

The Open Mind Common Sense Project

by Push Singh, MIT Media Lab

January 1, 2002
(references added July 3, 2003)

Originally published at KurzweilAI.net

Citation: Singh, Push. (2002). *The Open Mind Common Sense Project*. Available:
<http://www.kurzweilai.net/meme/frame.html?main=/articles/art0371.html>

Introduction

Why is it that our computers have no grasp of ordinary life? Wouldn't it be great if your search engine knew enough about life so that it could conclude that when you typed in "*a gift for my brother*", it knew that because he had just moved into his first apartment that he could probably use some new furniture? Or if your cell phone knew enough about emergencies that, even though you had silenced it in the movie theatre, it could know to ring if your mother were to call from the hospital? Or if your personal digital assistant knew enough about people that it could know to cancel a hiking trip with a friend who had just broken a leg?

Despite years of research into artificial intelligence, building machines that can think about ordinary things the way any average person can is still out of reach. To be sure, computers can now do many remarkable things. Recent increases in computer power have allowed many years of research into artificial intelligence to bear fruit. We now have programs that can play chess at the level of the very best players, that can solve the logistics problems of coordinating the most complex military deployments, and that can help us engineer the most intricate computer chips to the most complex airplane engines. There is no doubt that artificial intelligence has come into its own as a successful field with many practical applications.

Yet how is it that we can write software that can do such complex things as design airplane engines, but still we cannot build machines that can look at a typical photograph and describe what is in it, or that can read the simplest children's story and answer questions about it? We have been able to write programs that exceed the capabilities of experts, yet we have not been able to write programs that match the level of a three year old child at recognizing objects, understanding sentences, or drawing the simplest conclusions about ordinary life. Why is it that we can't seem to make computers that can think about the world as any person can?

The problem seems not to be that computer scientists lack ideas about how to write programs that can reason and solve problems. The field of artificial intelligence is a veritable gold mine of techniques. There are programs that can successfully diagnose medical symptoms, analyze and repair problems in space craft, plan the best route for your rental car to your desired destination, and transcribe speech into text. Each of these technologies employs a different type of reasoning method, and there are many more methods out there. We certainly do need new ideas about the mechanisms that underlie intelligent behavior. But it seems more likely that we have more than enough ideas about the ingredients of intelligence, and that the real problem lies elsewhere.

The real problem is that computers do not know anything about us! Our machines lack *common sense* – all that ordinary knowledge that people in our society share, things we understand so well we hardly even realize we know them. Computers do not know what we look like, how we typically behave, or what we are capable of. They do not know anything about the patterns of people's lives, the places we spend our time, or the kinds of relationships we have with each other. They know nothing of our hopes and fears, the things we like and the things we loathe, or the feelings and emotions that motivate and underlie everything we do. By giving computers ways to represent and reason about knowledge about people, they can be made to become helpful and eager participants in the human world.

Computers lack common sense

Why has it been so hard to give computers this common sense, the ability to reason about our ordinary goals and concerns? Why are we able to program machines like Deep Blue, that can beat the greatest chess players that have ever lived, yet we are not able to give computers the common sense that gives them the thinking abilities of a five year old child? Listen to Marvin Minsky in his book *The Society of Mind* [1]:

Common Sense: The mental skills that most people share. Commonsense thinking is actually more complex than many of the intellectual accomplishments that attract more attention and respect, because the mental skills we call "expertise" often engage large amounts of knowledge but usually employ only a few types of representations. In contrast, common sense involves many kinds of representations and thus requires a larger range of different skills.

In other words, giving computers common sense is the opposite of the problem artificial intelligence researchers have traditionally faced. Rather than it being a problem of how to give computers a great deal of knowledge about how to think about some particular area such as chess playing or circuit design, to give computers common sense we must program them with knowledge about many different areas: physical knowledge of how objects behave, social knowledge of how people interact, sensory knowledge of how things look and taste, psychological knowledge about the way people think, and more. Furthermore, each of these different facets of life requires its own specialized methods of reasoning about them. Giving computers common sense is not about figuring how to make some particular method of reasoning work over a particular type of knowledge, it is about how to make systems that are abundant with many types of knowledge and many ways of thinking about different things. There are two parts to this problem:

The first problem is how to give computers *commonsense knowledge*, the millions of ordinary pieces of knowledge that every person learns by adulthood. Many of these are so obvious we take them for granted:

- *Every person is younger than the person's mother*
- *A butcher is unlikely to be a vegetarian*
- *People do not like being repeatedly interrupted*
- *If you drop paper into a flame then the paper will burn*
- *You start getting hungry again a few hours after eating a meal*
- *People pay taxi drivers to drive them places*
- *People go to parties to meet new people*
- *People generally sleep at night*

We know many millions of such commonsense things, the knowledge about the world that most people in our culture share, and each of these pieces of knowledge is probably represented in the brain in several different ways.

The second problem is how to give computers the capacity for *commonsense reasoning*, ways to use this commonsense knowledge to solve the various problems we encounter every day, large and small. Having a database of millions of facts is not enough by itself (although we cannot do without it!) In addition to this sort of factual knowledge about the world, we need to give computers many different ways of using that knowledge to think: different methods of reasoning, planning, explaining, predicting, and all the other kinds mental skills we use for getting along. Some examples of the kinds of things we do when we think include these:

- *If you have a problem, think of a past situation where you solved a similar problem*
- *If you take an action, anticipate what might happen next*
- *If you fail at something, imagine how you might have done things differently*
- *If you observe an event, try to infer what prior event might have caused it*
- *If you see an object, wonder if anyone owns it*
- *If someone does something, ask yourself what the person's purpose was in doing that*

These kinds of things are themselves a type of knowledge, not about the outside world but about the world inside the mind. This knowledge is about ways to use other kinds of knowledge, in order to reason about the world in a way that helps us effectively navigate the problems of daily life.

Common sense is really the all those things that lets a person *think like a person*, which people are able to do because our brains are vast libraries of commonsense knowledge and methods for organizing, acquiring, and using such knowledge.

Common sense is a problem of great scale and diversity

How big is common sense? The scale of the problem has been terribly discouraging. Those courageous few who have tried have found that you need a tremendous amount of knowledge of a very diverse variety to understand even the simplest children's story. Many researchers regard the problem of giving computers common sense as simply too *big* to think about.

There have been several attempts to estimate how much commonsense knowledge people have. These attempts range from experiments that demonstrate that people can only acquire new long-term memories at the rate of a few bits per second, to counting the number of words we know and estimating how much we know about each word, to estimating the brain's storage capacity by counting neurons and guessing at how many bits each neuron can store. Several of these estimates produce approximately the same number, suggesting there are on the order of *hundreds of millions* of pieces of commonsense knowledge.

This is a huge number. What can we do, given that there is so much common sense out there that we need to give our machines? Is there anything we can do to accelerate the process of teaching computers all this knowledge? Building a robot that lives in the world and learns like a real child might eventually be made to work, and it is in fact a popular idea today to try to build such a "baby machine". But to date no one has figured out how to build learning programs that can learn the broad range of things that a child can, and that can keep on learning without getting stuck.

The problem seems to be one that was noted many years ago by artificial intelligence pioneer John McCarthy: *In order for a program to be capable of learning something, it must first be capable of being told it.* 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. There are many romantics who believe that this kind of prior understanding is not necessary to build a baby machine, and intelligence can be made to emerge from some generic learning device. But in our view, unless we can acquire some experience in manually engineering systems with common sense, we will not be able to build learning machines that can automatically learn common sense. So we must bite the bullet, and make a manual attempt to build a system with common sense. This is a problem as large as any other that has been faced in computer science to date.

There has been one major attempt to manually build a database of human common sense knowledge, the Cyc project led by researcher Doug Lenat [2]. Many people regard Lenat as foolhardy for taking on such a grand challenge, but to us he is doubly a hero – first, for having the courage to attempt this project at all, and second, for standing up to his critics over the years and persisting in the face of all that criticism. Doug Lenat and his team of knowledge engineers have worked for nearly two decades, at the cost of many tens of millions of dollars, painstakingly building up what is now a database of 1.5 million pieces of commonsense knowledge.

But as much as we admire the Cyc project, 1.5 million pieces of knowledge is still terribly far away from several hundred million. The Cyc project faces a challenge no single team could expect to succeed at. And indeed, they themselves admit they are still one or two orders of magnitude away from what is needed. We believe that a database this large simply cannot be engineered by any one group.

Turning to the general public

Is there a solution to this problem of scale? We believe that there is, based on one critical observation:

Every ordinary person has common sense of the kind we want to give our machines.

The advent of the web has made it possible for the first time for thousands or even millions of people to collaborate to construct systems that no single individual or team could build. These types of projects are known as *distributed human projects*. An early and very successful example was the Open Directory Project¹, a Yahoo-like directory of several million web sites that was built by tens of thousands of topic editors distributed across the web. The very difficult problem of organizing and categorizing the web was effectively solved by distributing the work across thousands of volunteers across the Internet.

Could the problem of giving computers common sense be cast in the same way, where thousands of people with no special training in computer science or artificial intelligence could participate in building the bulk of the database? Over a hundred million people have access to the web. If each of them contributed just one piece of knowledge, the problem could be solved!

Of course, it is easier said than done. But if this is indeed possible, then the reward would be tremendous – we would finally be able to solve this problem of giving computers common sense, making them far more understanding and useful than they are today. If we can find good ways to

¹ <http://www.dmoz.org>

acquire common sense from people by asking them questions, presenting them with lines of reasoning to confirm or repair, asking them to tell stories about everyday life, and so on, we may be able to accumulate the many of the types of knowledge needed to give our machines the capacity for commonsense reasoning. Even if we could only partly solve the problem this way, we would be a long way towards our goal.

Supporting thousands of on-line volunteers with Open Mind Common Sense

Could you build a web site that makes it fun and easy for everyone to collaborate and help teach computers about the everyday world? With this goal in mind, we built a web site called Open Mind Common Sense, which is located at <http://www.openmind.org/commonsense>. This web site has been running since the fall of 2000, and so far we have gathered many hundreds of thousands of pieces of knowledge from over eight thousand people. This is still only a quarter of what is in Cyc, but we acquired this knowledge in a fraction of the time and at a tiny fraction of the cost.

Open Mind Common Sense is a kind of second-generation common sense database. We learned much from Cyc, but as we were inspired by Cyc, we wanted to do things in a different way. After all, why are there not four or five Cyc sized projects in the world, each taking a different approach to attacking the critical problem of giving computers more common sense? How could we approach the common sense problem differently? In the following sections I will discuss in more detail our general approach in the Open Mind Common Sense project.

We need many types of knowledge and ways to represent that knowledge

One of the first questions we faced in this project was how to *represent* the common sense knowledge acquired from people. There have been many disagreements in artificial intelligence over this question. One popular view is that we should try to write down knowledge in the form of *rules*, things such as “if you are in a gravitational field, and you are holding an object, and you let go of the object, then that object will fall down.” Of course, every rule has its exceptions, but this approach has proven useful in building many practical reasoning systems.

Others have argued that more “concrete” representations, such as stories or descriptions of situations, are the right way to represent commonsense knowledge. In that approach, in order to solve a problem, you would look up a story or situation where you solved a similar problem, and do what you did then. For example, you might have the story “Bob was hungry. Bob ate a sandwich. Bob was less hungry.” Then if you were hungry and wanted to be come less hungry, you would be reminded of that story and get the idea that it might be useful to make a sandwich.

Still others have argued that much of common sense is less about “thinking” than it is simply routine behaviors that operate using less explicit types of knowledge, for example the kind of knowledge we use to walk around. Much of what we know seems to be implicit in this way, yet it is as much a part of common sense as the more explicit types of knowledge we can more easily articulate.

Our view is that when it comes to common sense thinking, *diversity* is the secret to success. It is not so much a matter of choosing between such representations as it is a matter of finding way for them to work together in one system.

Not only do we need many kinds of representations, we need many types of knowledge. To cope with the broad range of problems that can people solve, we need knowledge of a very diverse variety: spatial, physical, social, tactile, emotional, bodily, economic, and so on. The more types of knowledge available, the more likely it is you will have the appropriate type of knowledge to solve whatever commonsense problem you are faced with at the moment. For example, to understand the simple sentence “John saw the movie with Susan at the theatre”, you need to use all these types of knowledge, and many more:

- John was sitting near Susan while they watched the movie (*spatial*)
- John and Susan spent some money on tickets (*economic*)
- John and Susan know each other and are probably friends (*social*)
- John and his chair were in contact during the movie (*tactile*)
- Susan’s body was in a sitting posture during the movie (*bodily*)

The types of knowledge gathered by Open Mind Common Sense

A large part of building a commonsense knowledgebase is identifying the types of knowledge that are needed to solve different types of common sense problems. For each identifiable type of common sense knowledge, it might be possible to construct a web page specially designed to acquire that type of knowledge. And furthermore it might be possible to acquire that knowledge so that people untrained in artificial intelligence could supply it.

To this end, we constructed a variety of activities for gathering knowledge. Each activity tries to make it simple for the user to teach a certain kind of knowledge. At present we have about 25 activities. Some of the kinds of knowledge we collect include:

- What sorts of things are there? (*Cats are mammals.*)
- What are the likely effects of an action? (*Eating a sandwich makes you less hungry.*)
- What sorts of things do people want and not want? (*People want to be warm.*)
- Where are things typically found? (*You often find a hairdryer in a bathroom.*)
- What are objects for? (*Chairs are for sitting on.*)
- How do you solve problems? (*If you want to be less hungry, eat a sandwich.*)

One way to acquire this knowledge is by giving people explicit templates to fill in, such as:

- A hammer is for _____
- People want _____
- The effect of eating a sandwich is _____
- Somewhere you find a bed is _____
- A typical activity is _____

Another way is to give people simple stories and ask them to type in knowledge in a freeform way:

User is prompted with story:

Bob had a cold.
Bob went to the doctor.

User enters many kinds of knowledge in response:

Bob was feeling sick
Bob wanted to feel better
The doctor made Bob feel better
People with colds sneeze
The doctor wore a stethoscope around his neck
A stethoscope is a piece of medical equipment
The doctor might have worn a white coat
A doctor is a highly trained professional
You can help a sick person with medicine
A sneezing person is probably sick

We also collect knowledge that is larger than can be expressed in single sentences. For example:

Concrete situations: We have our users supply descriptions of photos in plain English:

Example: A mother is holding her baby. The baby is smiling. They are looking into each other's eyes. The baby is happy. The mother is happy.

Concrete episodes. We have our users supply stories, either to illustrate some existing fact like "flashlight light up places" or in response to a given title like "going outside at night":

Example: It was too dark to see. I went and got my flashlight. I turned on my flashlight. I could see much better.

Visual events. To allow more sophisticated forms of spatial reasoning, we allow users to annotate movies of simple iconic spatial events. Our goal is eventually to learn translation rules that let us produce event descriptions from movies and vice versa. This is a first step towards reasoning with pictures as well as with sentences.

Example: The small red ball rolls past the big blue ball.

This list only scratches the surface.

A modern attempt to build a commonsense database can start with natural language

Aside from our targeting the general public, the other major difference between Cyc and Open Mind Common Sense is that we have asked our users to supply knowledge in plain English sentences. Cyc expresses all knowledge in CycL, a precise but difficult language in which it is possible to state facts clearly and unambiguously, which makes it easier for computers to manipulate and use. We considered using CycL or something like it in Open Mind, but decided the average web user would have little interest in learning a whole new language to participate in our project. Using English as the representation may require using a simplified version of the language, but may be the only way to allow a wide range of people to participate in the project.

One reason to believe English input is a good idea is that modern natural language processing techniques are now good enough to extract the correct syntactic structure from a large proportion of the sentences people have supplied to our system. And even if we cannot build a system that can understand every sentence users enter right now, over time we can develop better and more accurate ways to represent the knowledge expressed in each statement.

Also, as pointed out earlier, one thing that distinguishes common sense from other types of artificial intelligence is the range of types of knowledge that are needed. We wanted to gather knowledge of many different types. No simple and easily learned language could possibly represent all the types of knowledge that someone might express in an English sentence. Any invented language would need to be at least as complex as English to express the range of common sense knowledge we were seeking.

There are problems with using English. The main reason that English is regarded as too unwieldy a representation is that vagueness and ambiguity pervade English, and computer reasoning systems generally require knowledge to be expressed accurately and precisely. But in our view, ambiguity is unavoidable when trying to represent the commonsense world. The core of the problem is that different people simply do not agree on the boundaries of ideas. One person's concept of a "vacation" will differ from the next person's. So if you want to take a knowledge engineering approach where many people enter knowledge into a system – which is true of Cyc as well, though they use fewer people – then you need a way to compensate for the fact that people have different ideas in their heads. The precision of most artificial intelligence languages may in fact be a problem for any approach that attempts to acquire knowledge from people. If ambiguity cannot be avoided, then we must learn to cope with it. So why not build on natural language?

Using the knowledge we have gathered

How can we use the knowledge that we have gathered? We can build software to help us in everyday life by actually understanding everyday life. At the Media Lab, we are exploring several kinds of applications. For example, we have built a search engine application that uses commonsense knowledge to reason about the true goal behind the search query [3]. When the user types "my cat is sick" into the search engine, the system reasons, roughly, that

*People care about their pets
& People want their pets to be healthy
& My cat is my pet
& I want my cat to be healthy
& A veterinarian heals sick pets
& A veterinarian makes sick pets healthy
& I want to call a veterinarian
& A veterinarian is a local service
Therefore: Search for a veterinarian in the user's area*

In other words, the system reasons that what the user really wants is for their cat to be healthy, and therefore they should seek a nearby veterinarian. None of the popular web search engines presently engage in this kind of reasoning about people's search queries.

Another application being developed by researcher Barbara Barry is what you might call an intelligent video camera [4]. Before people go to shoot an event with a camera, they have a good idea of what might happen at the event. But most of the time, they become so engaged with what they are looking at through the lens that they forget the types of shots that later would later convey a good story about the event to another person. They shoot a lot of pictures or video and when they get back to the desktop there are holes in their story. Retrieving knowledge from Open Mind about events can give the filmmaker a list of shots to take in order to build a complete story. Open Mind knowledge can be used for shot suggestions, to make sure the maker has the

ingredients to construct a coherent story. For example, here are some things that a marathon runner does, which could be seen as a simple list of shots that a person with a camera could capture during a marathon to tell a story about it in pictures:

- *Things found at a marathon: Crowds, runners, finish line, and paramedics*
- *A marathon runner eats pasta the day before the race*
- *A marathon runner gets their number at the starting line*
- *A marathon runner runs from the starting point to the finish line*
- *A marathon runner gets tired during the race*
- *A marathon runner will be very happy if they win the race*
- *People hand cups of water to marathon runners during the race*
- *The effect of winning a marathon is that the audience cheers*

Another application, under development by Henry Lieberman and his students Kim Waters and Hugo Liu, is *Aria*, a photo retrieval system that retrieves relevant photos as you are typing an e-mail to someone [5]. For example, let's say you have gone to your friend Mary's wedding and have taken a bunch of pictures which you have partly annotated. The following day you are writing an e-mail to another friend about the event. You type "I finally got to meet Mary's sister Judy." Even if you have no photos labeled "Judy", you would like to retrieve the photo "Mary and her bridesmaids." This is an example of where common sense is useful. If you knew that "At a typical wedding, the bridesmaid is often the sister of the bride", then the system might realize that photo was relevant, and suggest inserting it into the E-mail. Open Mind knows a lot about typical events such as weddings. Here are some:

- *A bride wears a wedding gown*
- *Weddings usually have parties afterwards*
- *A wedding ring is usually worn on the left hand*
- *Weddings often happen at chapels*
- *Wedding photographers are expensive to hire*
- *Something that is found at a wedding are the bride and the groom*

Eventually, we would like for such applications themselves to drive the acquisition of further knowledge. When the system makes some strange inference, the user can correct it, because they are motivated by the desire to improve the application. Or when the system fails to make an inference, they can add knowledge so that next time it will make it.

Achieving robustness through a diversity of heuristic reasoning methods

How can we reason with the knowledge we have acquired? We have explored using both rule-based and analogy-based methods of reasoning. For example, we have found that a combination of facts expressed in English and rules that match those facts can give us a simple way of reasoning with the data we have collected. It is difficult to produce only correct conclusions with these kinds of simple techniques, but for the applications we have explored, web search and photo retrieval, we have found that we do not need the kinds of accuracy that is needed in, say, building a medical reasoning system.

Of course, people make mistakes. Human reasoning is both unsound (we produce wrong conclusions from correct assumptions) and incomplete (we do not draw a conclusion even when it is "obvious" from the evidence.) This point has been powerfully demonstrated by the child psychologist Jean Piaget, who has shown that children do not just lack knowledge about the world, but that they lack the ability to make proper inferences given what they know. Reasoning

logically seems to be a skill that some people learn as they get older, but even adults fail to draw the correct conclusions when asked questions like “If most Canadians have brown eyes, and most brown eyed people have good eyesight, then do most Canadians have good eyesight?” You might guess yes, but in fact this is not true. (Imagine a world where most people have brown eyes and good eyesight, except for that small group of brown eyed Canadians!)

Yet, soundness and completeness are two of the most sought after properties researchers seek in building their inference systems. Of course these properties are desirable much of the time. You would not want your medical expert system making the wrong diagnoses. These methods have their place in building systems that, in narrow areas, engage in more robust reasoning than any human could.

But the trouble is that soundness and completeness requires carefully engineering our assumptions and rules of inference to prevent problems and conflicts, and while this may be possible in narrower domains of knowledge, it is unclear that we can ever achieve such perfection in the large commonsense systems we hope to build. No one has been able to build a system that can achieve the robustness of human reasoning in the commonsense domain. We need to give up on romantic dreams of pristine properties of soundness and completeness, and instead find ways to use databases that contain knowledge that has errors and biases, and use inference methods that are unsound. To scale up, we need to develop ways of reasoning where errors can be cancelled out or corrected by other information.

One idea we are fond of is that there are always multiple ways to make any particular argument. No single argument is always completely reliable. However, we think it is possible to improve the robustness of reasoning by combining multiple types of arguments. We have been exploring the idea that you can combine multiple inferences to produce a single, more robust inference. Robustness would come from combining multiple types of arguments, much as a table can be made to support a large weight not by giving it a single very strong leg, but by giving it many less strong legs that together can support more than any one leg could. And to continue the analogy, some of those legs might even be broken (just as individual pieces of knowledge and reasoning might be inaccurate) and yet the table could still stand up.

We need new kinds of architectures for building systems with common sense

We believe that commonsense databases will soon be ubiquitous. One way or the other, through directly acquiring knowledge from people, reading the web, or making use of robots that can see and touch things, we will find ways to acquire common sense about the world. But databases of commonsense knowledge by themselves are not enough to build systems with human-level thinking abilities. Adding a few types of reasoning methods helps, but this is still not enough. We need to develop new kinds of *architectures* – large-scale ways to organize and use knowledge – to use commonsense knowledge. Perhaps this is best illustrated by this scenario Marvin Minsky describes in his book *The Emotion Machine* [6]:

Joan is part way across the street on the way to present her finished report. While thinking about what to say at the meeting, she hears a sound and turns her head –and sees a quickly oncoming car. Uncertain whether to cross or retreat, but uneasy about arriving late, she decides to sprint across the road. She later remembers her injured knee and reflects upon her impulsive decision. "If my knee had failed, I could have been killed. Then what would my friends have thought of me?"

Minsky suggests that many kinds of thoughts go through Joan’s mind during this event:

Reaction: She reacted rapidly to that sound.
Representation: She constructed descriptions of things and ideas.
Attention: She noticed certain things rather than others.
Decision: She selected among alternative options.
Meta-Decision: She selected some method for choosing those options.
Embodiment: She was partly aware of her body's condition.
Intention: She formulated some goals and plans.
Language: She heard words or dialogs in her mind.
Imagining: She envisioned some alternative possible futures.
Planning: She considered various action-plans.
Reasoning: She constructed various arguments.
Recollection: She constructed descriptions of past events.
Identity: She regarded herself as an entity.
Reflection: She thought about what has she recently done.
Moral Reflection: She reflected upon what she ought to have done.
Self-Reflection: She reflected on what she was recently thinking.
Self-Imaging: She engaged certain models that she's made of herself.
Social Reflection: She considered what others might think about her.
Self-Awareness: She recognized some of her mental conditions.

Like Joan, a commonsense reasoning system will need to operate on all these levels. No single type of inference strategy, such as the application of a rule or the making of an analogy, is by itself up to the task of commonsense reasoning. The problem is not so much that these methods are not useful – on the contrary, they will form the foundation of any system that attempts to do commonsense reasoning. The problem is that even the most ordinary inferences may involve some combination of several different types of reasoning. The solution is to find ways to combine many kinds of knowledge and reasoning strategies to make more complex inferences than any individual strategy could make.

The problem of building architectures for common sense is, we believe, technically one of the most interesting challenges faced in artificial intelligence. It is a problem of a fundamentally different type from other problems in artificial intelligence. It is not so much about how to invent a simple algorithm or method for solving some simple problem. It is about how to use many kinds of methods together to solve hard and often poorly defined problems. But this is still a poorly understood area – how do you manage and coordinate multiple types of thinking processes? The solution to this problem may not be to embed it within the architecture itself. No simple way of arranging the elements of a commonsense system could possibly secure that all inference within it is productive and sound. Instead, we need to find ways to build systems into the mind that are experts at reflecting on, coordinating, and managing other knowledge-based processes.

The future of commonsense systems

We believe that the first systems to which people will attribute interesting human-like intelligence will result from finding ways to use and apply databases of commonsense knowledge, and that the best way to acquire these databases is through large collaborative efforts. Open Mind Common Sense is really a first attempt at realizing the idea that we might distribute the problem of constructing a system with common sense. We are now working on a second-generation version of the web site. The new site will focus on acquiring more structured forms of

knowledge, and will support more activities to repair and organize the existing knowledge. We also hope to forge a greater connection to the existing Cyc database.

We expect similar projects to Open Mind Common Sense to appear, perhaps taking very different approaches. Perhaps someone will build a system that will try to read the text on the web, but will have people help it comprehend the more difficult passages. Also, these days many people have web cameras attached to their computers. You could imagine a more visually-oriented version of Open Mind where thousands of people helped teach their computers to recognize the appearance and behavior of various kinds of objects. And we need not limit ourselves to the web. Everyone has a cell phone these days, which will soon come equipped with more powerful onboard computers, cameras, and global positioning systems. We could all start to teach computers the patterns of our everyday lives by letting them see and hear us as we actually do things in the world. There is a goldmine of opportunity for people who are willing to accept that there are countless people out there who would be willing to participate as volunteers in the effort to help artificial intelligence researchers achieve their dream of bringing a new kind of life form into this world, the class of computer software with common sense.

References

- [1] Minsky, M. (1986). *The society of mind*. New York: Simon and Schuster.
- [2] Lenat, D. B. (1995). CYC: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11), 33-48.
- [3] Singh, P. (2002). The public acquisition of commonsense knowledge. In *Proceedings of AAAI Spring Symposium: Acquiring (and Using) Linguistic (and World) Knowledge for Information Access*. Palo Alto, CA: AAAI.
- [4] Barry B., & Davenport G. (2003). Documenting Life: Videography and Common Sense. In *Proceedings of IEEE International Conference on Multimedia*. New York: Institute of Electrical and Electronics Engineers.
- [5] Lieberman, H. and Liu. H. (2002). Adaptive Linking between Text and Photos Using Common Sense Reasoning. In *Proceedings of the Second International Conference on Adaptive Hypermedia and Adaptive Web Based Systems*, (AH2002) Malaga, Spain.
- [6] Minsky, M. (forthcoming). The emotion machine.