PREFACE

The imagination is at the root of many things we see as social reality. A "phantasm" is a particular pervasive kind of imagination, one that encompasses cognitive phenomena including sense of self, metaphor, social categorization, narrative, and poetic thinking. This type of imagination influences almost all our everyday experiences, across diverse domains of experience, including art, entertainment, commerce, culture, and power relationships. This book presents an approach to understanding and designing computing systems that evoke these ubiquitous and often-unseen phantasms—with special attention paid to revealing oppressive phantasms and creating empowering phantasms.

Self-Reflection and Nostalgia

This book grew out of a childhood idea. Television, everyone said, was already compromised—it rotted minds and turned people into couch potatoes. So I wondered if there would be any way that I could help whatever medium came next to play a more ethical role in society. I cannot say that this book lives up to the dream that musing sparked. Yet this book might be considered a step in pursuing a more mature vision rooted in that early dream, with its dual aims of helping people see the unconscious forces that shape media images and helping them use computational media—the next media after television, one could say—to affect society for the better.

As a teenager, I learned of college majors such as "logic and computation," "symbolic systems," and "cognitive science." How wonderful it sounded to combine computer science with fields such as philosophy, mathematics, and psychology. Those majors also all involved artificial intelligence (AI). What a dangerous-sounding area of study—all wrapped up in the aesthetic sensibility of a computer game I loved called *Neuromancer*. By earning a hybrid computer science/mathematics/philosophy degree at the same time as a degree in art (which included learning to work with images, text, sound, video, computer-based media, and more), I would not have to leave any of my educational interests behind. As a small step before crossing the country to attend the university where I executed this plan, I joined a special-interest AI group with the engineers and academics of San Diego. I was the only youth, and a different tint and shade than the other members; I suffered a few askance looks, meanwhile their spouses offered me cookies.

In college, pursuing research in interactive narrative or AI agents seemed to be the best way to blend art and science. Interactive narrative research would allow the computer to help people tell stories set in dynamic worlds from multiple perspectives—we could see the world through the eyes of others in imaginative realms where tender moments constantly decay and die or memories emerge like old bones from tar pits. Research on AI agents would allow programmers to implement characters that have their own lives—characters who represent different modes of thinking. Artificial Intelligence is concerned with topics such as representing meaning, so artificial intelligence–based art seemed like the most artful of computer science fields. At Carnegie Mellon University, I was fortunate to be in a place where such research was pursued.

I chose interactive narrative from those possibilities; my path seemed set. Yet the types of meaning I was interested in were mostly left out of all that I encountered in that area. I felt that most of the work I saw in interactive narrative was focused on a single goal: the Holy Grail seemed to be immersing ourselves in virtual worlds and participating in stories as main characters. I had a passion for a different kind of question: could a computer help people imagine in new ways? When taking an introductory course on the history of art and civilization framed exclusively through the lens of European history, I daydreamed of implementing fantastic interactive narratives on the computer that incorporated what I was learning, creating scenes like this:

> Rising from the depths of a lake filled with thinned white oil paint, I encounter specters with faces that looked like Caravaggio's chiaroscuro figures, exaggerated

> shadows, intense and ecstatic gazes. I encounter other ghosts, long and streamlined like Brancusi's sculptures. Touching them transforms me; I am a spooky conglomeration of trans-European hobgoblins, cherubs, and Flemish nobles in a fluid white demesne. As I approach the surface, feeling more and more buoyant, I make one final choice that will ultimately determine my own image, reflected back down at me from the water's impending surface I see myself. I now look like . . .

Unusual daydreams like that were rooted in a reactionary sentiment[1] to a reoccurring experience I had: I was simultaneously inspired by, and excluded from, the cultural milieu in which I was immersed. My own interests in engaging intellectual and artistic traditions from a breadth of cultures were engaged only on the periphery—a single reading in a conceptual art course here, an independent study there. The transformative and inspiring references I found on my own during late nights in libraries were not part of the cultural canon. In computer science, interactive narrative research started with a limited range of types of interaction, rather than attempting to implement customized forms of interaction suitable to the meanings at hand. I asked myself: if poets can create finely crafted works based on evocative imagery, why should I not use the computer to create interactive narratives filled with evocative, poetic imagery, like the previously mentioned daydream? Much of the technology was already available: one could use a fluid dynamics model to simulate the viscous liquid of the translucent white lake, ray-tracing to address diffusion and refraction effects, a virtual camera to simulate movement continuously upward along the z-axis, and European art historical images as skins for 3D graphical models for nonplayer characters (NPCs), and one could build a simple finite-state machine to track how a player character transforms based on which NPCs it encounters. Sure, the models do not generalize, but could not that technology express the *meaning* and *subjective experience* I wished to convey through metaphor and narrative?

Beyond such abstract musings, I also had a different set of more conventional research questions. Could one algorithmically retell stories from multiple cultural perspectives? What about with different feelings? Different imagery? Could all of this be accomplished with the power, nuance, and imagination of a book such as Ralph Ellison's *Invisible Man*? Ellison's critically acclaimed novel continues to have an impact on U.S. society and has entered into the canon of world literature. It compels readers to share the experience of navigating an imaginative world of social invisibility, of being on the margins. The novel weaves together profound social critique,

experimentation with prose style and narrative structure, imaginative (even borderline science-fiction) events, and verbal imagery. When I read research papers and books about interactive narrative, the development of computational models engaged with imaginative and social meaning (modes of expression akin to Ellison's) was not considered a research aim. In short, I wanted to focus on content and content creation. Doing so meant really coupling representations of *meaning* (able to address issues of theme, social empowerment, emotion, culture, and aesthetics) with appropriate AI and interaction models. It was a focus that was not en vogue.

A breakthrough came when I discovered that some cognitive scientists were interested in the types of meaning that I was seeking. Mark Turner's book *The Literary Mind* was not computationally oriented, yet he wrote on topics such as narrative, parable, metaphor, and blending concepts. This book, and the books in cognitive linguistics to which it led me, connected the idea of artistic thinking to everyday thought and life. I learned that these research topics were the domain of the area of *cognitive semantics*—a subfield of cognitive linguistics, which is in turn a subfield of cognitive science. The field was little recognized in the AI, mathematical logic, and analytic philosophy I knew, nor was it commonly cited in my art training. I encountered criticisms about the field and from the field and defenses from those criticisms all around. I had unwittingly walked onto a battlefield within linguistics. Regardless, these were kindred spirits whose topics supported my needs. My challenge, then, was to figure out how this field's insights could be combined with the AI and logic that I knew well (and against which the cognitive semantics researchers so often seemed to argue) and then applied to artistic and cultural aims.

Urgently seeking, I came upon a theory called *algebraic semiotics*, invented by a computer scientist and mathematician named Joseph Goguen, whose research group was called the Meaning and Computation Laboratory at the University of California, San Diego (UCSD), a kindred spirit with whom I began to correspond. Algebraic semiotics combines insights from the area of algebraic semantics in computer science with semiotics and applies them to topics such as user interface design and cognitive science. Algebraic semiotics is unique as a theory because Joseph was always careful to discuss the limitations of mathematical-computational formalisms at the same time he was using them to think about and model the most poetic types of thought and language. Furthermore, UCSD was an oasis for cognitive semantics. Gilles Fauconnier, another visionary thinker about thought, was also there. Joseph, whose work investigated mappings in mathematics (among a multitude of other subjects), was friends with George Lakoff, whose work addressed one of the main types of meaning

I sought to learn more about: metaphor—mappings in cognition. Though loathe to return to beautiful, oppressive San Diego, where I grew up, to the very campus my parents attended, I went to UCSD to apprentice with Joseph. It was fortunate and tragic. We found great synergy and he embraced me as his Ph.D. student in computer science and engineering. However, he died before we could finish our work—just days after an international symposium held in his honor (occasioned by the publication of a Festschrift on Joseph's work called *Algebra, Meaning, and Computation*). Chapter 4 of this book is a tribute to Joseph's work. I hope it will help others see what I found so profound about algebraic semiotics.

Joseph was always a strong supporter, yet he finally conveyed to me a deep appreciation of the value and depth of my vision when we successfully developed a modest conceptual blending algorithm (which we later named Alloy) that I used to author a program capable of generating a new poem every time it was used. When I was invited to speak at a conference focused on issues of technology and social power relationships (with a feminist bent), I presented the system and its implementation to help make the argument that accounts of race and ethnicity needed to move beyond binary oppositions and instead focus on subjective and dynamic experiences. A year later, I sketched out an architecture diagram to generalize the operation of that first interactive poem, which grew into the GRIOT system.

The GRIOT system is a computer program developed as a platform to support the implementation of works of interactive narrative, poetry, and related types of interactive multimedia expression. It is a research system, a project initiated to explore a set of ideas about how to make storytelling on a computer more improvisational. In particular, central challenges were to develop technology focused on composing content that is meaningfully different each time and theory to explain just what "meaningfully different" means. It turned out that the outcome was not specifically focused on narrative, but rather on phantasms: outcomes of the more general aspects of imaginative cognition that include narrative imagining, poetic discourse, and figurative thinking.

Realizations and Hopes

Ultimately, I realized that AI-based narrative was not my main interest. Some of my work at the time looked like (interactive) surreal poetry; other work was more straightlaced user interface design, though driven by user values. Though some work

was highly narrative, other work was more metaphorical and associative (and, I hoped, evocative). Yet I realized that there was one strand that went through all of it: the imagination.

I deeply believe that the human capacity to imagine is our greatest means to freedom. I do not believe that great art and forms of culture should deal with only lofty issues of life, death, love, and beauty. I do believe in the functional value of art to change society and minds—and to understand how they should change. At the same time, I am concerned about social ills encountered in everyday life: violence, atrocities of war, small acts of unkindness, prejudice. These are not unique concerns; they are powerful phenomena that all sensitive people care about. I see these topics as part of everyday life, just like death, love, and beauty. Furthermore, I also see that the ways we conceive of these very real social ills are shaped by the human imagination. The ultimate function of the kind of cultural production I am interested in, whether one calls it art or engineering, is to reveal and shape the combination of imagination and real-world experiences that oppress or empower.

The focus on phantasms prompted and revealed by media is both an outcome of this journey and a lens through which others can begin to see my perspective on all of these topics. One of the things I work hardest at is making sure that my ideas are synthesized into a coherent whole rather than remaining discrete interests. I constantly chart, map, and plan ideas. This book has been a great challenge of synthesizing, charting, mapping, and planning, all wrapped up in the word *phantasm*. This topic is, perhaps, unheeding of many trends of current art practice and may seem rather "soft" to many engineers and scientists. Understanding and designing computing systems to prompt the imagination poses several challenges. In computer science, where researchers often prove ideas through making systems with techniques and technologies that rapidly become dated, one challenge is to produce theory that stands the test of time. Cognitive science is a field in which researchers propose theoretical models that also will one day be superseded by new theories with more rigorous empirical grounding as we learn more about the brain. Furthermore, this book is interdisciplinary and draws upon a range of ideas that I encountered and pursued as rigorously as I could, not only through traditional paths of achieving academic expertise but additionally just through reading and life experience (remember, that was part of the plan).

Despite these challenges, I hope that the book will be useful for researchers and other interested people who want to think about the computer, imagination, social reality, and cultural expression from a new perspective. I have given some attention

to making the book useful in a practical way as well as presenting new theory. Each chapter of the book offers a clear set of constructs (a model) that should be useful for enriching understanding and guiding design of computing systems.

I also realize a flaw of my childhood dream: the medium alone is not what is ethical or unethical. The computational medium that I work in, which has indeed come after television, is connected to a complex web of previous media, culture, behavior, preexisting power relationships, and more. From one perspective, I submit this book to you as a wistful, probably doomed, love letter—hoping for an "I love you too" (your theories make sense, and open up new possibilities), but I am prepared for my ideas to receive the fate of all words of unrequited lovers. That is, even if it ends up not laying a blueprint for the future, I strive for this book to be good poetry.

More assertively, I would say that this book is a manifesto arguing that great expressive potential of computational media comes from their ability to construct and reveal what I call phantasms. It argues for the importance of cultural content, diverse worldviews, and social values in computing. Designing and analyzing such epistemic functions of computational media is key to unleashing their expressive power. The expressive function of phantasms, those blends of cultural ideas and sensory imagination that the computer can so effectively conjure, is not restricted to purely aesthetic dimensions. More substantively, phantasmal media can express and construct the types of meaning central to the human condition.