The St. Thomas Commonsense Symposium

Marvin Minsky, MIT

Push Singh, MIT

May 13, 2002

Contents

Introduction

This report summarizes the conclusions of the St. Thomas Commonsense Symposium, held at the Ritz-Carlton, St. Thomas U. S. Virgin Islands, on April 14-16, 2002. The following people participated in the meeting:

- Larry Birnbaum
- Ken Forbus
- Ben Kuipers
- Douglas Lenat
- Henry Lieberman
- Henry Minsky
- Marvin Minsky
- Erik Mueller
- Srini Narayanan
- Ashwin Ram
- Doug Riecken
- Roger Schank
- Mary Shepard
- Push Singh
- Jeffrey Mark Siskind
- Aaron Sloman
- Oliver Steele
- Linda Stone
- Vernor Vinge
- Michael Witbrock

Thanks

This meeting was made possible by the J. Epstein Virgin Islands Foundation.

We would like to thank Cecile Dejongh for taking care of the local arrangements, and extend a very special thanks to Linda Stone for making this meeting really happen.

The Problem

The use of computers has not changed much over the past two or three decades.

Enthusiast: How can you possibly say such a thing? What about the Internet? What about the World-Wide-Web? Tens of Millions of people are using them!

Yes, there are many more people involved, but what they are doing is almost the same, such as typing out messages to their friends. In effect, these are just faster telegrams. True, only a thousand people or so could do that in the 70s. But let's not take quantity for quality.

Enthusiast: But some things we do now are different. We can buy things from many different stores. We can download music instantly. We can search the Web for documents that we could never access before.

Agreed, we can do more things and do them more quickly—but what about the quality? You can search for papers that have the right words—but those searches still do not know what you mean. The trouble is that we're still stuck with programs that don't understand what their users actually want. Our software systems have not got much smarter. What's happened is that more people use them, now.

And that's why those markets are flattening out. We've been only increasing the numbers of users, until it's become a substantial fraction of our total population. So now we're approaching the limit of growth. We can't keep increasing the numbers of users. So let's look at the real problem.

What is the problem?

As soon as the first computers appeared, their behavior became the subject of jokes. The slightest errors in what they were told would usually lead to useless results—or leave them entrapped in endless loops. From having no semblance of common sense, they would send out checks for absurd amounts, or wipe out their clients' bank accounts. And fifty years later they still behave in almost the same old inflexible ways! Why don't they learn from experience?

As time went on, these problems persisted, although there were some islands of progress. Now, computers can answer some difficult questions—but only in certain particular subjects. Today, robots make cars in factories, but none of them yet can make a bed, or clean your house or baby-sit. Programs now solve the most complex equations, and vanquish the best human players of chess—yet still cannot do what most children can do, such as understand a first-grade story.

Why could we get computers to work at the levels of specialized human "experts" before we could make them do everyday things that appear to us to be very much simpler? It was because each specialized program needs only to work with a small range of knowledge that applies to a small, tidy, circumscribed world. Commonsense thinking may seem less profound, but it works very well over much wider realms.

Instead of blaming them for their faults, we must make our computers more resourceful. We must furnish them with some common sense, so that when we tell them what we want, they'll be able to understand what we mean.

Our vision.

Our aim is to make an interface that understands our most common needs, and then helps us to achieve those things. Only then will computers be able to know what we mean by our requests. To be sure, some programs may seem to do this—but in the past, that's been done by tricks, which stop working whenever we change the context. However, we think that we've made enough progress toward theories of how to do such things that it now makes sense to start such a project. So imagine a future where computers arrive equipped with enough procedures and knowledge that they can engage in conversations with people who have no specialized training. This interface will have, from the start, a pleasant personality. But later, because that system can learn, each installation will go its own way, adapting to those who are using it. In other words, they'll extend themselves to know more about their environments, just as we people learn to improve the abilities of our bodies and minds.

What forms will these new systems have? In our vision, they'll be embodied in something like a single 'chip' that is connected to a computer's bus when the computer is shipped from the factory. That way, it will be like a network node that can see whatever is going on-and can intervene in various ways by sending appropriate signals. Eventually, it seems to us, it should take control of everything else, and replace the inflexible programs called 'operating systems' by 'managers' that understand more of what's happening.

A Project.

Today Artificial Intelligence is a chaotic array of attempts to exploit the advantages of knowledge in the forms of (for example) Neural Networks, Logic, Case-Based Reasoning, Genetic Programming, Statistical Linguistics, etc. However, our goal is not to promote any particular such way to represent or use knowledge, but to develop an architecture that knows how to exploit the virtues of each kind of representation. We hope to do this by using “higher level” knowledge about where each such scheme has advantages—and how to avoid their particular bugs. To make machines more versatile, they will need more knowledge about themselves, that is, on higher, more self-reflective levels.

We envision a future with computers that have enough commonsense knowledge and processes to solve everyday problems, understand typical texts, and converse in natural language about all sorts of ordinary subjects. To accomplish 'change of phase', we must cross a threshold that will come when we have systems that know enough about themselves to be able to improve themselves.

We held this symposium to discuss an architecture for such a Project—to build a machine with human level intelligence.

Executive Summary

We achieved consensus on these three points:

1. *The participants agreed that to achieve human-level intelligence we need an architecture that can integrate many different methods for representing, applying, and learning commonsense knowledge, because no single technique is powerful enough by itself.* Marvin Minsky offered the Six Level Model from his forthcoming book *The Emotion Machine* as an initial proposal for such an architecture. The key feature of this architecture is that it supports many "ways of thinking", each of which can be seen as a more specific architecture suited to solving problems of some specific type. The proposal was well received by the participants, and sparked a more general discussion of the three main aspects of the commonsense problem: architecture, knowledge, and reasoning. However, we decided that there would not be time at this meeting to discuss in detail the many technical questions that must be faced, and we postponed this discussion for a later meeting.

2. *The participants agreed on a problem domain in which we could situate and combine our efforts to build the architecture.* We proposed a realistic virtual environment containing several simulated "people" that could manipulate simple kinds of objects like blocks, balls, and cylinders. We chose this domain because solving problems within it would require combining reasoning in multiple mental realms: social reasoning, physical reasoning, self-reflective reasoning, communicating with language, and the other core skills that underlie our broader common sense abilities. The simulated people would solve problems like building physical structures, teaching each other things through natural language conversation, and finding new ways to learn from experience. We debated for hours the virtues and limitations of this problem domain, but we achieved a significant consensus that in order to solve harder or more practical problems requiring common sense, we first needed to solve the more restricted class of problems that show up in the simpler world we proposed.

3. *The participants agreed on a long-term application that would have great impact, that depended on a solution to the commonsense problem, and that we could build incrementally as we made progress on the commonsense problem.* We considered many applications likely to be attainable in the short-term, and for which there are already limited prototypes—such as smart search engines and commonsense-enabled personal information managers. We argued fiercely over what sort of commonsense-based application would have the most impact on the world, one that could not be fully realized without a solution to the commonsense problem. We eventually settled on the concept of a *personalized teaching machine* that would come to understand you so well that it could adapt to your particular circumstances, difficulties, and needs. We drew our inspiration from *The Young Ladies Illustrated Primer*, an AI-based children's book envisioned by science fiction author Neal Stephenson in his novel *The Diamond Age*.

The following diagram summarizes the consensus:

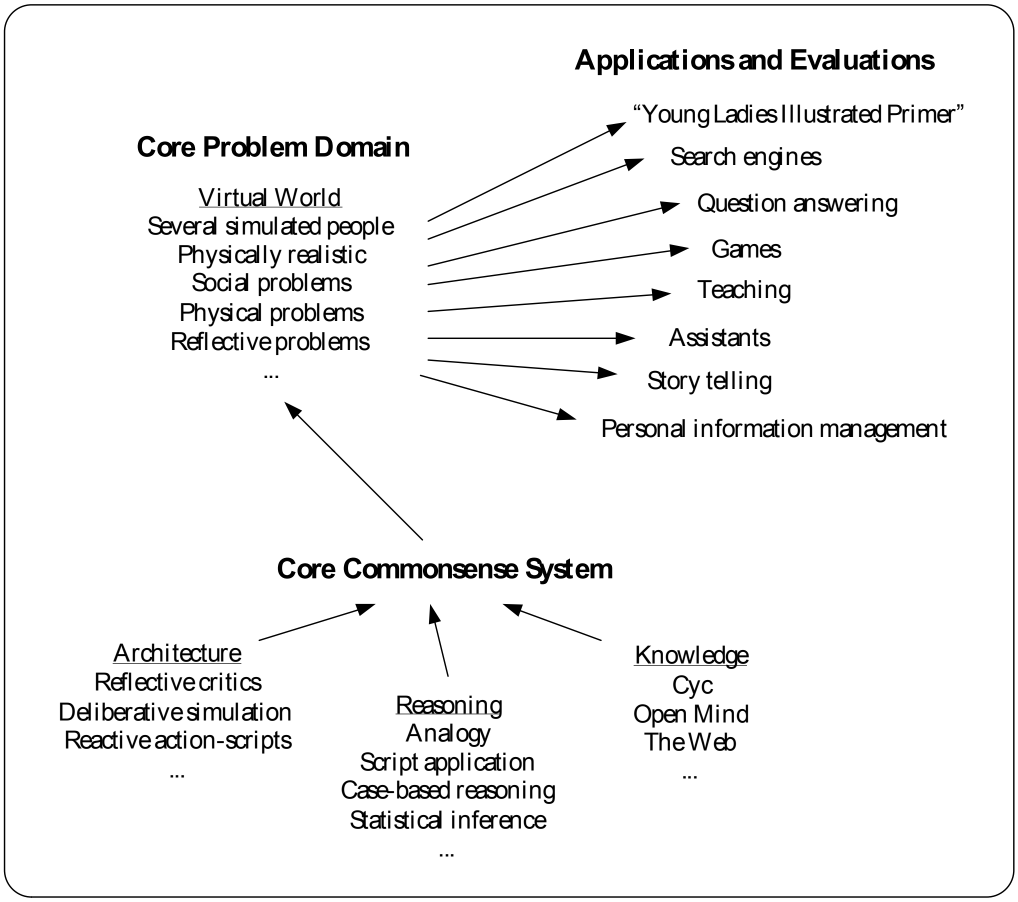

This diagram is based on one produced by Ben Kuipers at the conclusion of the meeting, which was crucial to achieving the consensus. It captures the dependencies between the three points of consensus: Practical applications depend on solving the more limited class of problems that show up in the core problem domain, and solving those depend on developing an architecture for commonsense thinking flexible enough to integrate a wide array of processes and representations.

The Proposed Architecture

We discussed the three main aspects of the common sense problem: architecture, reasoning, and knowledge. The participants agreed that the most popular methods currently being pursued by Artificial Intelligence community (for example, statistics, logic, neural networks, and genetic programming) could not by themselves achieve the goal of building a human level intelligence. They agreed that to achieve human-level intelligence we need an architecture that can integrate many different methods for representing, applying, and learning commonsense knowledge, as no single technique is powerful enough by itself.

Architecture: The Six Level Model

Marvin Minsky offered the Six Level Model from his forthcoming book *The Emotion Machine* as an initial proposal for such an architecture. This architecture is being developed jointly by himself and Aaron Sloman, and is based on several key ideas:

1. *Use several approaches, at once, to each problem*. When one method begins to fail the system can quickly switch to another. We represent each fragment of knowledge with several different representations. By always maintaining several viewpoints (in contrast to all traditional systems), our processes will rarely get stuck.

2. *Have many ways to recognize and respond to internal and external problems*. The architecture consists of layers of agents, where each layer is concerned with coordinating, managing, and responding to problems in the layers beneath it. Within each layer, there are 'critics' that detect types of problems in the layers beneath or in the outside world. These then turn on 'selectors' that invoke methods for resolving these problems.

3. *Support many different "ways of thinking"*. The most important high-level operation is mapping types of problems to large-scale "ways of thinking". Each way of thinking disposes the system to use certain types of knowledge, methods of reasoning, types of critics, and other kinds of resources to solve the problem at hand. This architecture is really a kind of meta-architecture, one that invokes *more specific architectures* in response to the kinds of problems the system encounters.

This proposal was well received by all the participants. Everyone agreed that to achieve human level intelligence we needed to develop more effective ways to combine multiple AI techniques. Ken Forbus suggested that we needed a kind of "component marketplace", and that we should instrument the components so that the reflective layers of the architecture had useful information to look at. He contrasted the Soar project as an effort to eliminate and unify components rather than accumulate and diversify them, as in our proposal.

Reasoning: Ways of Thinking

Ashwin Ram pointed out that despite the agreement over the architectural proposal it was still not clear what the components of a commonsense system would look like. He and Larry Birnbaum pointed out that we needed to think about what the *units of reasoning* would be. Larry Birnbaum stressed that we needed to expand on the many ways of thinking that constitute the methods and styles of reasoning that the architecture would support. The architecture should have many abilities, for example:

- Predicting what will happen next
- Explaining unexpected events
- Assigning credit for mistakes and successes
- Making analogies between concepts
- Generalizing from experience
- Recognizing the goals of other agents
- Inferring whether another agent needs help
- Describing problems to other agents

And so forth.

Knowledge: Mental Realms

Even the simplest commonsense problems require a broad range of commonsense knowledge. We proposed the idea of organizing this knowledge along *mental realms*. Each realm consists of the knowledge and methods of reasoning required to solve problems related to a particular aspect of the world. Consider the situation of a young child playing in a sandbox:

Even in this simple situation, the child may have concerns that span many mental realms:

- Physical: What if I threw the shovel?
- Bodily: Can I lift the shovel with only one hand?
- Social: Is my mother watching me?
- Psychological: I forgot where I buried my toy.
- Visual: Is that a rock or a clump of mud?
- Spatial: Will all that sand fit in the pail?
- Tactile: What would it feel like to squeeze sand?
- Reflective: Why didn't that work?

Each of these mental realms requires a vast array of commonsense knowledge. We estimate that an adult-level commonsense system will require 100 million units of commonsense knowledge. Such a resource is critical to the success of our project, because no matter how powerful the reflective and self-reflective levels of the system are, no learning approach is powerful enough to get to the level of even a young child starting from a blank slate, in a reasonable time frame.

Doug Lenat gave a presentation of the Cyc system, which is the furthest along at developing a useful such resource. Cycorp has made much progress on the problem of accumulating the vast array of world knowledge needed to give computers common sense. Their current direction is towards a distributed acquisition approach where eventually millions of teachers around the world will teach Cyc new commonsense knowledge. He described "clarification dialogues" as one method to allow non-logicians to participate in the complicated teaching and debugging processes involved in building up the Cyc knowledgebase.

Many of the participants agreed that Cyc would be useful, and some suggested we could even base our effort on top of it, but others were sharply critical. Jeffrey Siskind doubted that Cyc contained the spatial and perceptual knowledge needed to do important kinds of visual scene interpretation. Roger Schank argued that Cyc's axiomatic approach was unsuitable for the most important kinds of commonsense thinking—making the kinds of dynamic generalizations and analogies that a more case-based and narrative-oriented approach would support. Srini Narayanan worried that the Cyc project was not adequately based on what cognitive scientists have learned about how people make commonsense inferences. Oliver Steele concluded that while we disagreed about whether Cyc was 90% of the solution or only 10%, this was really an empirical question that we would answer during the course of the project.

For the purposes of the proposed project, we may only need a subset of the Cyc database. Marvin Minsky predicted that we would need to augment Cyc with many new kinds of procedural and heuristic knowledge:

- Problems that a knowledge item could help solve
- Ways of thinking that a knowledge item could participate in
- Known arguments for and against a knowledge item
- Means of adapting a knowledge item to new contexts

Push Singh stressed that knowledge about the world was not enough by itself. The problem remains that the programs we have for using knowledge are not flexible enough. Even a child can do far more with the little that it knows than our programs can today. This is why the proposed project is focused on developing the architectures and kinds of reflective and self-reflective knowledge needed for systems to use commonsense knowledge effectively.

The Proposed Domain

To build a commonsense system we need a concrete framework in which it can be developed, tested, and debugged. The participants agreed on a problem domain in which we could situate and combine our efforts to build the architecture.

A miniature approximation to the human world

It was difficult to determine a problem domain that set our sights high enough that we could not solve it by the application of shallow methods—ones that would not lead to full commonsense—but not so high that a project attempting to solve problems within it would likely fail. At the pre-meeting to this meeting held at IBM in March, we considered the target problem of getting a computer to understand children's stories. However, we concluded that by the time the child begins to read on its own it already has most of the common sense that it needs. Prior to learning how to read, children are speaking to one another, telling stories about their own situation and absorbing stories told by others.

We proposed that we should begin our work in a rich but simplified miniature of the real world, one that would be common to many of the researchers involved in the project. In this limited environment, many of the important problems would show up in a simplified way. We proposed a realistic virtual environment containing several simulated "people" that could manipulate simple kinds of objects like blocks, balls, and cylinders:

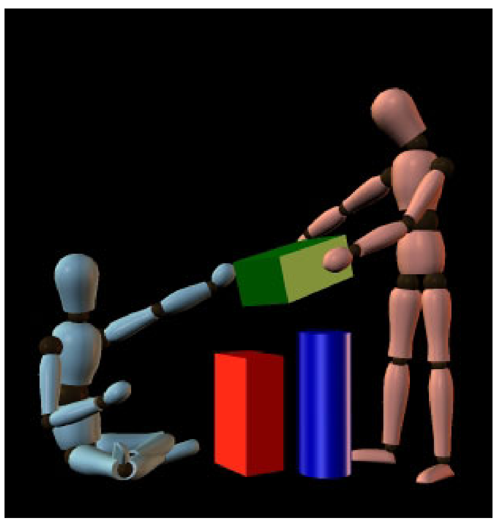

The simulated people would solve problems like building physical structures, teaching each other things through natural language conversation, and finding new ways to learn from experience. We chose this domain because solving problems within it would require combining reasoning in multiple mental realms: social reasoning, physical reasoning, self-reflective reasoning, communicating with language, and the other core skills that underlie our broader common sense abilities. Even within a limited world of objects, there is still a wide range of problems—in particular, many of the important realms show up in this simple world:

- Spatial: Reasoning about the locations of objects and path between them.
- Physical: Putting one object on top of another.
- Social: Dealing with the goals of the other people.
- Visual: The person will have a synthetic perceptual system.
- Reflective: The person can think about its own thoughts.
- Conversational: The people talk to each other about problems.
- Learning: The person can learn about the world from experience.

By first solving the key problems in this simplified world, we can build and debug the core of a commonsense system. Later, we can later expand it to solve other problems.

However, we should not assume that by restricting the world we will bypass the need for much commonsense world knowledge. We expect even this simple world to require hundreds of thousands or even millions of pieces of knowledge about space, physics, bodies, social relations, object appearances, and so forth.

Push Singh characterized this goal as that of building a Child Machine, because each of the people in the environment would have roughly the common sense of a child. The key idea is that young child has all the common sense it needs to keep on improving. Once you have built an equivalent to a five year old, you can send it to school, download a wider array of common sense knowledge into it, and so forth. Marvin Minsky distinguished the idea of programming a Child Machine with the commonsense of a five year old from building a Baby Machine that would learn everything on its own. It is a very popular idea to try to build an intelligent robot or agent that starts as a blank slate, and learns everything from physical experiences, reading the web, or conversing with people. His view is that we do not yet have enough ideas about how to represent, organize, and use much of commonsense knowledge, let alone build a machine that could learn all of that automatically on its own.

Objections to the micro-world approach

Many of the participants were enthusiastic about this proposal. Ken Forbus pointed out that the virtual reality and videogame communities would soon produce programmable virtual worlds that would easily meet our needs. Ben Kuipers appreciated that there would be challenging robotic, physical and spatial problems within this world.

Still, the participants had a heated debate about the adequacy of this proposed problem domain. The most common criticism was that this world is too limited in its scope—it does not contain enough of a variety of objects or richness of behavior for there to be enough interesting problems. Doug Lenat both complained about and suggested a solution to this problem, which was to embed the people within something like a typical house or office, as in the popular computer game The Sims. Doug Riecken suggested that we could build the 'core classes' within the more limited virtual world, and later add instances and specializations of those core classes to deal with a wider range of objects and phenomena.

One response to this criticism was that in order to focus on architecture we needed to simplify the problem domain. By limiting the world, the problems the system will encounter will be due not to a lack of world knowledge, but because there is some architectural problem where the system is not properly using what it knows.

Others argued that we should begin immediately within a practical application like a search engine, photo retrieval system, or conversational agent, and extend its commonsense abilities over time. Ben Kuipers pointed out that history has taught us that if you pick an application to start with, someone will come up with some shortcut that gives tolerable performance, without turning to general intelligence, and that does not lead to smarter systems—as is the case with present-day search engines. Marvin Minsky stressed that even if you had the goal of building a conversational system, most words are *about* things. To understand even the simplest sentence like “Jack gave Mary the book” you must be able to understand the physical, social, psychological, and other types of consequences of the action. By starting in the virtual world, you would be forced to address those problems.

We debated for hours the virtues and limitations of this problem domain, but we achieved a significant consensus that in order to solve harder or more practical problems requiring common sense, we first needed to solve the more restricted class of problems that show up in the simpler world we proposed. This world is challenging but tractable enough that we have a chance at getting a commonsense system working within it in a reasonable time frame. Once we get the core of the architecture functioning in this limited domain—in order to make sure that it can reason with a few realms and a small number of ways of thinking—we can go ahead and extend it to deal with a broader range of problems using a much broader array of commonsense knowledge.

Evaluating our progress

There was some discussion about how we could tell that we were making progress within this framework.

- Ben Kuipers suggested that there were lessons to be learned from the cognitive developmental of human children. He proposed that we could set our milestones relative to the abilities of real children, by determining some set of abilities that a one year old has, a two year old, and so forth. Michael Witbrock pointed out that it would be wrong to think about this system as a model of a real child, and we should instead regard it as a kind of AI approximation to *something like* a young child.

- Marvin Minsky suggested that it should be obvious when we are making progress, because we cannot solve problems that require common sense by simpler methods. If we solve one of these problems at all, it is because we have made real progress.

- Push Singh suggested that dividing common sense abilities along mental realms has the additional advantage of dividing up the evaluation task. Each realm suggests a set of questions or scenarios that the system should be able to deal with. Developing and defining realms will require that we pin down some set of problem types that occur within those realms, and that will clarify the evaluation task. Also, because there is a reflective layer, when the system fails on a problem the system itself should be able to notice. In other words, once the reflective layer is good enough, the system will become self-evaluating—it will be able to make its own judgments about its proficiency on various problem types.

Spin-Offs and Important Applications

How will giving computers common sense change the world? We considered a variety of important applications of this technology. The participants agreed on a long-term application that would have great impact, that depended on a solution to the commonsense problem, and that we could build incrementally as we made progress on the commonsense problem.

A killer application

Roger Schank suggested finding an application that everyone in the world would immediately recognize as important and useful. Marvin Minsky emphasized that it should be one that could not be built without an underlying system capable of human level commonsense reasoning. This application should be something very new and different, and not just some mild augmentation of an existing application.

After much argument, we decided what this application should be. The inspiration was drawn from science fiction author Neil Stephenson's *The Diamond Age*. In that novel he envisions a book—*The Young Ladies Illustrated Primer*—that, when given to a young girl, would immediately bond with her and come to understand her so well as to become a powerful personal tutor and mentor.

Our idea then is to build a *personalized teaching machine* that would adapt itself to someone's particular circumstances, difficulties, and needs. The system would carry out a conversation with you, to help you understand a problem or achieve some goal. You could pose it arbitrary problems, such as how to choose a house or car, how to learn to play a game or get better at some subject, how to decide whether to go to the doctor, and so forth. It would help you by telling you what to read, stepping you through solutions, and teaching you about the subject in other ways it found to be effective for you.

Textbooks will become obsolete, and replaced by systems that know how to explain ideas to *you* in particular. They will know your background, your skills, your learning styles, and tailor their content suitably. They will be able to construct example to pinpoint your particular strengths—and misunderstandings.

This kind of application could form the basis for a completely new way to interact with computers, one that bypasses the complexities and limitations of current operating systems. It would use commonsense in many different ways:

- It would understand human goals so that it could avoid the silliest mistakes.
- It would understand human reasoning so that it could present you with the right level of detail and avoid saying things that you probably inferred.
- It would converse in natural language so that you could easily talk to it about complex matters without having to learn a special language or complex interface.

This science fiction vision of a powerful tutor and problem-solving assistant was what rung most true among the participants. We felt that this goal could not be fully realized using anything less than a system with common sense, but that there could be an incremental path to this goal through more limited versions that would help you with more restricted classes of problems.

Other applications

We decided that we needed a small number of attractive applications, because if we picked just one then we might be tempted to invent some shallow method to achieve it. By picking several applications, it would be hard to find a single trick or technology that could make them all work. So in addition to the consensus we achieved around the teaching machine proposal, we considered a number of other possibilities. Marvin suggested these:

- A computer that can have real conversations. Our goal is build the first Natural Language system that can understand some of what it reads or hears. Speech recognition, grammar analysis, and text-to-speech are already almost good enough-but still make far too many mistakes to be useful to nonspecialists. This is because those programs still lack enough commonsense knowledge to understand what their user's words mean—or what those people want to know, or what goals they wish to accomplish. Our system will aim toward reasoning out what its user is doing—and why.

- Domestic robots. For decades we have all been hyped about 'robots for domestic use'—but all we've seen are clumsy things that purport to clean carpets and mow our lawns. But every such venture went down the drain—though mechanically, its robot was sound—for who wants a servant without any brain? Ours will keep track of, as well as maintain, everything that has entered your house—and will usually guess when you want it again. Ours will be able to understand scenes—and not only know what each object is, but also what those objects are *for*. These robots will be far more dexterous than today's because they will be solving meaningful problems, instead of obeying inflexible rules.

- Worker robots. It has been widely recognized by economists that we're approaching the retirement of the "baby boom" and that this, combined with increased longevity, will put great pressure on the decreasing proportion of productive workers. The answer to that seems quite obvious: more work should be done by intelligent robots! This can't be done now, simply because our machines don't yet have enough common sense—but our future project could remedy that. If we're all to enjoy those longer lives, this project becomes a necessity.

- Search engines with common sense. Commonsense-based knowledge retrieval will let professionals in every field put their data on the net so that you can discover new relationships. Everyone will benefit.

Other ideas included commonsense-enabled scheduling programs, games and toys, personal information managers, storytelling systems, and conversational agents. Henry Lieberman gave an impressive demonstration of ARIA, one of the few existing commonsense-enabled applications, a dynamic photo retrieval and annotation system that bases its retrievals on commonsense inferences from the annotations of the photos. We also discussed the possibility of programming systems by example—wouldn't it be great if we could show our robots and programs what we wanted them to do, rather than painstakingly encoding step-by-step procedures?

Applications depend on core common sense

We agreed that if we wanted to build a powerful teaching and helping machine, we would first need to give it knowledge about many more basic things. One example is that it would need to understand how to have a conversation. Ordinary conversation requires the ability to reason at least at a simple level about space, time, beliefs, plans, stories, mistakes, successes, relationships, and so forth. In other words, to achieve our longer-term goal of building powerful and important applications we must first address the problems that show up in the target problem domain of the virtual world. Ben Kuipers pointed out that *the hard part is getting the three-year-old working*. All attempts to find a shortcut around this have failed. Once we can do common sense reasoning at the level of a young child, building all of these other applications will become possible.

Towards a Collaborative Project

What comes next? There was some discussion of how we might organize a larger project to achieve the vision presented at this meeting.

Next organizational steps

We envision the project to take the form of a distributed AI lab. Erik Mueller and Push Singh suggested that we might build a lab along the following lines:

- Everyone would share the global vision of the virtual world problem domain, but we would not impose on the participants any specific method of representation or inference. We might ask that they agree to a set of general protocols that support the interoperation of a diverse array of methods.

- The lab would be modeled after the early MIT Artificial Intelligence Lab. Researchers would work on specific small projects to try out many ideas quickly. There would be a memo system, so that researchers could take some small chunk and develop it, and feed it back as a memo.

- There would be a core team that develops the meta-managerial layer. It would contain a knowledgebase "about" the different kinds of components that will exist in the system.

- Our goal would not be to build a 100 million item common sense knowledgebase, as Cycorp is trying to do, but we would develop tools to teach the system about the wide range of things it might need to know to deal with the simpler world that we have chosen as our target domain.

- As the lab would be distributed, we would arrange a series of meetings. There would be exploratory periods where people work separately, followed by convergence periods where we would try to integrate that work into a single system.

Next technical steps

Aaron Sloman listed some of the many technical steps that we will have take in this project:

- Develop protocols and representations for coordinating many different kinds of methods and representations.

- Determine the kinds of knowledge needed at the reflective level.

- Build the appropriate shared problem environment for the researchers.

- Identify the important mental realms, and develop representations and methods of reasoning for those realms.

- Build a database of the common sense knowledge needed within the target microworld.

- Support the larger Cyc effort to build a database containing the millions of items of knowledge one needs to model the everyday world. Eventually, have our system read this larger commonsense database as a fast way to get it up to speed on the rest of reality.

This list only scratches the surface.

Towards the future

There was a genuine sense of excitement at this meeting. Some participants went so far as to compare it to the historic Dartmouth conference where Artificial Intelligence was born. It is rare and special opportunity to plan an ambitious project of this nature—to focus the field of Artificial Intelligence once again on the grand goal of building a human level intelligence.