

SYNTIENCE INC.

"We Understand"

171 Main St., Suite 276
Los Altos, CA 94022

T +1 408 802-7229

Monica Anderson
monica@syntience.com
syntience.com

AI RESEARCH IN THE 21ST CENTURY

INTRODUCTION


Computers are Logic based devices. We used to believe implementing Intelligence in computers would be easy because we thought brains were built on Logic. AI researchers and programmers were especially skilled at logical, step-by-step Reasoning and analysis and when thinking about thinking, all they saw was Logic and Reasoning. But by the end of the 20th century it became clear that there was another component to Intelligence that had been severely neglected – intuitive Understanding.

REASONING REQUIRES UNDERSTANDING

Reasoning is a conscious, goal-directed and Logic-based step-by-step process that takes seconds to years. In contrast, Understanding is a subconscious, aggregating, Intuition-based and virtually instantaneous recognition of objects, agents, and concepts.

There is nothing mystical about Intuition. It is a straightforward process of recalling past experiences and matching them to the current situation. If you see a chair, your Intuition tells you it's a chair by matching your sensory input patterns to those of past occasions when you have seen chairs. You don't have to reason about it.

Our Intuition also allows us to understand relationships, complex patterns, trends, abstractions, contexts, meaning (semantics), and the big picture. With little effort we recognize the image of a chair, the flavor of an apple, the meaning of a sentence, or the face of a friend. Neither requires Reasoning, and humans can recognize and understand all of these without effort, and within milliseconds. Current computers and AI systems cannot do this.



More than 99.99%¹ of the processing in the brain is subconscious, which makes conscious Logic based Reasoning look like a **paint-thin layer** on top of our subconscious and Intuition-based Understanding.

In the 20th century, AI research was overmuch concerned with Reasoning. In the 21st century we must focus our AI research and resources on Understanding.

Our computers need **Artificial Intuition** in order to recognize sensory input data and understand concepts at low levels. Only after they understand something will they have something to reason about. The Understanding part could well be straightforward and easy to implement in computers compared to our attempts so far to automate Reasoning; we just need to do these two in the correct order.

¹ Tor Norretranders: "The User Illusion" (1999) discusses this and provides several references

REDUCTIONIST STANCE VERSUS HOLISTIC STANCE

Computer Science and a few other scientific disciplines such as Mathematics and Physics are permeated by a problem-solving strategy and philosophy called “The Reductionist Stance”. In this Reductionist paradigm we solve problems by extracting them from their environment and by splitting them into smaller problems. Once we understand enough of the problems we create Logic-based simplifications called “models” such as theories, equations, and computer algorithms. Models describe simplified and **context-free**, reusable and portable slices of reality that can solve entire classes of similar problems. But before models can be created, chosen, and applied, and the results interpreted in the current context we must **Understand** the problem domain.

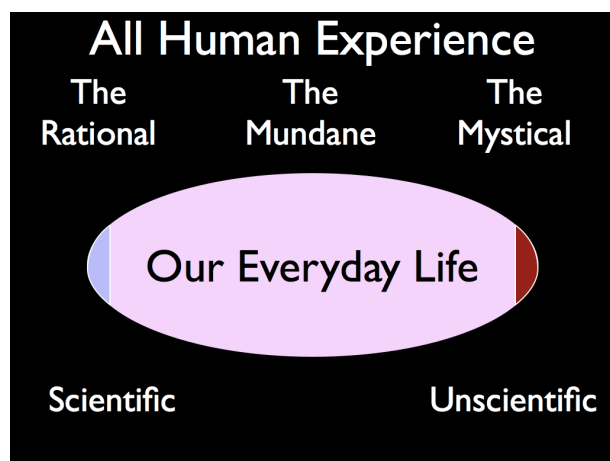
In contrast, when adopting a “Holistic Stance” we solve the actual current instance of the problem **directly**, without using models. We solve it in its **actual context** without attempting to bring it into the laboratory. We don’t attempt to subdivide it into smaller problems. We don’t look for known models that might fit the problem. Instead, we attempt to match a large set of simple and self-assembling patterns (and patterns of patterns) against anything and everything in the problem **including** its current context. The final assembly of patterns that matches **is** our Understanding of the problem. Our brains use sets of neurons to represent these patterns. The set of activated neurons **is** our Understanding. What else could it be? Note that this provides us with a very specific and useful definition of what Understanding means in humans which is also **implementable in computers**.

This application of patterns may seem like much more effort than the use of models and Reductionist methods. But it can be done effectively and mindlessly, without requiring a high level of Intelligence, by both brains and computers. It just requires a very large database of patterns; we call this database “experience”. Brains gather experience over a lifetime, and stores it as these self-assembling patterns. Computers could do the same.

Note the difference: models **require** Understanding (“Intelligence”) for creation and use. Patterns don’t. Instead, self-assembling patterns **provide** the Understanding that is required for all Reasoning, including model creation and use. Only **after** recognition, abstraction, and Understanding will you have something to reason **about**.

SCIENTIFIC VERSUS UNSCIENTIFIC PHENOMENA

Some people believe the world splits cleanly into the Scientific and the Unscientific. This is incorrect. Worse, this belief causes a blindness to entire spaces of approaches and solutions that has hindered progress in AI.



The majority of phenomena in the world, including almost everything in our mundane everyday lives, are neither scientific nor unscientific. Consider the task of walking a mile into town to buy a newspaper. People can easily do it, so it is not mystical, but it is also out of reach of science to reliably duplicate this task, for instance by building a robot. The Mundane world is deeply complex and changes too rapidly.

And the Mundane is exactly the domain of Artificial Intelligence. How to navigate the aisles of a grocery store if you are a robot; how to understand the slurred speech of the cashier; how to understand written language in a newspaper; how to understand a complex world with its many layers of meaning in nested contexts; how to deal with many other intelligent agents with goals often at odds with your own.

Brains solve all these little context-rich problems every day seemingly without effort. Brains can do it because they don’t operate logically and don’t use Reasoning for the majority of their functionality. They do it using a much simpler algorithm that instantly recognizes nested patterns. When we employ Holistic thinking, context isn’t a distraction to be removed. In fact, if a problem is too hard, we search for **more** context. More context means more patterns can be matched which means more powerful higher level patterns can be brought into play.

AI systems have to be able to deal with context in the same way. This means Understanding the mundane world Holistically. If we restrict our AI implementation to Logic based Reasoning then we will only be able to operate in the thin slice of rational problem domains. Such systems may be able to reason about math and Logic but they will not be able to deal with the world at large. Real world concepts cannot be “defined”, they must be learned.

Any sufficiently Reductionist AI effort is indistinguishable from programming. If a programmer attempts to describe the world to a Logic based AI, for instance by creating ontologies, he'll never finish the task. The world is too rich. The Cyc project – the largest and most famous AI project ever undertaken – has been trying to describe the world using predicate calculus for decades; it is the poster project for Reductionist approaches to AI. But Cyc will never approach anything worthy of the term "Intelligence". It has been told many things and can recite many definitions but Understands nothing. This is the difference between "Instructionist" top down education and "Constructionist" bottom up learning – a distinction poorly understood even in human education.

One of the biggest hurdles in the transition from a Reductionist to a Holistic stance is that the Reductionist stance works so well for simple problems, and thus is very seductive to beginners. But we must learn to fully understand the value of context in solving complex problems so that we can stay our hand before attempting futile reductions when facing an irreducible problem. This is a tradeoff. Let's look at what we are giving up and what we might gain, if we are AI researchers trying to adopt a Holistic stance:

VIRTUES OF REDUCTIONIST SCIENCE

Optimality	The best answer
Completeness	All the answers
Repeatability	Same answer every time
Efficiency	No waste of resources
Transparency	Understand the process of getting the answer
Scrutability	Understand the answer

These are some of the virtues we learn to appreciate as scientists. These are what we learn in school, especially as part of a Ph. D. education and especially in Computer Science. We internalize these virtues as obviously desirable properties of any investigative process. But this is also where our blindness begins.

At the lowest levels of the brain, scientific principles like Logic, Reductionist methods, proof, falsifiability, P-vs-NP, or the six virtues above simply don't enter into the picture. **Human minds neither provide nor need any of these values**; human problem solving in our mundane everyday life is nothing like science. When pulling up to a stoplight, you don't compute a differential equation to determine how hard to push the brake pedal. You use your Intuition and experience. The deceleration may be a bit jerky (not optimal), maybe you should have started braking earlier (incomplete), you did it slightly differently yesterday (not repeatable), and you have no idea why you did it the way you did (inscrutable); your brain activated untold neurons unnecessarily that did not contribute to the result (inefficient), and so on. But it is good enough to keep you alive, which is why Intuition-based Intelligence evolved in the first place.

We all use the Holistic stance when we are young. Through education (and by independently discovering in childhood how to build naive models) we learn the virtues of model use, of Reductionism, and science. We want to be solid Reductionists because our education always rewarded us for using Reductionist methods. But Understanding is more important than Reasoning. Enlightenment is the ability to see both kinds of solutions and to know which one is the most appropriate. We in the AI research community have repeatedly chosen poorly and we have paid the price. I have calculated that one million man-years may have been wasted on Reductionist AI.

DISADVANTAGES OF REDUCTIONIST SCIENCE

Limited applicability	Only works in problem domains where models can be created and used
Understanding	Requires Understanding the problem, the problem domain, and candidate models
Correct input data	Requires correct and unambiguous input data for correct results
Brittleness	Models may fail catastrophically if used in situations where they do not apply

Building **robust** and **useful** Reductionist models of complex dynamic systems such as physiology, drug effects and interactions in the body, of people, societies, industries, economies, stock markets, brains, minds, or human languages is **impossible**. These are “Bizarre” problem domains – defined as domains where models cannot be created or used and are discussed in detail at <http://artificial-intuition.com>.

Reductionist model-based attempts at AI will typically fail spectacularly at the edges of their competence but there may be no easy way to detect this failure since the borders of the model's applicability can be difficult to determine mechanistically. Humans know what they don't know and will at worst make human-like mistakes. Computer systems using inapplicable models make mistakes that make headlines and have shown us time and again how little computers really “know”. The AI community calls this “Brittleness” and it is a major problem.

The alternative is to get by without models, or to use weaker models with lesser requirements. The life sciences often do this since they may have no choice; life scientists may not be aware of the distinction. In a paper published 1935, Dr. Lionel S. Penrose, a pioneer geneticist working in the field decades before DNA was described (and incidentally, father of Sir Roger Penrose) first encouraged use of what he called “Model Free Methods”. If we examine various methods used in the life sciences we will find that many of these are Model Free or use Weak Models. Trial and error, pattern matching, table lookup, adaptation and other kinds of learning, search, evolutionary computation, markets, statistics, and even language use can be viewed as either Model Free Methods or Weak Models, depending on details, definitions, and circumstances. I discuss some of these issues in talks available as videos at <http://videos.syntience.com>. All truly Model Free Methods are Holistic, and all Holistic Methods are Model Free. AI research in the 21st century needs to learn from the life sciences:

BENEFITS OF A HOLISTIC STANCE

Applicability	Can be used anywhere, including in Bizarre problem domains
Ignorance	No need to understand the problem or the problem domain
Directness	Problem is solved directly. No need to interpret model output in current context
Learning	Many Model Free Methods provide learning and gathering of experience
Abduction	The ability to jump to conclusions (often correctly) based on incomplete input data
Novelty	Holistic Methods can provide true novelty; some depend on it for their function
Robustness	Graceful degradation. Failures are non-catastrophic, “human-like” errors.
Skepticism	Robustness against internal errors also provides ability to deal with erroneous input
Disambiguation	Provides context based ability to handle ambiguities in problem statement and input
Understanding	Provides saliency, abstractions, situation semantics, and language semantics

The last four benefits (in boldface) in the table above may be available as emergent effects when using Artificial Intuition or other advanced MFMs. The others may be available even in simpler MFMs.

DISADVANTAGES OF A HOLISTIC STANCE

Fallible	No guarantees of reaching a solution, or that found solutions are correct
Non-portable	Solving a problem may not be much help on other problems or in other domains
Wasteful	May be wasteful of computational resources and may require parallel processing
Inscrutable	Solutions don't necessarily provide any clues to how to create models for problem
Opaque	Solutions emerge. We cannot always analyze the details of how this happens.

There is no conflict between emergent Understanding and the inscrutability of the result. Direct application of simpler Model Free Methods such as evolutionary computation may well provide working but inscrutable solutions using an opaque process. Advanced Model Free Methods like [Artificial Intuition](#) can **provide**, as emergent effects, **Robustness**, **Skepticism**, **Disambiguation**, and **Understanding**. This is how MFMs can provide the foundations required for Reasoning and the ability to use Reductionist stance and Models. We could call this **Emergent Reductionism**. We may be able to create an AI that Understands and Reasons, but we may not understand exactly **how** it Understands.

Fallibility looks like a deal-breaker to a Reductionist. But given a Bizarre real world, where Reductionist Models cannot be built and a Reductionist approach could not even get started, we have no choice. Humans in the Bizarre Mundane Real World operate Holistically, and as our experience grows we find we fail very rarely in common everyday situations. Graceful degradation and emergent robustness allows the Holistic Methods to continue to work as the details of the mundane tasks change. The normal road to town may be blocked by construction. We add some more matching patterns to our Understanding of the situation and we spontaneously deal with the new context, adjust to it, and learn from it. The same kind of robustness will be available to our computer based AIs if we program them to use Holistic Methods.

AI SHOULD HAVE BEEN A LIFE SCIENCE

Not all scientific disciplines are dominated by Reductionism. In the life sciences, like Biology, Ecology, and Psychology, it has long been recognized that Reductionist Methods are insufficient. A few decades ago there was a kind of "Physics envy" in the life sciences. Every discipline was measured the way Physics was measured – by how well you could make long term predictions. This is easy when you are predicting movements of pendulums but very difficult when dealing with the population of muskrats in New England. The life sciences have achieved impressive results and have shaken off this inferiority complex by using alternatives to Reductionist Methods, including many Model Free Methods. Here is a clue to what has been wrong with AI research.

There has been a fatal mismatch between the properties of the problem domain and the properties of the attempted solutions. In the above I have argued (like so many others) that Intelligence is the ability to solve problems in their context. In contrast, programming deals with **discarding** context and breaking down complex problems into simple subproblems that can then be programmed as portable and reusable subroutines in the computer. Programming is the most Reductionist profession there is and has therefore mainly attracted the kinds of minds that favor Reductionist approaches.

That means that one of the most Holistic phenomena in the world – Intelligence – was being attacked by the most Reductionist researchers in the world. This is the mismatch; the biggest problem with AI research in the 20th century was that it was done by programmers.