Preface

Working Through Synthetic Worlds. The title and the content of this book are intended to initiate a new wave of research to advance the general state-of-the-art and use of immersive worlds. Before we begin, however, we need to describe what we mean by "working" and "synthetic worlds."

Over the past several years, a number of new Internet based immersive worlds (i.e., virtual worlds) have grown in popularity. Second Life (www.secondlife.com) logs millions of simultaneous users on a daily basis. Massively Multi-User Online Games (MMOGs) such as World of Warcraft are becoming ubiquitous. It is clear that games and social interaction are greatly facilitated by virtual worlds but we believe that the time is right to start taking a serious look at the future potential of these environments beyond games and social networks. We need to ask questions like: Just what are virtual worlds good for? What value is there in these virtual communities and what potential value could there be for doing real work? Where might one strategically invest in research with these technologies to influence their development to make people more productive?

So by "working" we mean: How could virtual world technologies be used as a tool to assist in transforming the traditional workplace? How can we use them to facilitate information sharing and collaboration across a diverse workforce that might be geographically and temporally separated? What does a virtual environment afford that offers a favorable value proposition compared to the real world?

A common general term for immersive world technology is the Metaverse which can be partitioned into four sub-areas:

1. Virtual Reality—represents a fantasy world for role-playing, gaming, and social networking.
2. Mirror Worlds—provides a digital representation of the real world.
3. Augmented Reality—overlays the real-world with digital information.
4. Life Logging—stores personal real-world events for later sharing.

Unfortunately, performing "real" work does not fit completely into any of these sub-areas even though some work may rely on virtual reality or mirror worlds. The authors will introduce work environments that require complete in-world control over time, size, and levels of abstraction. These dynamically created (or synthesized) worlds will provide the foundation for extending current immersive worlds into the work domain. So how do we move from fantasy worlds to worlds for work?

While the technology of virtual worlds is rapidly maturing, research related to transforming it into a useful tool is lagging behind. Most research is documenting what has been done, or speculating what could be done, rather than asking what *should* be done. It is in this context that we asked our professional colleagues, in academia, in

government, and in commercial ventures, "If you could work in a Synthetic World, what should it look like?"

We discovered immediately that our colleagues were brimming with ideas. Perhaps too many ideas! So, we started to constrain the question. We developed a set of instructions for prospective authors. These instructions included constraints on the formatting of the chapters, and the requirement to present their visions initially as a narrative. We redirected to other outlets a large number of proposals that sought to discuss gaming or social networking. About half of the initial proposals focused on the use of virtual environments for training. We encouraged these authors to look past the training or educational use of their synthetic world, towards a commercial application of the skill they sought to train. Since training is apparently a natural use of these technologies, we kept the best proposals of that genre, and we present them as one section of this book.

The result is a book that offers 20 distinct visions of the workplace of the future. And although the visions are completely original, and describe a large number of task domains, they share some remarkable similarities. We discuss one of these similarities, the ubiquitous presence of software agents, later in this preface.

We have designed this book to have a wide appeal. We expect that the information contained in this book will be accessible to academics, practitioners, managers, as well as people who are simply curious about what the future holds. Because the book deals with the application of technology to work, we expect that specific disciplines will have a special interest in the topic. These disciplines include both scientists and engineers involved with technology development (computer scientists, information scientists, computer engineers, human factors engineers, etc.) and also social, behavioral and cognitive scientists who are interested in the effects of technology on people and groups. We think this book would be an excellent resource for graduate seminars in any of these disciplines.

The format of the book was selected to maximize its accessibility to various user communities. We divided each chapter into three sections to enable efficient browsing for specific purposes. The first section of each chapter is a story about a day in the work life of a person in the year 2025. At some point in each story, the person is working in a Synthetic World. The stories present different visions of the ways that technology will influence our work lives. By browsing the stories, our readers can sample from among these visions, and decide for themselves which visions are likely to become reality. There are several common themes that emerged from these stories: What is the real value of a virtual environment for work? What roles will software agents play in this work environment? What is the unique contribution of the human?

The second section of each chapter is intended to be a scholarly examination of the vision illustrated by the story. In this section, the authors review the current literature related to the technologies described in their stories. In some cases, their stories describe tools and processes that are currently under development. In other cases, the technologies and affordances are novel. The goal is to provide the reader with estimates of which ideas are inevitable, possible, or unlikely. These second sections may be especially useful for graduate students, academics, and scientists who are looking for

summaries of current research and citations for refereed publications. The references are listed at the end of each chapter.

The final section in each chapter provides prescriptions for managers, scientists, and engineers. We seek to answer the question, "How do we get there from here?" The authors provide their expert opinions regarding the relative importance of various research and development projects. To convert the visions described in the stories to reality will require expenditure of resources. Some of that effort will come from commercial enterprises, some may come from university laboratories, and some may require government support. What is the best way to allocate our scarce resources for research and development? We expect that these sections will be especially useful for managers of research and development. Finally, the story format itself was chosen to promote certain cognitive benefits. In the first chapter, we provide a scholarly essay outlining the reasons why stories can facilitate encoding and recall of detailed information.

As noted earlier, the popular use of virtual worlds has been for gaming and social interaction. Our intent is to look beyond play and entertainment to present a view of the future where synthetic worlds are being used for significant work. As such, we have selected five work thrust areas for grouping the chapters. Readers interested in specific themes may find it useful to read only a subset of the chapters.

The first section, Forecasting, explores how synthetic worlds could be used to construct and visualize possible future situations in several different business domains. Messinger and Ge begin this section by looking at how market research techniques will be impacted by making online communications more personal and vivid. This first chapter in the Forecasting group is somewhat longer than the others because we asked the authors to provide an especially detailed review of the literature. Scholars and academics may find this chapter useful as a general reference. Brath et al. explore the collaboration aspect of synthetic worlds for analyzing and making projections in the financial industry. Osias et al. presents a mirror world situation for monitoring the security of data centers. Next, we turn our attention to the physical world where Bodnar et al. address biological research and Tignor explores simulating environmental change. This section concludes with a look at synthetic worlds for working the interrelated problems of intelligence analysis, smart collection, and mission operations by Trinko et al.

The second section, Forensic Analysis, looks at reconstructing and understanding past events. Miller begins this section by presenting an environment that is rich in information providing the user with total situational awareness. What may be especially interesting to our readers is that Miller's Miramar synthetic world actually exists, albeit not in the highly developed form described in his story. Waltz follows this up with a description of recreating foreign situations to aid in national security policymaking. Rice concludes this section by combining synthetic world technology with augmented reality for crime scene analysis.

The next section, Cognitive Amplifiers, looks at how synthetic worlds can be used to enhance, augment, or otherwise expand the participant's personal skills. Kisiel leads off this section by describing how synthetic worlds can be used to gain experience in faster than real-time. Bringsjord explores using synthetic worlds for visualizing storylines so

authors can directly interact with their characters during the writing process. Schultheis et al. show how synthetic world technologies can offer the opportunity for persons with disabilities to be accepted as part of the general public while minimizing the effect of the disabilities. Finally, Burley presents a case for using synthetic worlds to hone negotiating and consensus building techniques.

In the fourth section, we turn our attention to Training, an area that has already seen a lot of development over the past decade (e.g., serious games), but with a slight twist. Garretson and Denny describes a future where an airman enhances his skills using a combination of virtual and augmented reality. Goertzel and Cox et al. explore the training of avatars (supported by intelligent agent technology). McCuaig et al. conclude this section by exploring how synthetic worlds can help individuals learn and practice higher order thinking skills.

The final section, Infrastructure, takes a look behind the scenes at what it will take to ensure the security and usability of synthetic worlds. O'Connell et al. explore the future usability testing well before an actual product is built. In the final chapter of the book, Stanton reviews the history of multi-user worlds for clues to assist in identifying security issues in future virtual spaces.

One of the recurring themes that emerges in the visions for synthetic worlds is the extensive role for advanced automation, typically in the form of "agents". One of the more interesting aspects of agents, however, is that they mean different things to different people. It is worth a few moments here to consider how these different visions for agents are alike and different, how they relate to advanced automation, and ask: What should the role of agents be in synthetic worlds?

As you read through the book you will find some extremely ambitious roles for agents, they are often virtually indistinguishable from human entities working in a synthetic world. This vision of agents is not a new one, as it is a recurring theme in the science fiction literature. Unfortunately, there is still a huge leap of faith in bringing such a vision of agents into reality. On the other hand, there are other visions for agents that seem much more plausible. There are more limited concepts of agents where the agents serve as active context filters, where agents have an understanding of task context and selectively present changed information based on some level of understanding of likely relevance. Then there are agents as information gatherers; that is they are semi-autonomous constructs that look for information based upon information a user has previously used in performing a task or analysis. These agents are typically the logical products of workflow management tools.

Regardless of how we define agents, looking across the visions of synthetic worlds described in this volume, there are several issues that emerge in connection with them. The first is the relationship of the agents to the users they are working for or with. This is closely related to the work done in the late 1980's through the 1990's relating to philosophies of automation, and the appropriate degree of autonomy that should be granted to automation for various man-machine systems. Numerous attempts have been made to develop notional taxonomies for the application of automation in human-machine systems (e.g., Morrison and Gluckman 1993, Parasuraman et al. 2000, Hutchins et al. 1996). Parasuraman et al. (2000) proposed one of the first notional taxonomies for advanced automation.

In these articles, one of the recurring dimension of their taxonomies is the idea that Advanced Automation should refer to systems where the automation had equal status to that of the humans in the system, and had full autonomy to perform tasks with or without the knowledge of the human users of the system. Advanced automation gives the automation in a system the same status as the human users of that system. They share autonomy and authority.

In contrast, Decision Aiding is a concept where the automation is an active agent in facilitating the human processing of information by transforming it so as to highlight a predefined parameter of a complex system to elicit a required decision as determined by the system designer. In its simplest form, one could consider the advent of an automotive oil pressure "idiot light" as a decision aid designed to elicit a specific response from the operator as opposed to the pressure gauge that provides more subtle information. In the case of the warning light, the system designer has determined that the only thing the operator of the system needs to know is that there is a problem and the engine needs to be turned off.

Finally, Decision Support is a concept advocated in Hutchins et al. (1996), where the goal of the automation is to never take the decision-making away from the human decision maker. The goal of the automation is to get the decision-maker "in the ballpark", but never limit or force a specific decision. These differences in philosophy of automation are subtle but important. Hutchins et al. repeatedly found that effective decision support was much more accepted and useful to decision-makers in complex tactical weapons systems, even if the underlying algorithms were the same as those used in an automated or decision aiding system. In decision aiding, the automation has limited autonomy—there is usually an explicit contract between the human and the automation regarding roles and authorities. However, the automation can process data, use rules to infer new knowledge, and has the authority to take limited actions without explicitly consulting with the human users. Decision support may use many of the same underlying technologies as advanced automation, but the system is designed to never take decision-making away from the decision makers.

What does this suggest for agents? The view that agents will be synthetic, algorithmically based entities that wander around among humans in a synthetic world seems not just computationally difficult (extremely difficult), but also seems fraught with peril. It is unlikely that fully autonomous agents with equal status and authority are likely to be accepted in polite society—even a synthetic one. However, this same technology may prove critical to managing the astounding quantities of information implied by working in synthetic worlds. Agents derived from a user's activities, either explicitly created by the user or implicitly derived from workload, do seem plausible. Such agents would need to serve an advisory role where they track what has changed and perhaps seek out new information related to work a user has previously done in the synthetic world.

So, we suggest that as you read these different visions of the future, consider the implications of using automation in these various task domains. Does the use of automation create an unintended secondary task? Or, does the automation reduce the amount of low-level, repetitive work, thus freeing the human to perform more valuable

tasks? We hope you derive as much enjoyment from reading this book as we did from creating it.

C.A.P. Smith
Kenneth W. Kisiel
Jeffrey G. Morrison