

# Narrative Intelligence

Phoebe Sengers  
*Center for Art and Media Technology (ZKM)*  
*Institute for Visual Media*

## 1.0 Introduction

[C]ertainly it is the case that all biological systems.... [b]ehave in a way which just simply seems *life-like* in a way that our robots never do.... Perhaps we have all missed some organizing principle of biological systems, or some general truth about them. Perhaps there is a way of looking at biological systems which will illuminate an inherent necessity in some aspect of the interactions of their parts that is completely missing from our artificial systems.... [P]erhaps we are currently missing the *juice* of life. (Brooks 1997)

AI techniques have come a long way. We can now build agents that can do a lot for us: they search for information on the Web (Shakes *et. al.* 1997), trade stocks (Analytix 1996), play grandmaster-level chess (Hsu *et. al.* 1990), patrol nuclear reactors (Baker and Matlack 1998), remove asbestos (Schempf 1995), and so on. We have learned to use agents as powerful tools.

But one of the oldest dreams of AI is the ‘robot friend’ (Bledsoe 1986), an artificial being that is not just a tool but has its own life. Such a creature we want to talk to, not just to find out the latest stock quotes or the answer to our database queries, but because we are interested in its hopes and feelings. Yes, we can build smart, competent, useful creatures, but we have not built very many that seem complex, robust, and alive in the way that biological creatures do. Who wants to be buddies with a spreadsheet program, no matter how anthropomorphized? No matter how smart artificial creatures become, AI will not have completely fulfilled the dreams we have had for it until agents are not just intelligent but also *intentional* — living creatures with their own desires, feelings, and perspectives on the world.

How can we build creatures that are not just smart but visibly alive in this way? In this paper I will try to provide some answers by turning the question

inside out. In order to build creatures that seem intentionally alive, we can try to understand how *human beings* identify and interpret intentional behavior. If we understand what properties enable people to understand behavior as intentional, we may be able to generate a similar style of behavior in artificial agents. Such agents, being engineered to be intentionally understandable, are more likely to really seem alive than an agent optimized for formal intelligence or for use as a goal-oriented tool.<sup>1</sup>

Narrative psychology suggests that people understand the behavior of living agents by structuring visible activity into narrative (Bruner 1990). That is, people understand and interpret intentional behavior by organizing it into a kind of story. If this is the case, then our agents may appear more intentional if we build them so that their behavior provides the cues to be understandable as narrative. Here, I will describe the fundamental principles of narrative in order to explain (1) how current agent construction techniques actually undermine the appearance of intentionality and (2) how narrative principles can be applied in agent design to support the user in understanding the agent as an intentional being.

## **2.0 Principles of Narrative Psychology (or How We (Sometimes) Make Sense of Creatures)**

Artificial Intelligence attempts to generate intentional creatures setting up a correspondence between biological, living beings and automatic processes of the kind that can run on computers. That is, AI agents should ideally be understandable both as well-specified physical objects and as sentient creatures. But it turns out that there is a deep tension between these two views on agents. This is because human understanding of the behavior of humans and other conscious beings differs in important ways from the way we understand the behavior of such physical objects as toasters. Identifying the distinction between these two styles of comprehension is essential for discovering how to build creatures that are understandable not just as helpful tools but as living beings.

The way people understand meaningful human activity is the subject of narrative psychology, an area of study developed by Jerome Bruner (1986, 1990). Narrative psychology shows that, whereas people tend to understand inanimate objects in terms of cause-effect rules and by using logical reasoning, intentional behavior is made comprehensible, not by figuring out its physical laws, but by structuring it into narrative or ‘stories.’ This structure is not simply observed in the person’s activity; we generate it through a sophisticated process of interpretation. This interpretation involves such aspects as finding relations between what the person does from moment to moment, speculating

about how the person thinks and feels about his or her activity, and understanding how the person's behavior relates to his or her physical, social, and behavioral context.

Even non-experts can effortlessly create sophisticated interpretations of minimal behavioral and verbal cues. In fact, such interpretation is so natural to us that when the cues to create narrative are missing, people spend substantial time, effort, and creativity trying to come up with possible explanations. This process can be seen in action when users try to understand our currently relatively incomprehensible agents!

This sometimes breathtaking ability — and compulsion — of the user to understand behavior by constructing narrative may provide the key to building agents that truly appear alive. If humans understand intentional behavior by organizing it into narrative, then our agents will be more *intentionally comprehensible* if they provide narrative cues. That is, rather than simply presenting intelligent actions, agents should give visible cues that support users in their ongoing mission to generate narrative explanation of an agent's activity. We can do this by organizing our agents so that their behavior provides the visible markers of narrative. The remainder of this paper presents the properties of narrative and explains how they apply to agent construction.

### 3.0 Prolegomena to a Future Narrative Intelligence

There has recently been a groundswell of interest in narrative in AI and human-computer interaction (HCI). Narrative techniques have been used for applications from automatic camera control for interactive fiction (Galyean 1995) to story generation (Elliott *et. al.* 1998). Abbe Don (1990) and Brenda Laurel (1986, 1991) argue that, since humans understand their experiences in terms of narrative, computer interfaces will be more understandable if they are organized as narrative. Similarly, Kerstin Dautenhahn and Chrystopher Nehaniv (1998) argue that robots may be able to use narrative in the form of autobiography to understand both themselves and each other.

Michael Travers and Marc Davis developed the term Narrative Intelligence in the context of an informal working group at the MIT Media Lab to describe this conjunction of narrative and Artificial Intelligence. David Blair and Tom Meyer (1997) use the same term to refer to the *human* ability to organize information into narrative. Here, I want to suggest that Narrative Intelligence can be understood as the confluence of these two uses: *that artificial agents can be designed to produce narratively comprehensible behavior by structuring their visible activity in ways that make it easy for humans to create narrative explanations of them.*

In order to do this, we need to have a clear understanding of how narrative

works. Fortunately, the properties of narrative have been extensively studied by humanists. Bruner (1991) nonexhaustively lists the following properties:

- *Narrative Diachronicity*: Narratives do not focus on events on a moment-by-moment basis, but on how they relate over time.
- *Particularity*: Narratives are about particular individuals and particular events.
- *Intentional State Entailment*: When people are acting in a narrative, the important part is not what the people do, but how they think and feel about what they do.
- *Hermeneutic Composability*: Just as a narrative comes to life from the actions of which it is composed, those actions are understood with respect to how they fit into the narrative as a whole. Neither can be understood completely without the other. Hence, understanding narrative requires interpretation in a gradual and dialectical process of understanding.
- *Canonicity and Breach*: Narrative makes its point when expectations are breached. There is a tension in narrative between what we expect to happen and what actually happens.
- *Genericness*: Narratives are understood with respect to genre expectations, which we pick up from our culture.
- *Referentiality*: Narratives are not about finding the absolute truth of a situation; they are about putting events into an order that feels right.
- *Normativeness*: Narratives depend strongly on the audience's conventional expectations about plot and behavior.
- *Context Sensitivity and Negotiability*: Narrative is not 'in' the thing being understood; it is generated through a complex negotiation between reader and text.
- *Narrative Accrual*: Multiple narratives combine to form, not one coherent story, but a tradition or culture.

While these properties are not meant to be the final story on narrative, they stake out the narrative landscape. Taking narrative agents seriously means understanding how these properties can influence agent design. It will turn out that current AI techniques, which largely inherit their methodology from the sciences and engineering, often undermine or contradict the more humanist properties of narrative. Here, I will explain problems with current agent-building techniques, techniques already in use that are more amenable to narrative, and potential practices that could be more friendly to the goal of meaningful Narrative Intelligence.

One note of caution: the goal here is to interpret the properties of narrative with respect to agent-building. This interpretation is itself narrative. Since, as we will see below, the nature of narrative truth is different from that of scientific factuality, this essay should not be read in the typically scientific sense of stating the absolute truth about how narrative informs AI. Rather, I

will look at the properties of narrative in the context of current AI research, looking for insights that might help us to understand what we are doing better and suggest (rather than insist on) new directions. My conclusions are based on my particular human perspective, as a builder of believable agents in Joseph Bates' Oz Project with a strong interest and training in the cultural aspects of Artificial Intelligence.

### *3.1 Narrative Diachronicity*

The most basic property of narrative is its diachronicity: a narrative relates events *over time*. Events are not understood in terms of their moment-by-moment significance, but in terms of how they relate to one another as events unfold. For example, if Fred has an argument and then kicks the cat, we tend to infer that the cat-kicking is not a random event, but a result of his frustration at the argument. When people observe agents, they do not just care about what the agent is doing; they want to understand the relations between the agent's actions at various points in time. These perceived relations play an important role in how an agent's subsequent actions are understood. This means that, to be properly understood, it is important for agents to express their actions so that their intended relationships are clear.

However, it is currently fashionable to design behavior-based autonomous agents using action-selection, an agent-building technique that ignores the diachronic structure of behavior. Action-selection algorithms work by continuously redeciding the best action the agent can take in order to fulfill its goals (Maes 1989a). Because action-selection involves constantly redeciding the agent's actions based on what is currently optimal, behavior-based agents often display a kind of "schizophrenia" (Sengers 1996b). By schizophrenia I mean that they jump from behavior to behavior, without any kind of common thread that structures these behaviors into understandable sequences. Schizophrenic agents undermine the appearance of intentionality because agent action seems to be organized arbitrarily over time, or, at maximum, in terms of automatic stimulus-response.<sup>2</sup>

More generally, expressing the relationships between behaviors is not well supported in most behavior-based systems (a complaint also raised in (Neal Reilly 1996)). While these architectures do provide support for clear, expressive *individual* behaviors, they have problems when it comes to expressing relations *between* behaviors. This is because a typical behavior-based system (e.g. Blumberg 1994, Brooks 1986a, Maes 1991) treats each behavior separately; behaviors should refer as little as possible to other behaviors. Because of this design choice, a behavior, when turned on, does not know why it is turned on, who was turned on before it, or even who else is on

at the same time. It knows only that its preconditions must have been met, but it does not know what other behaviors are possible and why it was chosen instead of them. In most behavior-based architectures, behaviors simply do not know enough about other behaviors to be able to express their interrelationships to the user.

In this light, classical AI would seem to have an advantage over alternative AI, since it is explicitly interested in generating structured behavior through such mechanisms as scripts and hierarchical plans. However, classical AI runs into similar trouble with *its* modular boundaries, which occur not between behaviors but between the agent's functionalities. For example, the agent may say words it cannot understand, or clearly perceive things that then have no influence on what the agent decides to do.

Fundamentally, agent-building techniques from Marvin Minsky's *Society of Mind* (1988) to standard behavior-based agent-building (Maes 1991) to the decomposition of classical agents into, for example, a planner, a natural language system, and perception (Vere and Bickmore 1990) are all based on divide-and-conquer approaches to agenthood. Being good computer scientists, one of the goals of AI researchers is to come up with modular solutions that are easy to engineer. While some amount of atomization is necessary to build an engineered system, narrative intentionality is undermined when the parts of the agent are designed so separately that they are visibly disjoint in the behavior of the agent. Schizophrenia is an example of this problem, since when behaviors are designed separately the agent's overall activity reduces to a seemingly pointless jumping around between behaviors. Bryan Loyall similarly points out that visible module boundaries destroy the appearance of aliveness in believable agents (Loyall 1997).

The end result is that the seductive goal of the plug-n-play agent — built from the simple composition of arbitrary parts — may be deeply incompatible with intentionality. Architectures like that of Steels (1994) which design behaviors in a deeply intertwined way, make the agent design process more difficult, but may have a better shot at generating the complexity and nonmodularity of organic behavior. Less drastic solutions may involve the use of transition sequences to relate and smooth over the breaks between separately designed behaviors (Stone 1996) (Sengers 1998b). I use this strategy elsewhere as a cornerstone for the Expressivator, an architecture for Narrative Intelligence (1998a).

### 3.2 *Particularity*

Narratives are not simply abstract events; they are always particular. "Boy-meets-girl, boy-loses-girl" is not a narrative; it is the structure for a narrative, which must always involve a particular boy, a particular girl, a particular way

of meeting, a particular way of losing. These details bring the story to life. However, details do not by themselves make a narrative either; the abstract structure into which the details can be ordered brings meaning to the details themselves. A narrative must be understood in terms of tension between the particular details and the abstract categories they refer to; without either of these, it is meaningless.

This same tension between the abstract and the particular can be found in agent architectures. Agent designers tend to think about what the agent is doing in terms of abstract categories: the agent is eating, hunting, sleeping, etc. However, users who are interacting with the agent do not see the abstract categories; they only see the physical movements in which the agent engages. The challenge for the designer is to make the agent so that the user can (1) recognize the particular details of the agent's actions and (2) generalize to the abstract categories of behavior, goal, or emotion that motivated those details. Only with a full understanding at both the particular and the abstract levels will the user be likely to see the creature as the living being the designer is trying to create.

But AI researchers are hampered in this full elucidation of the dialectical relationship between the particular and the abstract by the valorization of the abstract in computer science. In AI we tend to think of the agent's behaviors or plans as what the agent is 'really' doing, with the particular details of movement being a pesky detail to be worked out later. In fact, most designers of agents do not concern themselves with the actual working out of the details of movement or action at all. Instead, they stop at the abstract level of behavior selection, reducing the full complexity of physical behavior to an enumeration of behavior names. Maes (1989b), for example, uses abstract atomic actions such as "pick-up-sander."

Similarly, the Oz Project's first major virtual creature, Lyotard, was a text-based virtual cat (Bates *et. al.* 1992). Because Lyotard lived in a text environment, his behaviors were also text and therefore high level: "Lyotard jumps in your lap," "Lyotard eats a sardine," "Lyotard bites you." Because we were using text, we did not need to specify action at a more detailed level. We did not have to specify, for example, how Lyotard moved his legs in order to jump in your lap.

Lyotard's successors, the Woggles (Loyall and Bates 1993), on the other hand, were graphically represented. As a consequence, we were forced to specifically define every low-level action an agent took as part of a behavior. The effort that specification took meant that we spent less time on the Woggles' brains, and as a consequence the Woggles are not as smart as Lyotard. But — surprisingly to us — the Woggles also have much greater affective power than Lyotard. People find the Woggles simply more convincingly alive than the text cat, despite the fact that Lyotard is superior

from an AI point of view. This is probably in part because we were forced to define a particular body, particular movements, and all those pesky particularities we AI researchers would rather avoid.<sup>3</sup>

If we look at animation (e.g. Thomas and Johnston 1981), the valorization tends to run to the other extreme: the particular is seen as the most essential. Animators tend to think mostly at the level of surface movement; this movement may be interpretable as a behavior, as evidence of the character's emotions, as revealing the character's motivations, or as any of a host of things or nothing at all. Animators make the point that any character is of necessity deeply particular, including all the details of movement, the structure of the body, and quirks of behavior. The abstract comes as an afterthought. Certainly, animators make use of a background idea of plot, emotion, and abstract ideas of what the character is doing, but this is not the level at which most of animators' thinking takes place.

Loyall (1997) points out that this focus on the particular is also essential to the creation of effective believable agents. A focus on particularity by itself, though, is not adequate for creating artificial agents. Agents are expected to interact autonomously with the user over time. In order to build such autonomous systems, we need to have some idea of how to structure the agent so that it can recognize situations and react appropriately. Because we do not know every detail of what will happen to the agent, this structure necessarily involves abstract concepts in such aspects as the modules of the agent, the classification of situations according to appropriate responses, abstract behaviors, emotions, goals, and so on.<sup>4</sup> We must design agents, at least partially, at an abstract level.

In order to build agents that effectively communicate through narrative, AI researchers will need to balance their ability to think at the abstract level with a new-found interest in the particular details their system produces, an approach that seems to be gaining in popularity (Frank *et. al.* 1997). Narrative Intelligence is only possible with a deep-felt respect for the complex relationship between the abstract categories that structure an agent and the physical details that allow those categories to be embodied, to be read, and to become meaningful to the user.

### *3.3 Intentional State Entailment*

Suppose you hear the following:

A man sees the light is out. He kills himself.

Is this a story? Not yet. You don't understand it.

After endless questions, you find out that the man was responsible for a light house. During the night, a ship ran aground off shore. When the man

sees that the light house light is out, he realizes that he is responsible for the shipwreck. Feeling horribly guilty, he sees no choice but to kill himself. Now that we know what the man was thinking, we have a story.

In a narrative, what actually happens matters less than what the actors feel or think about what has happened. Fundamentally, people want to know not just what happened but *why* it happened. This does not mean the causes of an event in terms of physical laws or stimulus-response reactions, but the reasons an actor freely chose to do what s/he did. The narrative is made sense of with respect to the thoughts and feelings of the people involved in its events.

This means that when people watch autonomous agents, they are not just interested in what the agent does. They want to know how the agent thinks and feels about the world around it. Instead of knowing only what the agent has chosen to do, they want to know *why* the agent has chosen to do it.

But in many autonomous agent architectures, the reasons for the decisions the agent makes are part of the implicit architecture of the agent and therefore not directly expressible to the user. Bruce Blumberg's Hamsterdam architecture, for example, represents the appropriateness of each currently possible behavior as a number; at every time step the behavior with the highest number is chosen (Blumberg 1996). With this system, the reasons for behavioral choice are reduced to selecting the highest number; the actual reason that behavior is the best is implicit in the set of equations used to calculate the number. The agent simply does not have access to the information necessary to express why it is doing what it does.

Instead of this emphasis on *selecting* the right action, Tom Porter (1997) suggests the strategy of *expressing* the reasons an agent does an action and the emotions and thoughts that underly its activity. This means organizing the agent architecture so that reasons for behavioral change are explicit and continuously expressed. By showing not only what the agent does, but why the agent does it, people may have an easier time understanding what the agent is thinking and doing in general.

A deeper problem with current architectures is that ethologically-based models such as (Blumberg 1996) presuppose that most of what an agent does is basically stimulus-response. But when we build agents that embody these theories, they often work through stimulus-response or straightforward cause-effect. This automaticity then carries forward into the quality of our agent's behavior> the agent seems nothing more than it is, an unthinking automaton.

More generally, as scientists, we are not interested in the vagaries of free will; we want to develop clearly-specified rules to explain why animals do what they do when they do it. In particular, in order to embody our ideas of agenthood in an automatically-running computational architecture, we must intentionally adopt what Daniel Dennett (1987) might call a 'non-intentional stance' – there is no algorithm for "and then the agent should do whatever it

feels like doing.” We therefore tend to develop theories of behavior that are fundamentally mechanistic.

But these mechanistic theories of agency often lead to mechanistic qualities of behavior in our generated agents. As a consequence, agents are not only non-intentional for us; they are also often reduced to physical objects in the eyes of the user. Narrative Intelligence requires agents that at least appear to be thinking about what they are doing and then making deliberate decisions according to their own feelings and thoughts, rather than simply reacting mindlessly to what goes on around them. We may be automatic; but we should not appear so.

### *3.4 Hermeneutic Composability*

Narrative is understood as a type of communication between an author and an audience. In order to understand this communication, the audience needs to go through a process of interpretation. At the most basic level, the audience needs to be able to identify the atomic components or events of the narrative. But this is just the beginning; the audience then interprets the events not in and of themselves but with respect to their overall context in the story. Once the story is understood, the events are re-identified and re-understood in terms of how they make sense in the story as a whole. In essence, this is a complex and circular process: the story only comes into being because of the events that happen, but the events are always related back to the story as a whole.

This property of narrative is another nail in the coffin of the dream of plug-n-play agents. If users continuously re-interpret the actions of the agent according to their understanding of everything the agent has done so far, then agent-builders who design the parts of their agents completely separately are going to end up misleading the user, who is trying to understand them dialectically.

More fundamentally, the deep and complex interrelationships between the things creatures do over time is part of what makes them come alive, so much so that when there are deep splits between the parts of a person — for example, they act very happy when they talk about very sad things — we consider them mentally ill. This kind of deep consistency across parts is very difficult to engineer in artificial systems, since we do not have methodologies for engineering holistically. In alternative AI, it is currently fashionable to believe that these deep interrelationships may come about *emergently* from separately designed pieces; whether this is wishful thinking or the foundation for a novel form of wholistic design is not yet clear. It may be that the best we can do is the surface impression of wholism; whether that will be enough remains to be seen.

### 3.5 *Canonicity and Breach*

A story only has a point when things do not go the way they should. "I went to the grocery store today" is not a story; but it is the beginning of a story when I go on to say "and you'll never believe who I ran into there." There is no point to telling a story where everything goes as expected; there should be some problem to be resolved, some unusual situation, some difficulty, someone behaving unexpectedly.... Of course, these deviations from the norm may themselves be highly scripted ("boy-meets-girl, boy-loses-girl, boy-wins-girl-back" being a canonical example).

It may be, then, that the impression of intentionality can be enhanced by making the agent do something unexpected. Terrel Miedaner's short story "The Soul of the Mark III Beast" (1981) revolves around just such an incident. In this story, a researcher has built an artificially intelligent robot, but one of his friends refuses to believe that a robot could be sentient. This continues until he hands her a hammer and tells her to destroy the robot. Instead of simply breaking down — the friend's canonical expectation — the robot makes sounds and movements that appear to show pain and fear of death. This shakes the friend so much that she starts to wonder if the robot is alive, after all. Watching the robot visibly grapple with its end, the friend is led to sympathy, which in turn leads her to see the robot as sentient.

More generally, people come to agents with certain expectations, expectations that are again modified by what they see the agent do. The appearance of intentionality is greatly enhanced when those expectations are not enough to explain what the agent is doing. That is, the agent should not be entirely predictable, either at the level of its physical actions or at the level of its overall behavioral decisions. Characters in a Harlequin romance — who inevitably fall in love with the man they hate the most (James 1998) — have nowhere near the level of 3-dimensionality of the complex and quirky characters of a Solzhenitsyn novel. Similarly, agents who always do the same thing in the same situation, whose actions and responses can be clearly mapped out ahead of time, will seem like the automatons they are, not like fascinating living creatures.

Making the creature do unexpected things may seem like a contradiction to one of the basic goals of Narrative Intelligence: making agent behavior more understandable. Stereotypicity may seem like a helpful step towards making agent behavior comprehensible. After all, if the agent always does the same thing for the same reasons in the same ways, the user will always know exactly what the agent is doing. But since users are very good at creating narrative, stereotyped actions bore the audience. In order to create compelling narrative, there needs to be some work for the reader to do as well. The agent designer needs to walk the line between providing enough cues to users that they can

create a narrative, and making the narrative so easy to create that users are not even interested.

### *3.6 Referentiality*

The ‘truth’ in stories bears little resemblance to scientific truth. The point of stories is not whether or not their facts correspond to reality, but whether or not the implicit reasoning and emotions of the characters feels right. A plausible narrative does not essentially refer to actual facts in the real world, but creates its own kind of narrative world that must stand up to its own, subjective tests of realism.

Similarly, extensive critiques have been made in AI about the problem of trying to create and maintain an objective world model (Agre 1997). Having the agent keep track of the absolute identity and state of objects in the external world is not only difficult, it is actually unhelpful. This is because in many situations the absolute identity of an object does not matter; all that matters is how the agent wants to or could use the object. As a substitute, Philip Agre has introduced the notion of deictic representation, where agents keep track of what is going on, not in any kind of absolute sense, but purely with respect to the agent’s current viewpoint and goals (Agre 1988).

While understanding the power of subjectivity for agents, AI in general has been more reluctant to do away with the goal of objectivity for agent researchers. AI generally sees itself for better or for worse as a science, and therefore valorizes reproducibility, testability, and objective measures of success. For many, intelligence is or should be a natural phenomenon, independent of the observer, and reproducible in an objective manner. Intelligence is not about appearance, but about what the agent ‘actually’ does. This reveals itself in the oft-repeated insistence that agents should not just appear but *be* ‘really’ alive or ‘really’ intelligent — anything else is considered illusionary, nothing more than a bag of tricks.

This ‘real’ essence of the agent is usually identified with its internal code — which is also, conveniently enough, the AI researcher’s view of the agent. As a consequence, the impression the agent makes on the user is often considered less real, and by extension, less important. This identification of the internal code of the agent as what the agent really *is* — with the impression on the user a pale reflection of this actual essence — has an unexpected consequence: it means that the subjective interpretation of the audience is devalued and ignored. The result is agents that are unengaging, incoherent, or simply incomprehensible.

This does not mean the AI community is idiotic. Most AI researchers simply have a scientific background, which means they do not have training in subjective research. But the accent on AI as a science, with the goals and

standards of the natural sciences, may lose for us some of what makes narrative powerful. I do not believe that 'life' in the sense of intentionality will be something that can be rigorously, empirically tested in any but the most superficial sense. Rather, generating creatures that are truly alive will probably mean tapping into the arts, humanities, and theology, which have spent centuries understanding what it means to be alive in a meaningful way. While intelligent tools may be built in a rigorous manner, insisting on this rigor when building our 'robot friends' may be shooting ourselves in the foot.

### 3.7 Genericness

Culturally supplied genres provide the context within which audiences can interpret stories. Knowing that a story is intended to be a romance, a mystery, or a thriller gives the reader a set of expectations that strongly constrain the way in which the story will be understood. These genre expectations apply just as well to our interpretations of everyday experience. The Gulf War, for example, can be understood as a heroic and largely victimless crusade to restore Kuwait to its rightful government or as a pointless and bloody war undertaken to support American financial interests, depending on the typical genre leanings of one's political philosophy.<sup>5</sup>

These genres within which we make sense of the world around us are something we largely inherit from the culture or society we inhabit. This means at its most basic that different kinds of agent behavior make sense in different cultures. For example, I once saw a Fujitsu demo of 'mushroom people' who would, among other things, dance in time to the user's baton. In this demo, the user went on swinging the baton for hours, making the mushroom people angrier and angrier. Finally, it was the middle of the night, and the mushroom people were exhausted, obviously livid — and still dancing. I thought this behavior was completely implausible. "Why on earth are they still dancing? They should just leave!" I was told, "But in Japan, that would be rude!" My American behavioral genre expectations told me that this behavior was unnatural and wrong — but in Japan the same behavior is correct.

Since cultural expectations form the background within which agent behavior is understood, the design of intentionally comprehensible agents needs to take these cultural expectations into account. In contrast, the current practice of building agents tends not to consider the specific context in which the agent will be used. Patricia O'Neill Brown (1997) points out that this is likely to lead to agents that are misleading or even useless. This means an understanding of the sociocultural environment in which an agent will be inserted is one important part of the agent design process. In fact, O'Neill Brown goes one step further: not only does cultural baggage affect the way agents *should* be designed, it *already* affects the way agents are designed. That is, the way designers think of agents has a strong influence on the way we build them to start out with.

This tension can even be seen within American culture (Sengers 1994). In particular, the American tradition of AI has included two competing visions of what it means to be an agent. Classical AI on the one hand tends to favor representational, deliberative, rational, cognitive agents. In contrast, alternative AI tends to argue for nonrepresentational, reactive, situated, and embodied agents.

From within AI, these two conceptions of agents can seem to stem purely from technical imperatives. With a broader perspective, they can be traced back to the culture in which AI is situated, which has a number of different traditions of conceptualizing what it means to be human. In the West, human beings have traditionally been thought of through what cultural theorists call the Enlightenment model of consciousness: the mind is separated from the body, it is or should be fundamentally rational, and cognition divorced from emotion is the important part of experience. This form of agency is in many ways fundamentally equivalent to the notion of agent proposed by classical AI. At the same time, in the last 30 or 40 years this view of humanity has been challenged by the ‘schizophrenic’ model of consciousness (see e.g. Massumi 1992).<sup>6</sup> This model considers people to be immersed in and to some extent defined by their situation, the mind and the body to be inescapably interlinked, and the experience of being a person to consist of a number of conflicting drives that work with and against each other to generate behavior. Alternative AI is clearly inspired by this notion of being human.

The conclusion from this and similar analyses of the metaphors behind AI technology (e.g. (Wise 1998)) is that *AI research itself is based on ideas of agenthood we knowingly or unknowingly import from our culture*. Given that this is the case, our best bet for harnessing the power of culture so it works for AI instead of against it is what Agre calls a critical technical practice: the development of a level of self-reflective understanding by AI researchers of the relationship between the research they do and culture and society as a whole (Agre 1997).

### 3.8 Normativeness

Previously, we saw that a story only has a point when things do not go as expected, and that agents should similarly be designed so that their actions are not completely predictable. But there is a flip side to this insight: since the point of a story is based on a breach of conventional expectations, narratives are strongly based on the conventions that the audience brings to the story. That is, while breaking conventions, they still depend on those same conventions to be understood and valued by the audience.

Intentional agents, then, cannot be *entirely* unpredictable. They play on a tension between what we expect and what we do not. There needs to be enough familiar structure to the agent that we see it as someone like us; it is only against this background of fulfilled expectations that breached expectation comes to make sense.

### 3.9 Context Sensitivity and Negotiability

Rather than being presented to the reader as a *fait accompli*, narrative is constructed in a complex interchange between the reader and the text. Narrative is assimilated by the reader based on that person's experiences, cultural background, genre expectations, assumptions about the author's intentions, and so on. The same events may be interpreted quite differently by different people, or by the same person in different situations.

In building narrative agents, on the other hand, the most straightforward strategy is context-free: (1) choose the default narrative you want to get across; (2) do your best to make sure the audience has understood exactly what you wanted to say. The flaw in this strategy is that narrative is not one size fits all. It is not simply presented and then absorbed; rather, it is constructed by the user. In assimilating narrative, users relate the narrative to their own lived experience, organizing and understanding it with respect to things that have happened to them, their generic and conventional expectations, and their patterns of being. Narrative is the interface between communication and life; through narrative a story becomes a part of someone's existence.

This means the 'preformed narrative' that comes in a box regardless of the audience's interests or wishes is throwing away one of the greatest strengths of narrative: the ability to make a set of facts or events come to life in a meaningful way *for the user* — in a way that may be totally different from what someone else would see. Rather than providing narrative in a prepackaged way, it may be more advantageous to provide the *cues* for narrative, the building blocks out of which each user can build his or her unique understanding.

And if narrative is not the same for everyone, then narrative agents should not be, either. If narrative is fundamentally user-dependent, then inducing narrative effectively means having some ideas about the expected audience's store of experience and typical ways of understanding. Just as the author of a novel may have a typical reader in mind, the designer of an agent needs to remember and write for the people who will use that agent, relating the agent's projected experiences to the lived experience of the desired audience.

And just as the author of a novel does not expect every possible reader to understand its point, the author of an agent does not necessarily need to be disappointed if only some people understand what the agent is about. The statistical testing of an agent's adequacy over user population may miss the point as much as using bestseller lists to determine the quality of novels. It may be that making the point well with a few users is better, from the point of view of the designer, than making the point adequately with many users.

### *3.10 Narrative Accrual*

Generally speaking, narratives do not exist as point events. Rather, a set of narratives are linked over time, forming a culture or tradition. Legal cases accumulate, becoming the precedents that underly future rulings. Stories we tell about ourselves are linked together in a more-or-less coherent autobiography.

The mechanism by which narratives accrue is different from that of scientific fact. We do not find principles to derive the stories, or search for empirical facts in the stories to accept or reject according to a larger paradigm. Stories that contradict one another can coexist. The Bible, for example, first cheerfully recounts that, on the 7th day, God made man and woman at the same time; a little later, God makes man out of mud, and only makes woman after man is lonely (Various 1985). Similarly, we do not necessarily have a problem reconciling two stories, in one of which Fred is mean, and in the other he is nice. The process of reconciliation, by which narratives are joined to create something of larger meaning, is complex and subtle.

The ways in which stories are combined — forming, if not a larger story, at least a joint tradition — are not currently well understood. Once we have a better understanding of how this works, we could use these mechanisms in order to modulate the effects of our narrative agents as they move from episode to episode with the user. As Dautenhahn (1997) has suggested, agents are understood by constructing ‘biographies’ over the course of prolonged interaction. By investigating the mechanisms whereby the user constructs these biographies from the mini-narratives of each encounter, we stand a better chance of building our agent so that over time it makes an appropriate impression on the user.

## **4.0 Narrative Intelligence as a Research Tradition**

Now that we have learned about the properties of narrative and what they might have to do with agents, we can return with a broader perspective to the questions raised in the introduction: why, as Rodney Brooks wonders, our agents simply do not seem alive in the way that natural creatures do. Having looked at the nature of narrative in detail, I believe that at least part of the problem in AI research is its inheritance from scientific research traditions, through which — mostly unconsciously — the life in our creatures can be drained away. In this final section, I will first explain why an overreliance on science can undermine narrative, and then look at the implications of this problem for Narrative Intelligence as a field.

#### 4.1 Narrative and Atomization

In general, narrative involves understanding the wholistic relationships between things: the relationship between the different events in the story, the relationship between the events and how the actors feel about the events, the relationship between what the author tries to tell and the way in which the audience constructs what it hears, the relationship between the audience member and his or her cultural background, and so on. With layer upon layer of interdependency, this narrative view of the world can become extremely complex.

In contrast, the scientific worldview tends to value simplicity through black-boxing, a process I have elsewhere (1998a) termed *atomization*. Atomization is the process of splitting something that is continuous and not strictly definable into reasonably well defined, somewhat independent parts. In classical AI, for example, we atomize agents by dividing them into relatively independent functionalities, such as a planner, an executor, and a perception system. In alternative AI, we atomize agents by dividing them into relatively independent behaviors.

We do this for a good reason: atomization is a way of getting a handle on a complex phenomenon, a way of taking something incoherent, undefined, and messy and getting some kind of fix on it. It is only through atomization that we can understand something clearly enough to be able to engineer a working system of any complexity. In this sense, atomization is essential to AI.

But atomization as used in science is not a transparent methodology. In many ways, its properties are the exact opposite of those of narrative. This can be seen more concretely by inverting each of the properties of narrative:

1. *Structure over time*: Narrative structure is diachronic; it is about how events relate to one another. Atomistic structure is statistical. Patterns of events over time are simply correlated with one another.
2. *Essence*: Narrative is interested in particular events; it matters which person a story is about. Atomization is interested in finding salient properties so that events can be generalized as parts of a system; individual water molecules, for example, are not differentiated. Narrative sees events as essentially particular; atomization, as essentially abstract, with specific features seen as noise.
3. *Components*: Narrative is interested in events mainly in terms of how the actors involve understand and interpret them. Scientific atomization is interested in the facts that can be established independently of any one person's experience.
4. *Combination*: Narrative is wholistic; the act of bringing its components together changes the components themselves. In atomization, the

combination of events is seen as the sum of the parts.

5. *Relation to expectation*: Narrative must contain elements that are unexpected; things cannot go as planned. In contrast, the goal of scientific atomization is to be able to predict and control with reasonable certainty the outcome of events.
6. *Referentiality*: Narrative is fundamentally subjective; it is about how different people come to interpret it in different situations. Scientific atomization is meant to be objective. Its laws hold in every situation, independent of context and interpretation.
7. *Dependence on culture*: Similarly, while narrative is largely dependent on culturally bound norms and expectations, scientific atomization is thought of as culturally universal, true for everyone.
8. *Audience judgment*: The audience must use its judgment for narrative to be realized; but audience judgment is considered to be unimportant for determining the truth of scientific atoms.
9. *Application*: The way in which narrative is used depends on context; atomic facts are meant to be absolute.
10. *Accrual*: Narratives combine to form a not necessarily coherent tradition. Atomic facts are combined by comparing them and finding a logical structure that subsumes them. Facts that are inconsistent are thrown away. These aspects are summarized in Table 1.

Aspect	Scientific Worldview	Humanist Worldview
structure over time	statistical	diachronic
essence	abstract	particular
components	factual	experiential
combination	additive	wholistic
relation to expectation	predictable	creative
referentiality	objective	subjective
dependence on culture	culturally universal	culturally variable
audience judgment	unimportant	essential
application	absolute	context-sensitive
accrual	logical coherence	tradition

*Table 1. Relations between scientific (atomistic) and humanist (narrative) worldviews.*

Clearly, these statements are too absolute. Not all scientific work is, for example, interested purely in statistical properties of events. Many forms of science have shaded over to the narrative end of the spectrum (e.g. Goldstein 1995). Psychiatry and neurology, for example, often depend heavily on case studies, which chronicle the particular life history of individual patients. While

science, being a heterogeneous set of practices, cannot be absolutely identified with the purely atomistic end of the spectrum, scientific values and practices do cluster towards atomization. Similarly, the humanities are not unanimous in being placed at the purely narrative end, but humanistic projects do tend to have more of the narrative attributes. This means the division of atomization from narrative is meaningful, at least heuristically.

#### *4.2 Atomization, Narrative, and AI*

Atomization is an essential tool for AI. As Brooks (1986b) states,

There are many possible approaches to building an autonomous intelligent system. As with most engineering problems they all start by decomposing the problem into pieces, solving the subproblems for each piece, and then composing the solutions.

But because atomization is closely linked with mechanicity, its value must be called into question when the goal is building truly intentional beings. As narrative psychology has demonstrated, when humans try to make intentional behavior meaningful, they use a fundamentally different procedure from that of atomization and the scientific method. Rather, humans create meaning by structuring their experience according to narrative, in the tradition of the humanities. This difference between the atomistic standpoint of the agent designer and the narrative viewpoint of the eventual agent audience can undermine the designer's ability to construct intentionally understandable agents.

To understand how this works, consider the situation where a designer is building an autonomous agent on scientific principles. S/he starts with a living agent — or idea of an agent — that the designer wants to copy. The designer tries to understand the dynamics of this agent's behavior by finding out its atomic attributes. For example, the designer may try to find out the typical activities in which the agent engages, the conditions under which each activity is likely to occur, and the length of time the agent tends to spend on various activities. Each of these attributes reflects one measurable, objectively verifiable property of the dynamic which the agent designer would like to reproduce.

The designer may objectively and accurately measure many different attributes of the agent. Using these facts, the designer can then construct a system that has the same attributes; for example, the new agent will have the same typical activities, the same probability of being in any particular behavior, and spend the same amount of time on each activity. Once the system can generate behavior that closely approximates the finite list of atomic attributes with which the designer has measured the agent, the designer is satisfied that the agent is a reasonable facsimile of the living agent.

Scientifically speaking, the designer is correct.

But now consider the user's point of view. Rather than being interested in the empirically determinable individual attributes of the creature, the user focuses on how the creature's activities seem to meld together into a whole. Rather than caring about whether the creature engages in a particular set of typical activities, for instance, the user may want to know how the activities relate to each other. Rather than caring about the percentage of time spent in any particular activity, the user may wonder why the agent seems to prefer some activities to others at different moments. Rather than caring about the objectively measurable attributes of the original agent, the user will probably care about whether the subjective part of the original agent — its soul or personality — seems to come across.

These *narrative* attributes of the agent's activities — the extent to which the agent's behavior is not simply the sum of predictable parts — are precisely what the scientific copy of the creature has left out. This means that *even if the designer succeeds in making an accurate copy according to scientifically measurable properties, from the point of view of the user the living creature is fundamentally different from the constructed agent.*

If we are to build agents that truly appear intentional, then, we need to include narrative properties in our design of artificial creatures. Currently, many (though by no means all) AI techniques fall on the scientific end of the spectrum in Table 1. This atomistic worldview reflects itself not only in the internal code of the agents, but also in the quality of the externally observable behavior that forms the basis by which audiences try to understand the agent. The challenge for an AI that wants to build, not just intelligent tools, but intentional agents, is to find ways of moving AI methodology towards the values embodied in narrative. The point is not that narrative is good and science as embodied in current AI is bad, but that we need *both* narrative and AI techniques set in relationship to one another. With the development of a dialectical relationship between narrative and AI, we may finally be able to reach our goal: the development of artificial beings we can truly call intentional.

## Notes

1. If the goal of behavior that is not just "life-like" but alive seems trivial or useless to you, you may take heart in the following side-effect of this goal: if we know how humans try to understand intentional behavior, we may be able to build creatures that are more understandable.
2. This is unfortunate, since the original idea of constantly redeciding behavior came in work explicitly interested in diachronic structure. Philip Agre and David Chapman (1987) focus,

not on the design of the agent per se, but on the ongoing dynamics of the agent and the environment. The goal is to construct action-selection so that, when put in a particular environment, agents *will* tend to have particular diachronic structure in their behavior. Continuous redecision is part of this work because it keeps the agent's actions closely tied to the agent's context, a property that is also important for narrative, as we will see below. However, the concept of the action-selection algorithm itself tends to undermine diachronic structure, especially when it is used to simply design the agent's behavior alone, rather than for designing a complex interaction between an agent and a dynamically changing environment.

3. Similar arguments may hold for robots. The Sony robotic dogs at Agents '97 were a compelling demonstration that robots may have much greater affective power than even graphically represented agents (Fujita and Kageyama 1997).
4. It may be that one day we can use machine learning to develop this structure instead; whether this learned agent must also be structured abstractly remains to be seen.
5. A similar perspective is used to automatically generate ideologically-based understanding of news stories in (Carbonell 1979). For a humanist example of the effect of generic ways of thinking on the actions we take in our everyday lives, see (Sontag 1979).
6. This is inspired by, but not the same as, the notion of schizophrenia used in psychiatry.

## References

- Agre, Philip E. 1988. *The Dynamic Structure of Everyday Life*. PhD thesis, Massachusetts Institute of Technology Artificial Intelligence Laboratory, Cambridge, MA.
- Agre, Philip E. 1997. *Computation and Human Experience*. Cambridge, UK: Cambridge University Press.
- Agre, Philip E. and David Chapman. 1987. Pengi: An implementation of a theory of activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, 268-272.
- Analytix, Inc. 1996. <http://www.analytix.com/>.
- Baker, Stephen and Carol Matlack. 1998. Chernobyl: If you can make it here... *Business Week*, March 30, 1998, 168-170.
- Bates, Joseph, A. Bryan Loyall, and W. Scott Reilly. 1992. Integrating reactivity, goals, and emotion in a broad agent. Technical Report CMU-CS-92-142, Carnegie Mellon University. Also appeared in the *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, Indiana.
- Blair, David and Tom Meyer. 1997. Tools for an interactive virtual cinema. In Robert Trappl and Paolo Petta (eds), *Creating Personalities for Synthetic Actors: Towards autonomous personality agents*, number 1195 in Lecture Notes in Artificial Intelligence. Berlin: Springer Verlag, 83-91.

- Bledsoe, Woody. 1986. I had a dream: AAAI presidential address. *AI Magazine*, 7(1):57-61.
- Blumberg, Bruce. 1994. Action-selection in Hamsterdam: Lessons from ethology. In *Proceedings of the 3rd International Conference on the Simulation of Adaptive Behavior*, Brighton.
- Blumberg, Bruce. 1996. *Old Tricks, New Dogs: Ethology and interactive creatures*. PhD thesis, MIT Media Lab, Cambridge, MA.
- Brooks, Rodney. 1986a. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2:14-23.
- Brooks, Rodney. 1986b. Achieving Artificial Intelligence through building robots. Memo 899, MIT AI Lab.
- Brooks, Rodney. 1997. From earwigs to humans. *Robotics and Autonomous Systems*, 20(2-4):291-304.
- Bruner, Jerome. 1986. *Actual Minds, Possible Worlds*. Cambridge, MA: Harvard University Press.
- Bruner, Jerome. 1990. *Acts of Meaning*. Cambridge, MA: Harvard University Press.
- Bruner, Jerome. 1991. *The narrative construction of reality*. *Critical Inquiry*, 18(1):1-21.
- Carbonell, Jaime. 1979. *Subjective Understanding: Computer models of belief systems*. PhD thesis, Yale University Department of Computer Science, New Haven, CT.
- Dautenhahn, Kerstin. 1997. Could be you – the phenomenological dimension of social understanding. *Cybernetics and Systems Journal*, 28(5): 417-453.
- Dautenhahn, Kerstin and Chrystopher Nehaniv. 1998. Artificial life and natural stories. In *International Symposium on Artificial Life and Robotics (AROB III)*, volume 2, Beppu, Oita, Japan, 435-439.
- Dennett, Daniel. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.
- Don, Abbe. 1990. Narrative and the interface. In Brenda Laurel (ed.), *The Art of Human-Computer Interface Design*, Reading, MA: Addison-Wesley, 383-391.
- Elliott, Clark, Jacek Brzezinski, Sanjay Sheth, and Robert Salvatoriello. 1998. Story-morphing in the affective reasoning paradigm: Generating stories semi-automatically for use with ‘emotionally intelligent’ multimedia agents. In Katia P. Sycara and Michael Wooldridge (eds), *Proceedings of the Second International Conference on Autonomous Agents*, New York: ACM Press.
- Frank, Adam, Andrew Stern, and Ben Resner. 1997. Socially intelligent virtual petz. In Kerstin Dautenhahn (ed.), *Proceedings of the 1997 AAAI Fall Symposium on Socially Intelligent Agents*, Menlo Park, CA: AAAI Press. AAAI Technical Report FS-97-02.
- Fujita, Masahiro and Koji Kageyama. 1997. An open architecture for robot entertainment. In W. Lewis Johnson (ed.), *Proceedings of the First International Conference on Autonomous Agents*, New York: ACM Press, 435-442.

- Galyean, Tinsley. 1995. *Narrative Guidance of Interactivity*. PhD thesis, MIT Media Lab.
- Goldstein, Kurt. 1995. *The Organism: A holistic approach to biology derived from pathological data in man*. Zone Books, NY.
- Hsu, Feng-hsiung, Thomas S. Anantharaman, Murray S. Campbell, and Andreas Nowatzyk. 1990. Deep thought. In T. Anthony Marsland and Jonathan Schaeffer, (eds), *Computers, Chess, and Cognition*. New York: Springer Verlag, 55-78.
- James, Ellen. 1998. *Her Protector*. Harlequin Superromance. Toronto: Harlequin Books.
- Laurel, Brenda. 1986. Interface as mimesis. In *User-Centered System Design*, Hillsdale, NJ: Lawrence Erlbaum Associates, 67-85.
- Laurel, Brenda. 1991. *Computers as Theatre*. Reading, MA: Addison-Wesley.
- Loyall, A. Bryan. 1997. *Believable Agents: Building interactive personalities*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA. CMU-CS-97-123.
- Loyall, A. Bryan and Joseph Bates. 1993. Real-time control of animated broad agents. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Boulder, Colorado.
- Maes, Pattie. 1989a. The dynamics of action-selection. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, volume 2, pp. 991-997, Detroit, MI: Morgan Kaufman.
- Maes, Pattie. 1989b. How to do the right thing. AI Memo 1180, MIT AI Laboratory.
- Maes, Pattie. 1991. Behavior-based Artificial Intelligence. In Jean-Arcady Meyer and Stewart W. Wilson, (eds), *From Animals to Animats 2*, Cambridge, MA: MIT Press, 2-10.
- Massumi, Brian. 1992. *A User's Guide to Capitalism and Schizophrenia: Deviations from Deleuze and Guattari*. Cambridge, MA: MIT Press.
- Miedaner, Terrel. 1981. The soul of the Mark III Beast. In Douglas R. Hofstadter and Daniel C. Dennett, (eds), *The Mind's I: Reflections on self and soul*. New York: Basic Books, Inc., 109-115.
- Minsky, Marvin. 1988. *The Society of Mind*. New York: Simon and Schuster.
- Neal Reilly, Scott. 1996. *Believable Social and Emotional Agents*. PhD thesis, Carnegie Mellon University. CMU-CS-96-138.
- O'Neill Brown, Patricia. 1997. Setting the stage for the culturally adaptive agent. In Kerstin Dautenhahn (ed.), *Proceedings of the 1997 AAI Fall Symposium on Socially Intelligent Agents*, Menlo Park, CA: AAI Press. AAI Technical Report FS-97-02. 93-97.
- Porter, Tom. 1997. Depicting perception, thought, and action in *Toy Story*. In *First International Conference on Autonomous Agents*. Invited Talk.
- Schempf, Hagen. 1995. BOA: Asbestos pipe-insulation removal robot system. Technical Report CMU-RI-TR-95-15, Carnegie Mellon University Robotics Institute, Pittsburgh, PA.

- Sengers, Phoebe. 1994. Using culture to understand agents. In *AAAI Believable Agents Workshop* (Palo Alto, CA).
- Sengers, Phoebe. 1996a. Socially situated AI: What it is and why it matters. In H. Kitano (ed.), *Proceedings of the AAAI-96 Workshop on AI / A-Life and Entertainment* Menlo Park, CA: AAAI Press. AAAI Technical Report WS-96-03.
- Sengers, Phoebe. 1996b. Symptom management for schizophrenic agents. In *AAAI-96*. Menlo Park, CA: AAAI Press, vol. 2, 1369.
- Sengers, Phoebe. 1998a. *Anti-Boxology: Agent Design in Cultural Context*. PhD thesis, Carnegie Mellon University Department of Computer Science and Program in Literary and Cultural Theory, Pittsburgh, PA. CMU-CS-TR-151.
- Sengers, Phoebe. 1998b. Do the thing right: An architecture for action-expression. In *Proceedings of the Second International Conference on Autonomous Agents*, ACM Press, 24-31.
- Shakes, Jonathan, Marc Langheinrich, and Oren Etzioni. 1997. Dynamic reference sifting: A case study in the homepage domain. In *Proceedings of the Sixth International World Wide Web Conference*, 189-200.
- Sontag, Susan. 1979. *Illness as Metaphor*. New York, NY: Vintage Books.
- Steels, Luc. 1994. The Artificial Life roots of Artificial Intelligence. *Artificial Life*, 1(1-2):75-110.
- Stone, Brian A. 1996. Dynamically sequencing an animated pedagogical agent. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Portland, OR, 424-431.
- Thomas, Frank and Ollie Johnston. 1981. *Disney Animation: The illusion of life*. New York: Abbeville Press.
- Various. 1985. Martin Luther (trans.), *Die Bibel*. Deutsche Bibelgesellschaft, Stuttgart.
- Vere, Steven and Timothy Bickmore. 1990. A basic agent. *Computational Intelligence*, 6:41-60.
- Wise, J. MacGregor. 1998. Intelligence agency. *Cultural Studies*. 12:3.