Back into Equilibrium: Balancing the Ordinary and the Extraordinary

Henry Lieberman

Media Laboratory
Massachusetts Institute of Technology
Cambridge, MA, USA

Introduction

Science and Art are both about the balance between the *ordinary* and the *extraordinary*. Science and Art both begin by removing us from our everyday experience. The concepts they introduce are often unfamiliar at first, and surprise or even shock people when first encountered. But with these new concepts, Science and Art then bring us back to everyday experience, albeit with a new perspective. Therein lies their value.

Interdisciplinary inquiry searches for the connections between Science and Art, and searches for new ways to increase the synergy between Science and Art. For this, I think it is necessary to understand the ways in which both Science and Art connect to our experience of everyday life, and also take us beyond it.

Scientists and artists usually stay focused on the specialized work that constitutes their contributions. They have the reputation in society of being somewhat absent-minded, sometimes neglecting to pay attention to the "common sense" that dominates the life of non-specialists.

However, a small group of scientists in the field of Artificial Intelligence have taken it upon themselves to conduct a detailed study into the nature of the body of human knowledge that we call "Common Sense". Common Sense is that knowledge that we expect an ordinary person to know. Common Sense is, by its nature, elusive. Since it is what everybody knows, you rarely have to communicate it to another person, so it is hard to articulate it. But understanding it is key to understanding how humans operate in the world, which is why Artificial Intelligence is so interested in it.

And understanding Commonsense knowledge is also key to understanding how Science and Art operate. As computers are increasingly used as tools by scientists and artists, they become partners in the activity of bringing everyday experience into their creations and creating new ways of looking at this experience. In this article, I hope to show how a better understanding of Commonsense knowledge can help build tools that help creators of art and of scientific knowledge in their search for discovery.

The Quest for the Extraordinary in Art

Many memorable pieces of art have as their goal to get us to understand the ordinary world in extraordinary ways. Art has always struggled against a naïve conception of its mission as a reflection of the real world. Berthold Brecht said, "Art is not a mirror held up to reality; it is a hammer with which to shape reality.".

Literal and symbolic paintings

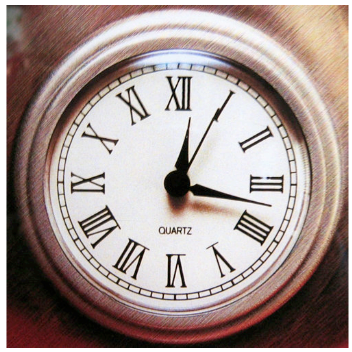

Art, above all, gets us to notice things. We all get used to seeing the objects and activities of everyday life; so much so that things that exhibit regularities start to disappear. We direct our attention to that which is different, novel, and unexpected. Everything else disappears.

Artists are keen observers of the world, and more so than most, they have the ability to capture and understand the recurring patterns of our existence. By first being able to see the world as it is, they are able to go beyond literally reproducing that world, and add their own perspective. That perspective then becomes the message.

Even artists who create abstract, non-representational art often start out in more directly representational forms, and often define their abstractions by a contrast with, comment upon, or rebellion from, the prior forms. It is significant that the rise of abstract art coincided with the popularization of photography; photography took over the mission of reproduction that was formerly art's, freeing art to more directly represent its own perspective. Subsequently, photography developed artistic perspectives of its own.

Artistic movements are often founded around concepts that embody particular perspectives; the abstractions themselves are not directly visible, but they provide ways of presenting material that causes people to change their perception. Impressionism in painting caused people to challenge the nature of visual images; rap music redefined the meaning of poetry; science fiction as a literary form gave authors a way to speculate about the future.

Artistic innovators sometimes show familiar scenes, but in a way that calls attention to their portrayal. Individual artists start by being observant of the world around them. They tend to notice things about the visual appearance of the world that go unnoticed by their peers. They then reflect on the differences between their own perspective and conventionally accepted perspectives on the subject matter. When the artist senses that he or she has a unique perspective to communicate, they go on to create original, creative works of art. The tension between directly reflecting the world and imposing a new perspective on it has been a motivating force throughout the history of art [Gombrich].

The Quest for the Extraordinary in Science

Science, too, proceeds by understanding the relationship between the ordinary and the extraordinary. Like art, it takes as its ground truth the shared experience of ordinary life. But it interprets that everyday life through new and abstract theories that interpret observed phenomena through a perspective of mathematical description, that itself cannot be directly observed.

Image of the sun and visualization of its magnetic field

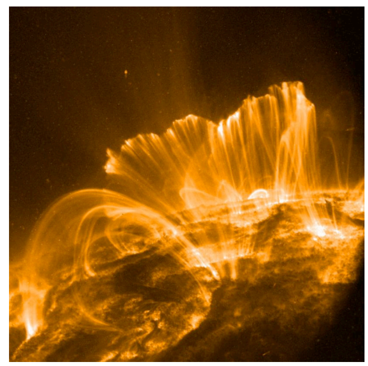

Science, like art, can come to surprising and counterintuitive conclusions about our interpretation of everyday experience. Scientists are also keen observers of everyday life, through their experience measuring things that remain unnoticed to the majority of us. They then try to interpret the empirical data according to theoretical frameworks that seek to explain it, by accounting for the data by appeal to general scientific principles such as the laws of electromagnetism.

The concepts introduced by science are often not directly observable. We cannot see that ordinary matter is made of atoms and quanta. We cannot see the DNA mechanisms that express genes in a cell. These concepts are invisible. It is only by the fact that they explain observable phenomena that they gain their power.

Where a scientist sees a discrepancy between his or her personal perspective and established convention and explanation, he or she perceives an opportunity to make an original and creative contribution. This often takes the form of a new theory that explains old data in new ways, as relativity displaced Newtonian mechanics.

Science and Art share common methodology

So we can express some of the commonalities between the methodology of scientists and artists by the following steps.

- ***Observe*** the real world;

- ***Invent*** new concepts that reference elements of the real world, but may not be directly visible to the untrained observer;
- ***Re-interpret*** real world phenomena using those concepts;
- ***Communicate***, first with peers, then the general public, to increase understanding of how these new perspectives increase understanding and/or provide pleasure;
- ***Absorb*** the ideas into the general culture, where they eventually become taken for granted, and provide the foundation for the next scientific or artistic movement.

Kuhn's Structure of Scientific Revolutions [Kuhn 62] details the social processes that surround theory generation. Indeed, the comparison between Kuhn's description of the social processes surrounding science and Gombrich's view of the history of art show striking parallels.

Disruptive ideas upset people's Information Balance

New artistic movements or scientific perspectives often have the effect of upsetting people's *Information Balance* – the tradeoff between the practical process of dealing with everyday life as we observe it, and the way we conceive it conceptually with scientific theories or artistic perspectives. It upsets this balance because it threatens with the prospect of "everything you know is wrong". Things we think we've figured out are suddenly called into question. Our Commonsense intuition, which served us well in the past, is suddenly wrong, seemingly for no good reason.

Both Kuhn's description of the history of science and Gombrich's description of the history of art are rife with descriptions of resistance to new ideas. The impressionists were rejected by the artistic establishment of the time. Quantum physicists were initially ridiculed, even by Einstein. Introduction of many digital technologies are now meeting resistance from technology-naïve segments of the public.

Underestimation and overestimation

The fact is that people always, and inevitably, both underestimate and overestimate new ideas. They underestimate the impact of revolutionary new ideas because they are trying to fit those ideas into their existing categories, interpreting them as only an incremental change. The digital revolution often gets underestimated as people tend to understand it by reference to previous technologies. You hear things like "E-mail is just another kind of mail" or "electronic books" (the latter phrase always seems to me just as oxymoronic as "horseless carriage").

Likewise, new technologies also get overestimated, and here Artificial Intelligence itself is often served as a cautionary example. When one starts to speak of intelligence, it invites the assumption that a super-intelligence that can answer any question is what is being proposed. People compare artificial intelligence projects to fanciful portrayals in science fiction, and are inevitably disappointed. This often dampens enthusiasm for the very real accomplishments of the field of AI.

Convergence leads to understanding

Many caution that fields like AI should take greater care not to over-promise what they can do, and tone down their claims of being able to represent and reproduce human intelligence. Indeed, the field has in the past been guilty of over-hyping what it can accomplish in the short term. But I would be sad if the field abandoned speaking of its long-term dream of modeling human thought processes. We don't yet know how far this can take us.

I think the key is understanding that both overestimation and underestimation of new ideas occur simultaneously, and both are inevitable. The saving grace is the fact that, with repeated exposure, more familiarity, and more thorough exploration as these ideas permeate more and more situations, perceptions tend to converge. Seeing the new perspective as a cure-all tends to diminish as its limitations and tradeoffs are illuminated. As the new perspective sparks consequences throughout society in many different situations, its impact is more fully appreciated, correcting the initial underestimation.

The final stage is where the new perspective is fully integrated by the public, and the idea becomes fully accepted, taken for granted, and finally, people forget what the world was like before that idea was introduced. We finally declare that perspective common sense, at which point it is fully invisible. That's how new Commonsense knowledge comes about.

The Quest for the Ordinary in Artificial Intelligence

The field of Artificial Intelligence has as its goal to get computers to reproduce some of the kinds of reasoning and understanding performed by humans. In order to do that, we need to create symbolic representations of the underlying knowledge that people have concerning everyday life. That knowledge, acquired by most people by the age of ten or so, is widely referred to as "Common Sense" knowledge.

A characteristic that makes Commonsense knowledge difficult to characterize is its invisibility. It is invisible because it is so simple that people do not usually take the time to express it to one another. It took a while for the field to realize the essential importance of collecting Commonsense knowledge.

This point was driven home to early researchers in natural language understanding and translation, when they tried to interpret simple children's stories such as the following.

> Mary invited John to her birthday party.
> He thought she would like a kite.

An idealized birthday party

Understanding this story involves understanding the following Commonsense facts:

> A birthday is traditionally celebrated with a party
> It is traditional to give gifts to a person on their birthday.
> A kite is a toy.
> A toy can be a gift.

This implicit knowledge is expressed nowhere in the story, but is essential for understanding it. Computers will not be able to understand human language until they are able to store the background knowledge necessary for them to perform the same kind of implicit inference.

Diapers and beer

In recent decades, a success of the field of Artificial Intelligence has been the technique of *Data Mining*; or as the field would now prefer to be known, *Knowledge Discovery*. Fixed-structure digital repositories of knowledge, often called "databases", were set up so that users could use them to answer specific questions. Knowledge discovery, though, aims and discovering the *questions* that should be asked, not, at least directly, the answers.

A story often told about the early success of knowledge discovery, is the "diapers and beer" story. In the USA, it is now common for supermarkets to offer affinity cards to their customers that offer discounts, in exchange for the consent of the customer to allow the store to track and record his or her purchases. For the first time, supermarkets could see who bought what, when.

When data miners started to examine which products were often bought together, they got some surprises. As expected, they found that people who bought diapers often also bought milk. But they also found that people who bought diapers frequently bought beer. Why?

A little reflection comes up with the scenario that yields the answer. Imagine a new mother taking care of a baby and realizing the diaper supply is running low. She sends her husband out to the store to buy diapers, and what happens? Now, many supermarkets place stacks of beer cans next to the diaper aisle.

The key to understanding this is knowledge of Commonsense scenarios of everyday life. Data mining itself is simply the process of searching for mathematical regularities or statistical correlations in a database. It can simply tell you when a mathematical relationship occurs – it can't tell you whether that relationship is important or surprising. For that, human judgment compares the statement of the relationship to the preconceived expectations that they bring to the situation. The unusual can only be discovered by contrast to the usual. Collecting and analyzing Commonsense knowledge can help us create a deeper understanding of the mundane.

Collecting and using Commonsense knowledge

Several projects in Artificial Intelligence set out the goal of creating collections of such Commonsense knowledge in order to provide a computer with a basis of shared understanding for communication between people and machines. The centrality of the Commonsense problem was identified by McCarthy as early as 1959 [McCarthy 59]. McCarthy adopted an approach based on mathematical logic, which remains popular in the field to this day [Mueller 07].

But there are other approaches, the search for which was motivated by the apparent mismatch between the precision of mathematical logic and the relative imprecision and flexibility of natural language expressions of Commonsense knowledge. Researchers were also frustrated by the slow progress of the logicists, who often fought bitter battles over minor differences between formalisms.

The most famous of the large-scale Commonsense projects is Doug Lenat's Cyc project [Lenat]. Cyc, short for "encyclopedia", has been collecting Commonsense knowledge steadily for over 20 years, entered by a team of professional knowledge engineers, and has roughly 3 million assertions.

In our lab at MIT around 2001, Push Singh and collaborators decided to harness the power of the Web, and created the Open Mind Common Sense site, where Web users were asked to enter Commonsense knowledge in English sentences and simple templates rather than formal mathematical languages. Since that time, it has grown to about 800,000 sentences. Natural language processing techniques pick out patterns of 22 distinguished relations, things like "A _ is used for _", or "You find a _ at _".

Below, the current version of our site, Rob Speer's Open Mind Commons.

Open Mind Commons

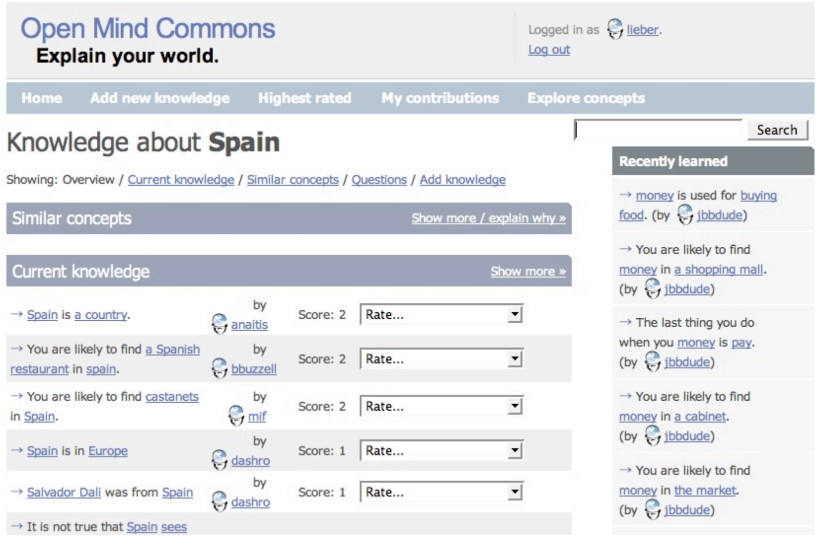

We have also started Commonsense collections in other languages and cultures. We have started to use it for both language translation, and, perhaps more interestingly, cultural translation by finding analogous concepts in different languages.

Several years later, myself, Hugo Liu, Barbara Barry, and others developed a methodology for using Commonsense knowledge in interactive applications. Because the knowledge is not complete and consistent, care must be taken in using it. It is not so good for accurate factual inference, but it can be excellent for helping an application determine the likely context of a user request; for making analogies; for providing plausible defaults; and many other uses. We have developed a wide range of applications in areas like predictive typing, speech recognition, management of media libraries, online help, educational applications, and much more. Some of these applications are surveyed in [Lieberman, et. al 05].

Surprisingly, Commonsense knowledge is useful in helping analyze user input for intangible and difficult-to-compute qualities such as affect or point of view. Recently, Rob Speer, Catherine Havasi and Jason Alonso have developed a technique called AnalogySpace, which can analyze a range concepts according to linear dimensions such as "good vs. bad" or "urban vs. rural".

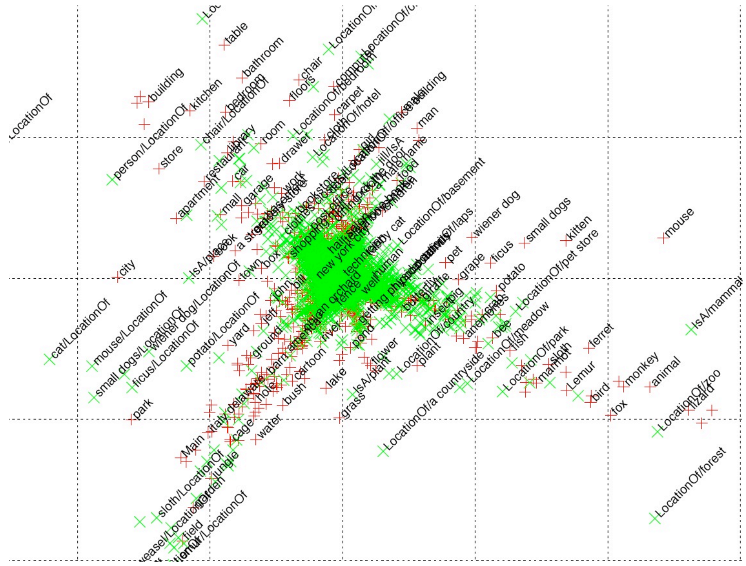

AnalogySpace graph: inside/outside vs. animate/inanimate

Commonsense and Storytelling

In particular, understanding the relation between ordinary and unusual knowledge has been useful to us in applications that involve storytelling. In the process of constructing illustrated textual stories, or documentary video, we have built systems that help retrieve media elements relevant to story composition in real time, as the story is being constructed by the user. Both storytelling and story understanding essentially depend on shared Commonsense knowledge between the author and the audience. Commonsense knowledge is used as the bridge to match up explicit elements in the story with the narrator's intent.

In Aria, a textual story is authored by the user, and relevant annotated photographs are automatically retrieved as the user is typing. In Barbara Barry's Mindful Camera, Commonsense knowledge is retrieved as real-time suggestions in a video camera, to reminder the cameraperson to remember to record scenes relevant to the story before they are lost. These applications are described in [Lieberman et. al 05].

Our latest effort in this area is Edward Shen's Storied Navigation, an intelligent video editor for documentary video, that connects the narrative intent of the story with specifics of characters, events, and dialogue presented in the video clips. Storied Navigation goes beyond simple relevance, and incorporates notions of story structure, themes, expressed emotions, and other concepts. This can be useful as a story construction tool, but also for creating nonlinear or branching stories in which a user can dynamically navigate and make choices that affect the progression of the story.

Storied Navigation

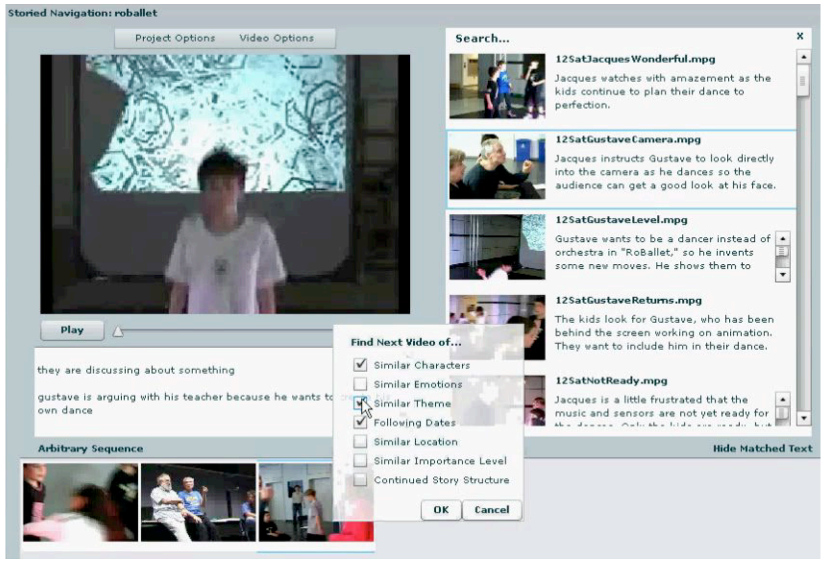

Using the ordinary to find the unusual

Sometimes, storytelling in Storied Navigation involves using Commonsense to find material directly relevant to a story, as in Aria. Sometimes, though, the intent of the author is to find something that is deliberately *not* ordinary, or that violates a Commonsense expectation. This is to create suspense, to create questions in the mind of the viewer that later get answered or resolved, or to direct the user's attention to a surprising element of the story. Thus the value of Commonsense is to be able to distinguish the ordinary from that which is not.

Schank, and others, have noted the theory of *expectation violation* for story understanding, and for general problem solving [Schank 90]. This says that, when humans interpret input, it sets up expectations due to Commonsense knowledge. When we hear of a birthday party, we know a certain sequence of events that will be likely to take place, a script: gifts, a cake, singing, etc. But perhaps the familiar elements are missing or wrong in a particular story. If subsequent input violates the assumptions, problem solving activity is immediately directed first to noticing the discrepancy, then to constructing an explanation for it. This creates a learning episode that we then store for future reference. Much of our problem solving capability comes from retrieving past stories and explanations, and modifying them as new conditions arise. Storied Navigation is built on finding story analogies.

Commonsense knowledge in tools for science and art

Understanding Commonsense knowledge is the first step towards computer understanding and computer processing of knowledge that goes beyond Commonsense. If we understand the ordinary, we are in a position to begin to understand the extraordinary, or understand the ordinary in extraordinary ways.

Scientists and artists who now use computers as tools for representing their work will find that representations of Commonsense knowledge will become essential in their quest to understand and develop new ideas. Scientists now use computers to keep track of data, and perform analyses on it. They use it both to test hypotheses and to develop the intuition from which new hypotheses are formed.

Artists use computers as "digital paintbrushes" for creating images, edit digital video for filmmaking, control sculpting machines, author new forms of interactive fiction, and other uses. They both manage the ingredients from which art is formed, and also serve to collect and access the raw material from which ideas for new pieces are generated.

But the next generation of both scientific and artistic tools will go beyond just collection, production, and visualization of data. Computer vision, speech recognition, natural language understanding, knowledge representation and inference will permit computer tools for artists and scientists to engage with the content as well as the form of their materials. They will help us both in our observation of the ordinary, and in our creation of the extraordinary.

References

Gombrich 60 — *Art and Illusion: A Study in the Psychology of Pictorial Representation*, E. H. Gombrich,. Oxford: Phaidon Press, 1960.

Kuhn 62 — The Structure of Scientific Revolutions. T. S.,Kuhn, University of Chicago Press, Chicago, 1962.

Lenat 95 — CYC: A large-scale investment in knowledge infrastructure, Douglas Lenat, Communications of the ACM, 38(11): 33-38, 1995.

Lieberman, Liu, Singh, Barry 05 — Beating Common Sense into Interactive Applications, Henry Lieberman, Hugo Liu, Push Singh, and Barbara Barry, *AI Magazine*, Association for the Advancement of Artificial Intelligence, Winter 2004.

Liu and Lieberman 02 — Robust photo retrieval using world semantics, Hugo Liu and Henry Lieberman, Proceedings of the 3rd International Conference on Language Resources And Evaluation Workshop: Using Semantics for Information Retrieval and Filtering (LREC2002), Las Palmas, Canary Islands.

Liu, Lieberman, Selker 03 — A Model of Textual Affect Sensing using Real-World Knowledge. Hugo Liu, Henry Lieberman and Ted Selker, International Conference on Intelligent User Interfaces (IUI 2003). Miami, Florida.

McCarthy 59 — Computers with Common Sense, John McCarthy, reprinted in V. Lipschitz, ed., *Formalization of Common Sense*, Ablex, 1990.

Mueller 07 — *Commonsense Reasoning*, Erik Mueller, Morgan Kaufmann, 2007.

Schank 90 *Tell me a story: A new look at real and artificial memory*. New York, NY: Charles Scribner, 1990.