A pragmatic approach to using AI for Business Management Analytics

Combining human and machine strengths in neural networks to produce usable knowledge for complex business management challenges

Jesse Glick *and* **Joseph Glick**

Abstract

As our scientific understanding of the human and animal nervous system has grown, computer scientists have achieved amazing results by imitating in machines what biologists discover in living organisms. However, the biggest leaps forward in the field of AI are in imitating things that humans, even small children, do effortlessly and without much thought (understanding speech, recognizing people and places, etc.). While these advances enable a variety of tactical solutions for security, marketing and consumer product enhancement, they offer little value to managers who must address the Volatility, Uncertainty, Complexity and Ambiguity[1] that drive the opportunities and risks of the real world. Two decades of experience in applying cognitive science to business solutions have taught us how to combine proven methods commonly used in IT for decades with approaches successfully used by human subject-matter experts for millennia to discover high-value knowledge previously beyond reach.

What is a neural network?

In biology, a neuron is a cell that takes input signals from other neurons, and processes them according to some rule in order to produce an output signal that in turn is consumed by other neurons. The **intelligence** of this system lies less in the rules used by each neuron to process its inputs than in the **arrangement of the connections**, or *synapses*, by which the output of one node feeds into the next. This arrangement of connections becomes the **context** for interpreting new data input.

In computer science, a *neural node* is a tiny piece of software that operates similarly: takes inputs from other nodes, and processes them to produce an output that in turn serves as an input to still other nodes. These nodes form the building blocks for structures known as *neural networks*, which can range from simple structures of a dozen or so nodes to complex structures of hundreds of thousands. Their processing, however, differs from the Boolean algebra at the heart of today's computers.

Biological neurons produce binary, or true/false, outputs: either the neuron is firing or it is not. The inputs are either stimulating or suppressing, and the neuron fires when the stimulating inputs exceed the suppressing ones by a given amount. This is optimized for pattern recognition: a neuron to recognize dogs, for example, might be stimulated by a canine tail, paws, fur, muzzle, teeth, etc., might be weakly suppressed by catlike features such as pointy ears (which some dog breeds have), and would be strongly suppressed by features such as scales. The idea of something being a **better or worse match to a concept**, rather than definitely meeting or failing to meet a precise definition, is key to the way the

[1] VUCA – acronym coined by US Army

human brain works. In the 1990s, this concept was well known by the term “fuzzy logic,” [2] though ideas on how to apply it in practice could also be (and often were) described as “fuzzy.”

Our approach at Expertool has been to focus on the stimulation/suppression process to arrive at a better or worse conceptual match, simplifying it for the sake of easy and auditable computation to four states: impossible[3], possible, likely, and confirmed.

Another aspect of biological nervous systems is that their structure changes to record correlations learned by experience. If two neurons are repeatedly triggered at the same time, a new synapse forms between them, so that one stimulates the other. A classic example of this is Pavlov’s famous experiment where he trained dogs to associate bells with food.

In computer science, this is imitated by searching for statistical correlations. Perhaps the most powerful engine to date for doing this is IBM’s Watson. It famously won a game of *Jeopardy!* against the two most successful human players ever by searching a large base of textual materials for phrases that commonly occurred alongside phrases in the clues.

Can computers be superstitious?

The statistical approach can be very powerful, and the knowledge that it produces can be valuable. However, this knowledge is not always reliable: in its *Jeopardy!* win, after sealing the victory in the “Double Jeopardy” round, in “Final Jeopardy” Watson then proceeded to reply “What is Montreal?” to a clue in the category “U.S. Cities.”

The same thing happens in the biological equivalent. When a random, unexpected event happens at the same time as something important, we tend to want to correlate the two. Classical philosophers call this the *post hoc ergo propter hoc* fallacy. If it becomes rooted in human thought, the way Pavlov’s dogs continued to salivate in response to the bell long after he stopped feeding them along with it, we call it a superstition.

Because of this uncertainty in purely statistical methods, at Expertool we combine them with the cognitive suppression approach. If Watson had been endowed with this supplementary approach, then the category of "US Cities" would have produced an "impossible" state for the answer of "Montreal".

Our methodology uses an additional approach to reduce random patterns and false positives. In the past few decades, scientists have realized that natural phenomena tend to follow “fractal” patterns, which look similar at both large and small scales. Because of this, we explore the different possibilities using relatively small data samples, which are much less demanding of computational resources and produce initial results without waiting for an integration project to be completed. With cognitive suppression and human judgement, this may be sufficient to drive strategic decisions. Otherwise, the results may suggest an approach for designing a “big data” analysis. If the “big data” analysis finds a similar pattern, this fractal-like pattern (similar on large and small scales) can establish the business case for a major operational solution.

[2] http://plato.stanford.edu/entries/logic-fuzzy/
[3] i.e. – impossible within the defined context

Tried and True IT

Unless they're made of real biological neurons, "neural networks" are, fundamentally, computer software. As such, we believe they can be most powerful (and integrate the easiest with existing IT systems) if they incorporate IT concepts and methods that are well known and commonly used, and have proven their usefulness over time.

Relational Data

If one technology can be considered the lifeblood of information systems, it is relational data. It has been around in something like its present form since the 1980s, and its roots go back earlier. One way or another, it is included in nearly every IT project big or small since then.

Entity Relationship Models, the conceptual designs of relational databases, can form a rich model of reality using the same web of connections as a neural network. Because of this, at Expertool we like to use relational data to form much of the structure of our neural models.

Word Searches

Almost anyone who has used a computer in the past twenty years has at one time or another typed a word into a search box to find the information they need. Another method we use to structure our neural networks is the use of word indexes – we form connections between nodes representing concepts and those representing elements of text containing the corresponding words or phrases.

Arithmetic and Statistics

Another IT staple is the spreadsheet – data laid out in rows and columns, with new cells calculated arithmetically from the original data. At Expertool, we supplement our "relevance" nodes simulating biological stimulation/suppression with more classically-computing arithmetic and statistical nodes, which perform computations not only along rows and columns of data, but also along interwoven neural pathways.

Tried and True Human Reason

Artificial Intelligence is all about imitating human intelligence, and pushing the capabilities of machines beyond those of standard computing methods. Neural networks in particular are designed to imitate the human brain. Because of this, we believe some of the same methods we humans use to get the most out of our brains can be used to get the most out of computerized neural networks.

Instinct and Education

Humans (like many animals) like to explore the world, make discoveries, and draw conclusions. Using neural networks, especially with some sort of mechanism for Pavlovian correlations, researchers have built robots that do the same. Most of what humans accomplish, however, is not by reinventing the wheel. Some approaches are instinctual and improve through practice, such as simple arithmetic or tracking and intercepting a moving object. However, significant achievements begin with learning what others already know – leveraging existing expertise.

At Expertool, we design each neural network with a specific purpose in mind, using available expertise to frame and define content and methods. We expect the constructed model to make new discoveries, not just to tell us what we already know, but the basic methods used to explore are built-in – you could call them "instincts," and they are "taught" by additional domain or organization-specific expertise.

Collaboration

It has been recognized for some years that the most powerful solutions come from human-computer collaboration.[4] The "relevance" nodes of an Expertool model are designed to collaborate with humans. They usually produce either a "possible" or "impossible" answer, and wait for a human to confirm. Sometimes they might produce a "likely" answer and wait for a human to deny. Many Expertool neural networks have a "navigation" section, where the scope of analysis is determined by the process of a human selecting from a list of options, leading to some suggestions by the neural network, from which the human chooses leading to more suggestions – truly collaborative discovery.

Integration

Perhaps the biggest source of the power of the human brain is the holistic architecture - everything in it seems to be connected, directly or indirectly, to everything else. When solving a problem, we don't just stick to one method or information source (unless it's an artificial practice or test problem). We go through our whole toolbox of methods, our whole library of reference materials, and our whole network of human collaborators to find what we need to solve the problem.

Similarly, Expertool models are built using a holistic methodology, incorporating potentially relevant data sources and analytical methods. A neural structure may be partly based on relational data, partly on word indexing, and partly hand-built. It may include relational data from different databases, or a variety of data structures such as tables, trees or hierarchies, and dictionaries. It can mix statistical methods with rigid binary ones, and abstract concepts with quantitative values[5]. It can scope fully automated methods with a "navigation" section driven by human selections. The only thing uniform within a model is its purpose.

Concluding Summary

In the *Volatile, Uncertain, Complex* and *Ambiguous* world of business management, our approach to artificial intelligence has rapidly produced the most useful knowledge at global entities faced with perplexing challenges. The pragmatic solution has been using neural networks that -

- Mimic the stimulation/suppression logic of biological neurons
- Use learning by correlation but do not over-rely on it
- Apply well-known, tested and successful computing methods by:
  - Leveraging relational data and Entity Relationship Modeling
  - Using word indexing across textual sources to create new connections
  - Performing arithmetic and statistical computations across newly discovered pathways
- Imitate successful human problem solving by:
  - Being "educated" by human designers in the approach to solving particular problems
  - Interacting and collaborating with humans throughout the analytical process
  - Integrating a variety of analytical methods, data sources, and data structures into a single neural network

The above approach to AI blends cognitive science, computer science, domain expertise and managerial judgement into a pragmatic balance to address real-world business problem analysis challenges.

---

[4] *Integrating Machine and Human Computation, Frontiers of Engineering* – National Academy of Sciences 2013

[5] An abstract concept node may be linked to an array of potentially relevant textual and mathematical nodes