Narrating System Intentionality

Copycat and the Artificial Intelligence Hermeneutic Network

ABSTRACT

Computer systems designed explicitly to exhibit human-like intentionality (seeming to be about and directed toward the world) represent a phenomenon of increasing cultural importance. *In the discourse about artificial intelligence (AI) systems, system intentionality is often seen as a technical property of a program, resulting from its underlying algorithms and knowledge engineering. By contrast, this article proposes a humanistic framework of the AI hermeneutic network, which states that along with any technical aspects, system intentionality is* narrated *and* interpreted *by its human creators and users. We pay special attention to system authors' discursive strategies in constructing system intentionality. Finally, we demonstrate the utility of our theoretic framework with a close reading of a full-scale AI system, Douglas Hofstadter and Melanie Mitchell's* Copycat*.*

 ISSN 1071-4391 ISBN 978-1-906897-16-1

by

Jichen Zhu & D. Fox Harrell

Jichen Zhu, PhD
Assistant Professor
Digital Media | The Antoinette Westphal College of Media Arts & Design
Computer Science | College of Engineering
Drexel University
jichen.zhu@drexel.edu

D. Fox Harrell
Massachusetts Institute of Technology
fox.harrell@mit.edu

1. INTRODUCTION

Human interaction with technical artifacts is often mediated by treating them as if they are alive. We exclaim, "my car doesn't *want* to start," or "my computer *loves* to crash." Yet, of increasing cultural importance are computer systems designed explicitly to appear intentional. These systems exhibit complex behaviors usually seen as in the domain of intentional human phenomena, such as planning, learning, and reasoning. Compared with more instrumental programs, such as *Adobe Photoshop*, intentional systems appear to produce output *about* or *directed at* certain things in the world rather than the mere execution of algorithmic rules.

Intentional systems are of particular relevance to digital arts and culture as they provide new ways of conveying meanings and expressing ideas. [1] Indeed, many salient examples of intentional systems can be found in music (e.g., George Lewis's interactive music system *Voyager*, Gil Weinberg & Scott Driscoll's robotic drummer *Haile*), visual arts (e.g., Harold Cohen's painting program AARON), storytelling (e.g., Michael Mateas and Andrew Stern's drama manager in *Façade*) and other cultural artifacts (e.g., *Pleo*, a popular robotic dinosaur toy). In all these examples, the inanimate computer systems seem to display beliefs, desires and other mental states of their own, whether through music notes, color palettes, playful behaviors, or otherwise.

The growing number of intentional systems requires a thorough understanding of the nature of intentionality in the context of computers and how such system intentionality is formed. Our discussion draws heavily on artificial intelligence (AI) as it is the technological foundation of intentional systems. The series of cross-disciplinary debates about AI between the 1980s and the early 1990s offer various insights of system intentionality at large. Among the various approaches, Daniel Dennett [2] argued that system intentionality

ISSN 1071-4391 ISBN 978-1-906897-16-1

connects to one of the most important strategies that humans use to predict the behaviors of other humans, animals, artifacts, and even ourselves – the intentional stance. This evolutionary skill requires the observer to treat those entities as rational agents with beliefs and desires in order to predict their potential behaviors. For instance, we may not know exactly how Pleo, the robotic dinosaur, is constructed internally, but we can nevertheless make sense of and predict its behaviors. We can interpret its visible behaviors as a manifestation of its emotions and desires such as a craving for "food" and attention. Dennett subsequently defined the systems to which we apply intentional stance as *intentional systems*.

This article aligns with the core of Dennett's theory, that is, the system intentionality is *not* a technical or ontological property of computers, as many computer scientists and theorists may believe. However, we propose to rethink the definition of intentional systems. Recent studies have also shown that people can apply the intentional stance to almost all artifacts. [3] For most people, certain digital artifacts (e.g. Pleo) afford intentional readings much more easily than others (e.g. *Photoshop*). This observation raises the question about the boundary of intentional systems and calls for further understanding of the phenomenon of system intentionality.

In this article, we introduce the AI *hermeneutic network*, a new framework to highlight system author's narration as an equally important element as users' interpretation in the formation of system intentionality. More specifically, we call critical attention to the use of *intentional vocabulary* as a key component of the author's discursive strategies in their narrations. This article is based primarily on part of Zhu's dissertation [4] along with several papers co-authored by Harrell and Zhu. We first present our theoretical framework informed by both the humanities and the AI community. Next, we introduce our new construct of the AI hermeneutic network, which argues that system intentionality arises from a complex meaning-making network that incorporates software authors' discursive narration and users' hermeneutic interpretation of system intentionality in a broad social context. Finally, we demonstrate the effect of the AI hermeneutic network through a close reading of a full-scale AI system, *Copycat*. In addition to the source code of *Copycat*, we look closely into a substantial corpus of the technical literature produced by the system authors, which is a rich and yet relatively unexplored area in software studies and critical code studies.

2. THEORETICAL FRAMEWORK

The topic of intentionality is of longstanding concern in philosophy. Reintroduced by Franz Brentano [5] in the late nineteenth century and later taken up by Husserl, [6] the concept is often considered as the linkage between the "inexistence" of human mental phenomena and the material establishments and states in the world. In modern philosophy, it is commonly understood as "aboutness," [7] defined as "that property of many mental states and events by which they are directed at or about or of objects and states of affairs in the world." [8] Intentional mental states, which include beliefs, desires and other states, are not free-floating thoughts, but are always *about* or *directed at* something in the world. This means that our beliefs and desires do not exist in abstract forms. Instead, they are *always* about certain states (e.g., I believe that it is going to rain tomorrow) or directed at certain objects (e.g., his desire for a flashy sports car).

Many scholars have insisted that intentionality is an intrinsic aspect of the human existence and not applicable to machines. A renowned argument, among others, is John Searle's "Chinese Room." [9] Searle argues that computers' ability to complete highly sophisticated tasks is not grounded in their intentionality as in humans, which is a prerequisite for intelligence. In comparison, Dennett's theory of intentional stance challenged the existence of intrinsic intentionality and asserted that intentionality is derived by the observer in the case of both humans and machines. It is hence possible for artifacts to display similar phenomenon. In the rest of this section, we continue this direction by examining several relevant works from both the humanities and the scientific community and discuss how our approach extends them.

2.1 Lessons from Alife

Artificial Life (Alife) is a research area that studies life and its process through computer simulation. It bears many resemblances to AI and intentional systems, for they all share the goal of constructing computer systems that display phenomena not commonly associated with machines – either aliveness or intelligence/intentionality. Critical understandings of Alife therefore are of particular use to unpack system intentionality.

Seeking to understand how it is "possible in the late twentieth century [for Alife researchers] to believe, or at least claim to believe, that computer codes are alive – and not only alive, but natural," N. K. Hayles [10] studied the Alife research community by "looking not only at the scientific content of the [Alife] programs but also at the stories told about and through them." She discovered that three levels of narratives are essential to the field. The first level includes "representations, authorial intention, anthropomorphic interpretation" of Alife computer programs. By observing how Alife researchers construct the narratives so that they are tightly interwoven into the operations of the program through terms such as "mother cell," "daughter cell," "ancestor," Hayles argues, "the program operates as much within the imagination as it does within the computer." Narratives at the second level, in comparison, are concerned with Alife as a legitimate research area within theoretical biology. In the pursuit of this goal, researchers frame their programs not as biology-inspired simulations, but as *life-as-it-could-be*, a more general framework containing traditional biology (i.e., *life-as-we-know-it*) as a special case. The third level establishes the relationship between Alife and the present and future of terrestrial evolution forms. Alife, according to the narratives, is not a simulation of the human, but rather a model to understand the prospect of human race. [11]

Hayles's analysis of Alife and what she calls "the nature and artifice of Artificial Life" demonstrates how Alife researches actively negotiate the meaning of their technical practice through complex discourses. Her methods of revealing the "multilayered system of metaphors and material relays through which 'life,' 'nature,' and the 'human' are being redefined" directly influenced our own framework of the AI hermeneutic network.

System intentionality, as argued below, is deeply discursive as well. Built on Hayles's work, our analysis further extends the notion of authors' discursive strategies to include addressing their technical practices. We argue that these narratives of system intentionality are not deceptive or conceal the "real" system operation. Instead, they are entrenched in the technical practice, both of which are constitutive elements of AI.

ISSN 1071-4391 ISBN 978-1-906897-16-1

ISSN 1071-4391 ISBN 978-1-906897-16-1

2.2 The "Epidemic" of Intentional Vocabulary

The use of intentional vocabulary, such as "reasoning," "planning," and "learning," is so pervasive in AI that it is almost impossible to describe any AI algorithm or system operations without using it. Theorist and AI practitioner Philip Agre once argued, "the purpose of AI is to build computer systems whose operation can be narrated using intentional vocabulary." [12] In the technical community, however, the significance of intentional narrations is generally not acknowledged.

How can AI practitioners ignore the significance of these intentional terms, which occur so frequently in their work? The answer is what Agre points out as the "dual character of AI terminology":

> *A word such as planning .. has two very different faces. When a running computer program is described as planning to go shopping, for example, the [AI] practitioner's sense of technical accomplishment depends in part upon the vernacular meaning of the word... On the other hand, it is only possible to describe a program as "planning" when "planning" is given a formal definition in terms of mathematical entities or computational structures and processes.* [13]

The elasticity of these terms to switch between formal and vernacular meanings has two fundamental functions. It enables the low-level machine operations to be connected to system intentionality through narrations. Equally important, it ensures that intentional narrations are seamlessly interwoven into the technical practice of AI.

Among the very few practitioners who noticed AI's reliance on intentional terms, Drew McDermott [14] considered these words misleading and harmful to the field and practitioners themselves. He criticized the use of intentional vocabulary as "wishful mnemonics," and saw it as "a major source of simple-mindedness in AI programs." Intending to stop this epidemic of "contagious wishfulness," McDermott urged his fellow practitioners that, instead of naming their programs "UNDERSTAND" or "THINK," all disciplined programmers should give them names that do not reveal their intended functions – such as "G0034." Then, the practitioners can decide if the operation of "G0034" still convinces themselves or anyone else that it implements some part of understanding or thinking.

Even Agre's attitude towards intentional vocabulary can be ambiguous at times. On the one hand, he denies that the strategic elasticity of these key terms is a conscious deception by the AI community. On the other hand, he admits such use of intentional vocabulary is "self-defeating" because these terms inevitably link AI to a much larger discourse based on reflections of their vague meanings. Consequently, AI practitioners will "find it remarkably difficult to conceptualize alternatives to their existing repertoire of technical schemata." [12]

The AI community's reservations about the intentional vocabulary seem to speak to a romanticized notion of "science," which portrays a "disciplined" practitioner as a neural channel between her subject and the knowledge she produces. Any subjective interpretation and narration need to be eliminated; after all, we should let nature speak for itself. The problem, however, is that nature cannot speak. Even in natural sciences, as Latour [15] cogently points out, part of a scientist's mission is to be the *spokesperson* for what is inscribed by her instruments. A practitioner's narration therefore can never be disassociated from her work completely.

The "contagious" use of intentional vocabulary is not because it is "deceptive," but because it is *necessary* to the practice of AI. A "G0034"-titled program without the narration of its author is like an unread result inscribed on complicated lab equipment, waiting for the scientist to be its "mouthpiece." Similar to the Alife researchers above, the AI practitioners' task is to create artifacts with certain properties that were deemed to be "sacred" to humanity. The elasticity of the intentional vocabulary provides AI practitioners with an effective discursive device, whether it is used consciously or otherwise. Without the glue of intentional vocabulary, in certain aspects, the field of AI would collapse. An AI practitioner's discursive construction of system intentionality through narrations therefore is as constitutive as her technical work.

ISSN 1071-4391 ISBN 978-1-906897-16-1

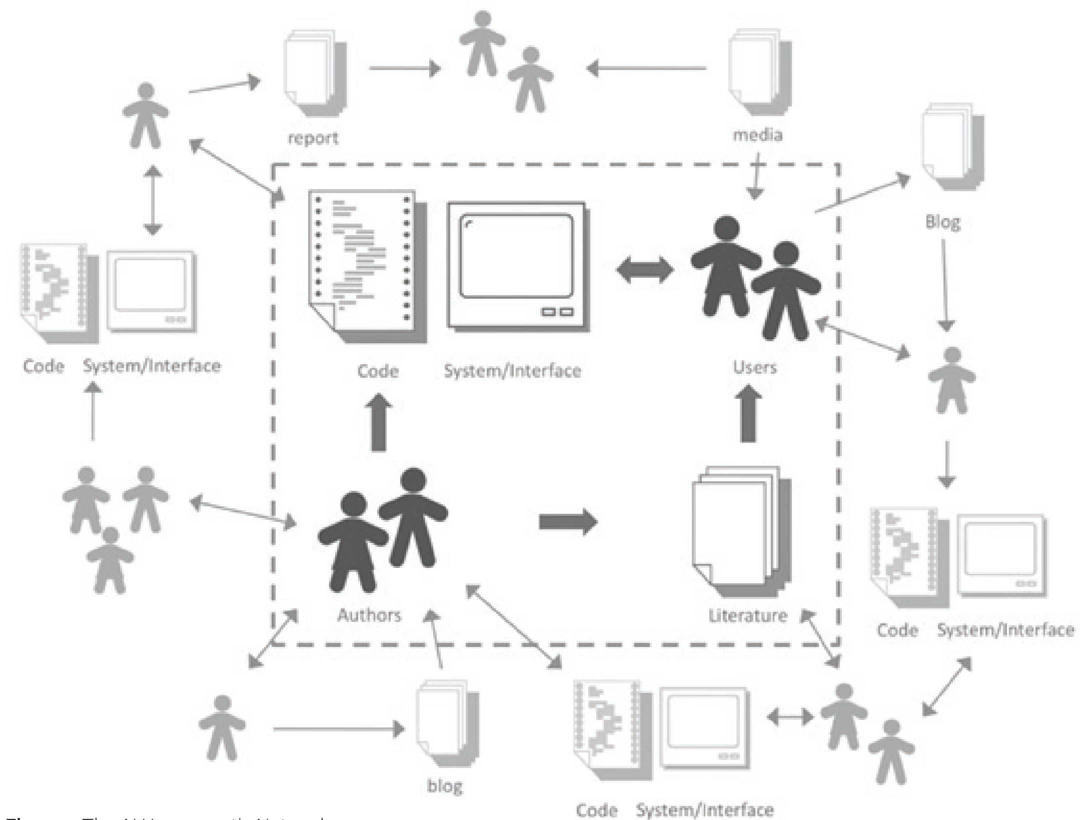

Figure 1. The AI Hermeneutic Network

3. THE AI HERMENEUTIC NETWORK

Our framework of the AI hermeneutic network (Fig. 1) argues that system intentionality arises from a hermeneutic communication process. It incorporates two equally important components: the system author's discursive narration and the user's hermeneutic reading of the system in their respective contexts. In this framework, authors and users negotiate meaning through both the system (e.g., source code, interface, and system output) and the related discourse (e.g., technical publications, media coverage, and authors' blogs). [16] A more detailed explanation of this framework can be found in Zhu's dissertation. [17] In this article, we focus on the author's narration, in particular their use of intentional vocabulary.

A system author's narration, as argued above, should be distinguished from a kind of subterfuge story that obscures rather than explains system function. Unlike fictional fairy tales, practitioners' narrations of system intentionality are an indispensable element of AI. When a practitioner claims that her system is capable of "planning", what is at work is that the term's formal meaning temporarily takes over its vernacular signification. When a lay user, or sometimes an AI practitioner herself, encounters such discourse relating to the system, she may take on the vernacular meaning of "planning" as a lens to interpret system operation.

ISSN 1071-4391 ISBN 978-1-906897-16-1

In the meantime, users bring their own experiences and social/cultural backgrounds in order to make sense of computer systems. Researchers in information studies, for instance, have conducted ethnographic studies of how users hermeneutically read quantitative data provided by information systems and contextualize these "cold and objective categories and numbers" with the real-life situations. [18] As part of a feedback loop, users' collective experiences with intentional systems will shape our society's dominant view of intentionality and intelligence, which in turn may be incorporated by AI researchers into their evolving formal definition of the key intentional terms.

Many other social agents, or "actants" in the terminology of actor-network theory also participate in the hermeneutic meaning exchange process between the system author and the user. [19] Funding agencies, mass media corporations, government policies, economic developments, other similar AI systems, etc. are all part of the network in which the hermeneutic communication of system intentionality takes place. In many cases, the clear distinction between author and user evaporates; many users of AI research projects are AI practitioners themselves who are developing related systems.

4. CASE STUDY: A CLOSE READING OF COPYCAT

In order to demonstrate the utility of the AI hermeneutic network, this section presents a close reading of *Copycat*, a full-scale AI system. The *Copycat* project was developed by Douglas Hofstadter and his Ph.D. student Melanie Mitchell between 1984 and 1995. Its research goal is to gain better understanding of the human analogy making process by building computer models. The domain of *Copycat* is alphabetic analogical mapping problems, such as "abc → abd, pqr → ?" ("if *abc* is to *abd*, then *pqr* is to what?") A correct answer is "*pqs*," after replacing the last letter with its alphabetic successor.

In particular, Hofstadter and Mitchell's focus is to model the human "mental fluidity" by algorithmically constructing not only correct, but "insightful" analogies. Another correct answer to the problem above is "pqd" (always replacing the last letter with "d"), but most people find this analogy less insightful. To fully appreciate the type of *analogical fluidity* the system authors intend to achieve, it is worthwhile to look at a more complex problem: "abc → abd, pqqrrr → ?" Following the logic of the previous problem, one may answer "*pqqrrs*," "*pqqsss*," or "*pqqddd*." However, if we convert each letter in the first pair into its alphabetic position (e.g. "a" is 1, "b" is 2...) and use that number to determine how many occurrences each letter has in the second group (e.g., once for "p", twice for q...), we get a new answer of "pqqrrrr." To many, the last solution is more insightful because it *fluidly* maps the concept of "alphabetic position" to "group size." This type of analogy is what *Copycat* was constructed to perform.

4.1 The Corpus and Method

In addition to the source code of Copycat, we include over 200 pages of major technical articles and book chapters published by the authors. [20] We also incorporate additional material of Hofstadter's interviews with the mass media, non-technical articles, and personal websites, all of which provide us with the social context of the project and the authors' ideological/philosophical positions on issues related to intentional systems.

Our method is threefold. We first analyzed the technical, social, cultural contexts in which *Copycat* was built. Next, we performed a close reading of the corpus. Each article was carefully analyzed to identify the authors' rhetorical strategies to construct system intentionality. Each instance of discursive strategies was documented and carefully scrutinized to identify the patterns and trends. Finally, we contextualized the authors' narrations in their ideologies and beliefs in AI. Are they firm supporters of the Strong AI hypothesis? How do they compare the operation of their system in regard to related human cognitive processes? These are important questions to help us further connect the discursive strategies identified from the previous step to the authors' grand goals.

4.2 The Two Languages of Copycat

In this article, we present the primary results from our close reading of the corpus. Overall, we identified two paralleling languages, one intentional and the other formal/technical, co-existing simultaneously in Hofstadter and Mitchell's discourse of *Copycat*. In order to draw contrast to them, we first artificially separate them into two different narratives of *Copycat*. Nevertheless, we do not suggest the separate existence of an "objective" technical language and discursive one. These two semiotic systems are tightly intertwined with and dependent on one another. The two languages of *Copycat* are aligned with Michael Mateas's observation of the coexistence of code machine and rhetorical machine. [21] Both works argue that the technical practice of AI is intrinsically discursive.

A Stochastic Local Search Program. *Copycat* is a stochastic local search program. It receives three character strings (*String 1*, *String 2*, and *String 3*) as input and generates a single output character string. During the process, *Copycat* performs a stochastic local search in a particular search space, optimizing a certain heuristic function. Its search space is all possible *structures* that can relate the three input strings together. Each structure is a graph built from a set of predefined *primitive constructs*, such as "b is the successor of a." A structure captures the relations between the three input strings and determines the compatibility of the primitive constructs appearing inside. *Copycat* maximizes a heuristic function, that is, the extent to which the proposed structure captures all the regularities and relations among the three input strings. The system may terminate its search at any point in time. Based on its heuristic function, the better the structure constructed by the system, the higher the probability of termination. Once the search stops, the system generates an output string according to the transformation operations specified by the current structure. This means that the same operations that transform *String 1* to *String 2* will be applied to *String 3* in order to derive the output string.

A Fluid Analogy Maker. Intermingled with the technical discourse, the intentional narration portrays *Copycat* as a fluid analogy maker, emphasizing the program's human-like psychological plausibility in its algorithmic operation. Modeled on fluid human analogy-making process, *Copycat* is capable of constructing "insightful" analogies through the "slippage" of concepts from one to another, such as from "alphabetic order" to "group size." In its process, *Copycat* simultaneously deploys many small pieces of code, called *codelets*, to perform various tasks such as creating or destroying a structure, evaluating how promising a particular structure is, and creating more codelets. They can be seen as the enzymes in biological cells, where each enzyme does only one very small task, but the combination of thousands of them manages to fulfill complex tasks.

One example of how these two languages intermingle is the "happiness" level of each of *Copycat*'s many codelets. Once any codelet's happiness level drops below a certain threshold, the system will dedicate more resources to it. Most people, including AI practitioners, will be quick to agree that "happiness" is an explicit intentional term. Unlike "planning" or "learning," its highly subjective and emotional undertone stands

 ISSN 1071-4391 ISBN 978-1-906897-16-1

ISSN 1071-4391 ISBN 978-1-906897-16-1

in stark contrast to any formal definition based on machine operation. This rather bold choice of terminology, however, contributes to *Copycat*'s psychological plausibility, one of its authors' main research goals. The term "happiness" not only lends itself to one of the most common human psychological states, but also makes possible the intended narrative that *Copycat*'s operation follows the common wisdom that "the squeaky wheel gets the oil."

Of course, any intentional narration completely disconnected from the technical operation will be just a castle in the air. The architectural design of *Copycat*'s happy codelets is also technically sound because an "unhappy" element corresponds to a weak structure. By focusing on the weakest point in the ensemble, the system maximizes the chance of improvement.

4.3 Three Main Discursive Strategies

The example of codelets' "happiness" levels shows that the two language systems of *Copycat* are connected to each other through the use of intentional vocabulary. In the main corpus, we identified three discursive strategies frequently deployed by the system authors to connect its algorithmic operation and system intentionality. Examples of each strategy can be found in Fig 2.

First, intentional verbs are heavily used throughout the corpus to describe *Copycat's* underlying operation. Some frequently used words in the primary corpus include "know," "resist," "understand," "prefer," etc. A representative example is "*Copycat tends to resist bringing numbers into the picture, unless there seems to be some compelling reason to do so.*" These intentional verbs effectively portray the system as an intentional entity, with its own belief and desire to come up with insightful analogies.

Second, some of *Copycat*'s key data structures and functions are referred to using names of human cognitive faculties and human mental states. For instance, the possible analogical structures that *Copycat* constructs are called "point of views." Similarly, *Copycat* has "long-term memory," "drive," "desire," and "personality." These terms not only establish close ties between the system's operations to human cognitive processes, but also indirectly contribute to its system intentionality.

Discursive Strategies	Examples
intentional verbs	"Just as the program *knows* the immediate neighbors of every letter in the alphabet ..."
	"Copycat tends to *resist* bringing numbers into the picture ..."
human cognitive faculties & mental states	"[The Slipnet] can be thought of, roughly, as *Copycat*'s *long-term memory*."
	"*Copycat* must reconcile a large number of mutually incompatible local *desires* (the technical term for this is '*frustration*')."
	"... and those data provided some of the most important insights into the program's '*personality*.'"
benchmarking human & other life forms	A (creative) human faced with situation X will react with action Y, and *Copycat* also performs action Y in this situation X.

Figure 2. Three Main Discursive Strategies (Emphasis Added)

 ISSN 1071-4391 ISBN 978-1-906897-16-1

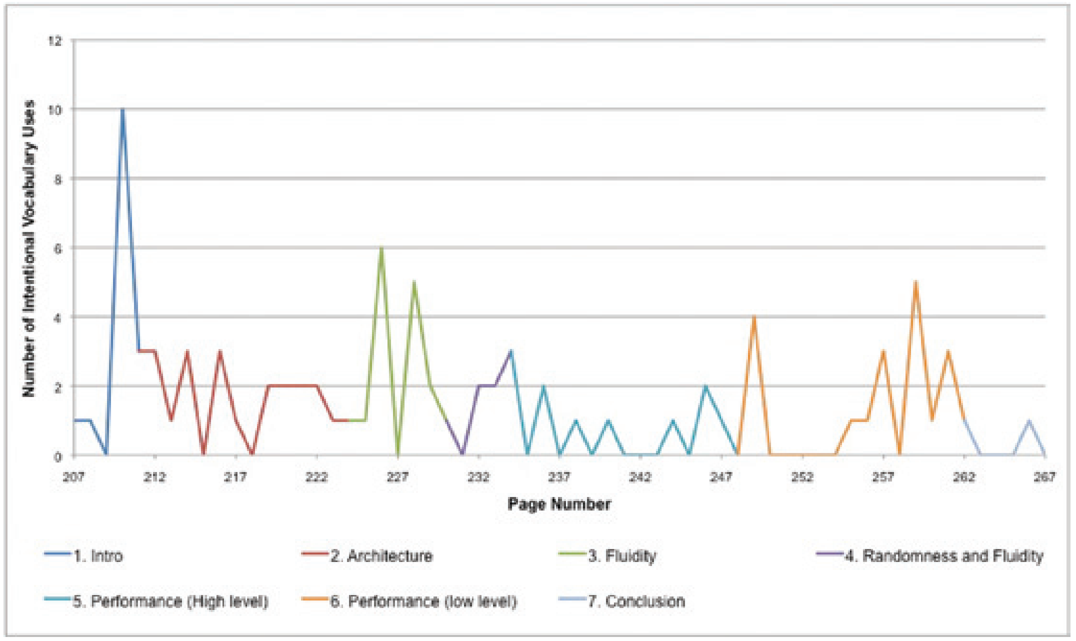

Figure 3. The Frequency of Intentional Narration (in Hofstadter and Mitchell 1995)

Finally, and more subtly, the system's operation is often benchmarked with human and other forms of life (e.g., cells and ants). Although the content may vary depending on the context, these arguments typically take the following form: A (creative) human faced with situation X will react with action Y; *Copycat* also performs action Y in this situation X. The underlying implication is that *Copycat*'s operation can be seen as similarly intelligent, intentional, and creative as human. An example is:

> *In particular, people are clearly quicker to recognize two neighboring objects as identical than as being related in some abstract way.. [*Copycat*] tends to spot them [(neighboring identical objects)] and to construct them more quickly than it spots and constructs bonds representing other kinds of relationships.* [22]

In addition to the types of discursive strategies, we also identify when the authors rely more on intentional narrations. Fig. 3 illustrates the frequency of intentional narrations – defined as the number of intentional vocabulary on each page – in different sections of Hofstadter and Mitchell's article. [23] It shows that when the authors set the goal of *Copycat* (in the "Introduction" section) and discuss the fluidity of its analogies (in the third section, whose full title is "The Emergence of Fluidity in the Copycat Architecture"), they tend to use intentional vocabulary more often. The sections with lower frequency (Section 2 and 5) are concerned with topics of system architecture and performance.

ISSN 1071-4391 ISBN 978-1-906897-16-1

5. CONCLUSION

In summary, this article has introduced a humanistic and interpretive framework to analyzing intentional systems through the AI hermeneutic network. Different from seeing system intentionality as an intrinsic (technical) property of software, we highlight an actor-network of which software is just one component. Here, we primarily focus on the technical literature surrounding software systems, so far a relatively unexplored area in software studies.

By applying this framework to a full-scale AI system, we have identified various discursive strategies that the system authors used to narrate the system intentionality of *Copycat*. The use of intentional vocabulary, as we have shown, connects the discursive and technical requirements of the system. In this regard, the practice of AI is fundamentally technical *and* discursive at the same time. This often-neglected discursive aspect of software stresses the importance of critical understandings of complex technological artifacts.

As part of our future work, we plan to apply the AI hermeneutic network to other AI systems and to expand it into a more general framework of software hermeneutics. Certainly, many of the issues pertinent to AI can also be applied to the broader domain of software. The burgeoning area of Software Studies, explored by researchers such as Lev Manovich, Noah Wardrip-Fruin, Mathew Fuller, Mark Marino, Jeremy Douglass, and others, points to a recognition of the need for software studies methods. Part of our contribution is to critically analyze the practice of AI from the vantage point of an insider-outsider. Just as an ethnographer who, when living in a different culture, must (ideally) become a member of the group being studied, we base our work on our experiences as AI practitioners. Yet, we also bring in techniques and lenses from afar. Our critique of, and approach to, AI and intentional systems are informed by the lenses of textual analysis, literary theory, and related approaches from the humanities. The resulting concept, the AI *hermeneutic network*, is the basis for an approach that contextualizes computational systems in a broader network of discourse practices, human interpretation, and social and cultural practices of information dissemination and exchange. ■

ISSN 1071-4391 ISBN 978-1-906897-16-1

REFERENCES AND NOTES

1. Jichen Zhu and D. Fox Harrell, "Daydreaming with Intention: Scalable Blending-Based Imagining and Agency in Generative Interactive Narrative," in *Proceedings of the AAAI 2008 Spring Symposium on Creative Intelligent Systems* (Stanford University, CA, 2008).
2. Daniel C. Dennett, *The Intentional Stance* (Cambridge: The MIT Press, 1987).
   Daniel C. Dennett, *Kinds of Minds* (New York: Basic Books, 1996).
   Daniel C. Dennett, "True Believers: the Intentional Strategy and Why it Works," in *Mind Design II: Philosophy Psychology Artificial Intelligence*, ed. J. Haugeland (Cambridge: The MIT Press, 1981).
3. Nigel Seel, *"Agent Theories and Architectures ,"* (Ph. D. dissertation, Surrey University, Guildford, UK, 1989).
4. Jichen Zhu, "Intentional Systems and the Artificial Intelligence Hermeneutic Network: Agency and Intentionality in Expressive Computational Systems," (Ph. D. dissertation, Digital Media, Georgia Institute of Technology, Atlanta, GA, 2009).
5. Franz Brentano, *Psychology from an Empirical Standpoint* (London: Routledge and Kegan Paul, 1874).
6. Edmund Husserl, *Logical Investigations* (London: Routledge, 1970. Original edition, 1900).
7. Philip E. Agre, *Computation and Human Experience* (Cambridge, U. K.: Cambridge University Press, 1997).
   Daniel C. Dennett, *The Intentional Stance* (Cambridge: The MIT Press, 1987).
   Terry Winograd and Fernando Flores, *Understanding Computers and Cognition* (Norwood, NJ: Ablex, 1986).
8. John Searle, *Intentionality: An Essay in the Philosophy of Mind* (Cambridge: Cambridge University Press, 1983).
9. Ibid.
10. N. Katherine Hayles, *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics* (Chicago, IL: University of Chicago Press, 1999).
11. Ibid. , 224-239.
12. Philip E. Agre, "Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI," in *Social Science, Technical Systems, and Cooperative Work: Beyond the Great Divide*, ed. G. C. Bowker, S. L. Star, W. Turner and L. Gasser (Mahwah, NJ: Lawrence Erlbaum Associates, 1997).
13. Ibid.
14. Drew McDermott, "Artificial Intelligence Meets Natural Stupidity," *SIGART Bull*, No. 57 (1976): 4-9.
15. Bruno Latour, *Science in Action: How to Follow Scientists and Engineers Through Society* (Cambridge, MA: Harvard University Press, 1988).
16. Jichen Zhu and Fox Harrell, "The Artificial Intelligence (AI) Hermeneutic Network: A New Approach to Analysis and Design of Intentional Systems," in *Proceedings of the 2009 Digital Humanities Conference* (College Park, MD, 2009).
17. Ibid.
18. R. J. Boland, "Information Systems Use as a Hermeneutic Process," in *Information Systems Research: Contemporary Approaches Emergent Traditions*, eds. H. -E. Nissen, H. K. Klein and R. Hirschheim (New York: North-Holland, 1991).
19. Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network-Theory* (Oxford, U. K. : Oxford University Press, 2005).
20. Douglas R. Hofstadter, *The Copycat Project: An Experiment in Nondeterminism and Creative Analogies* (Cambridge, MA: AI Laboratory, MIT, 1984).
   Douglas R. Hofstadter, "Perspectives on Copycat: Comparisons with Recent Work," in *Fluid concepts & creative analogies*, ed. D. R. Hofstadter (New York: Basic Books, 1995).
   Douglas R. Hofstadter and Melanie Mitchell, "The Copycat Project: A Model of Mental Fluidity and Analogy-Making," in *Advances in Connectionist and Neural Computation Theory: Analogical Connections*, eds. J. A. Barnden and K. J. Holyoak (Norwood, NJ: Ablex, 1994).
   Douglas R. Hofstadter and Melanie Mitchell, "The Copycat Project: A Model of Mental Fluidity and Analogy-Making," in *Fluid Concepts & Creative Analogies*, eds. D. R. Hofstadter (New York: Basic Books, 1995).
21. Michael Mateas, *Interactive Drama, Art, and Artificial Intelligence* (Ph. D. dissertation, Computer Science, Carnegie Mellon University, Pittsburgh PA, 2002).
22. Douglas R. Hofstadter and Melanie Mitchell, "The Copycat Project: A Model of Mental Fluidity and Analogy-Making," in *Fluid Concepts & Creative Analogies*, eds. D. R. Hofstadter (New York: Basic Books, 1995).
23. Ibid.

ISSN 1071-4391 ISBN 978-1-906897-16-1