-based (a variant of behaviour-based architectures where each agent calculates a vector based on fields of attraction and repulsion for each active scheme, and normalize a result to determine a final action vector [Pirjanian, 1999]). Perceptual schemas observe the environment and determine conditions of interest (including examining objects in the environment that may or may not be stigmergic indicators left by others). Motor schemas encapsulate the actions necessary to achieve a particular behavior (e.g. *follow-wall, avoid-obstacle,…*) and each is stimulated by output of perceptual schemas and provide as output suggested alterations of agent effectors. These schemas can also be connected and grouped into implementations of higher-level behaviours.

Like any other communication, stigmergy operates by having agreement on meaning between sender and receiver. Stigmergic communication is performed in this implementation by allowing behaviours to to leave environmental markers under specific conditions, and by allowing similar agents to recognize those same environmental markers as signals for information. We may agree on certain features that will distinguish an stigmergic marker from naturally occurring phenomena, and perceptual schemas must be in place implementing these distinctions, as well as motor schemas that cause the appropriate distinctive markers to be left when communication is desired. For example, in group navigation a marker may be left when an agent perceives a particular danger (a dead-end corridor for example, or a likely antagonist's hiding place) or useful landmark. Agents perceiving these can be influenced by these instructions to perform particular activities (in navigation, for example, by perceiving a landmark and indicating it is ahead, it can an agent give advance knowledge and allow it escape local minima situations, such as moving toward a less desireable landmark when a better one is just out of perceptual range [Wurr, 2002]). This implementation performs navigation in Half-Life[1] environment. This environment is a rich (in terms of sophistication of agent movement and object interaction, as well as in terms of interacting with other agents) environment real-time software simulation. The original intent of this software is for individual gaming against computer-generated opponents, but like other such environments [Laird and van Lent, 2001], it is increasingly being adopted by AI as a rich, controllable experimental setting for intelligent agent research. The results of this work are both applicable to improving teamwork of human opposition in such gaming environments and demonstrations of the principles of stigmergic teleautonomy as well.

Stigmergic communication is accomplished physically in this domain by dropping objects that are similar to physical game objects (weapons, ammunition, etc.). In a physical robotic domain, this could be accomplished through leaving visible cues (e.g. paper cutouts) or with greater sophistication, incorporating recent work on scent as an attractor in robotics (e.g. [Hayes et al., 2002]). Like the natural fading of scent, a software environment can simulate the fading of stigmergic cues by making them less visible as time passes.

Just as other members of a team can perform stigmergic communication in this fashion, a human can enter the environment and perform similarly, allowing this approach to support both multi-agent indirect communication and human teleautonomy. We are currently experimenting with direct comparisons of the efficiency of stigmergic vs. direct communication in performing group navigation of a domain in both the presence and absence of antagonistic agents.

From the standpoint of supporting teleautonomy, this approach not only moves upward (significantly in terms of the communication dimension), but also supports a breadth of teleautonomy in comparison to that of Ali and Arkin [7] or Chapman [5]. We have also used this schema-based approach without the stigmergic schemas [Wurr, 1999] to implement teleautonomy analogous to use of this approach in more direct stigmergy analogous to Ali and Arkin [2000]'s. Here, schemas such as follow-leader could be used to allow a human controlling a half-life agent to

[1] Half-Life is a trademark of Sierra Entertainment, Inc.

directly guide artificial agents (that is, act directly on their internal mechanisms as opposed to via communication). Beyond earlier approaches, this can also be used to have agents react in different manners to different stigmergic signals, and take the importance of these signals in context. It thus covers both levels 0 and 1 on the interpretation scale, and provides greater agent flexibility than previous teleautonomous approaches. The significant aspect of this is that it does so while still employing the same, simple behaviour-based strategies of earlier work.

3.1. Intention-Based Teleautonomy

4. Discussion

In this paper we have presented a novel breakdown of the dimensions of teleautonomy for research purposes and a description of our efforts in covering a breadth of these dimensions.

Also interested in stigmergy and examining situations where implicit communication can be incidental to task performance rather than a deliberate and separate action.

Bibliography

[1] Kitano, Hiroaki, "Nausicaä and the Sirens: A Tale of Two Intelligent Autonomous Agents", *IEEE Expert*, December, 1996, pp. 60-61.

[2] Negroponte, Nicolas, "From Direct Manipulation to Delegation", in Bradshaw, J. (Ed.), *Software Agents*, Cambridge: MIT Press, pp. 57-66.

[3] Norman, Donald, "How People Might Interact with Agents", in Bradshaw, J. (Ed.), *Software Agents*, Cambridge: MIT Press, pp. 49-55.

[4] Barker, Ken, Mark Evans, and John Anderson, "Quantification of Autonomy in Multi-Agent Systems", *Proceedings of the AAAI Workshop on Cooperation Among Heterogeneous Agents,* San Jose, California, July, 1992, 6 pp.

[5] Chapman, David, *Vision, Instruction, and Action* (Cambridge, MA: MIT Press), 1991, 311 pp.

[6] Agre, Philip E., *The Dynamic Structure of Everyday Life*, Ph.D. Dissertation, Department of Computer Science and Electrical Engineering, MIT, 1988. 282 pp.

[7] Ali, Khaled, and Ronald Arkin, "Multiagent Teleautonomous Behavioral Control'", *Machine Intelligence and Robotic Control*, Vol. 1, No. 2, 2000, pp. 3-10.

[8] Anderson, John, and Mark Evans, "Constraint-Directed Improvisation", *Proc. of the 11th Biennial Canadian Society for the Computational Studies of Intelligence Conference (AI-96)*, Toronto, May, 1996, pp. 1-13 (*Lecture Notes in AI* 1081)

[9] Anderson, John, "Providing a Broad Spectrum of Agents in Spatially-Explicit Simulation Models", In Gimblett, H. R (Ed.), *Integrating GIS and Agent Based Modeling Techniques for Understanding Social and Ecological Processes.* Oxford: Oxford University Press. 2001.

[10] Goldberg, Ken (Ed.), *The Robot in the Garden: Telerobotics and Telepistemology in the Age of the Internet* (Cambridge, MA:MIT Press), 2000. 366 pp.

[Evans et al., 1992] "Achieving Flexible Autonomy in Multi-Agent Systems using Constraints" M. Evans, J. Anderson, and G. Crysdale, *Applied Artificial Intelligence*, Volume 6, Number 1, 1992, pp. 103-126.

"Supporting Flexible Autonomy in a Simulation Environment for Intelligent Agent Designs", J. Anderson and M. Evans, *Proceedings of the Fourth Annual Conference on AI, Simulation, and Planning in High Autonomy Systems*, Tucson, AZ, September, 1993, pp. 60-66.

[Wurr, 1999] Wurr, Alfred, "Situated Multi-Agent Teamwork in Dynamic Environments", Technical report, Department of Computer Science, University of Manitoba, May, 1999. 37 pp.

[Wurr, 2002] Wurr, Alfred, " Robotic Team Navigation in Complex Environments Using Stigmergic Clues", forthcoming M.Sc. Dissertation, Department of Computer Science, University of Manitoba, 2002.

T. Hayes, A. Martinoli, and R. M. Goodman. Distributed Odor Source Localization. IEEE Sensors. To appear. © IEEE Press.

[Laird and van Lent, 2001] Laird, J. and Van Lent, M. "Human-Level AI's Killer Application: Interactive Computer Games", *AI Magazine*, pp. 15-25, Summer 2001.