
CIO's Guide to Adopting Cognitive Technology

If you are not CIO of a progressive IT organization, this is going to be a very boring read. Our apologies. See you next time.

If you are CIO of a progressive IT organization, you are thinking about adopting Cognitive Technology in the near future. You realize that cost is just the starting point. Once a process is automated, it should be able to scale to unprecedented breadth and/or depth. It should also make the organization super-responsive. You know that the right way to think about this is to ask your business managers what could they do if they had 1,000 interns somehow telepathically connected. You expect cognitive technology to come close to that scenario. In the least, it should be completely transformational to your business process.

We agree. In fact, we thought we should share some of our some-say-considerable experience. If you have recently acquired cognitive technology, do leave a comment here with your own.

Accept the fact that Cognitive Technology is not 100% accurate.

In 2011, IBM's Watson Supercomputer played the popular game show, Jeopardy, against two of its most successful human contestants, and won. This event is often said to be a great milestone in development of cognitive computing. It is important to note that even though Watson won the game, it did not get every question right. The more you learn about cognitive technology, the more you will realize that anyone promising 100% accurate answers is lying outright.

Lets take the process of extracting nouns or noun-phrases out of any sentence. This natural language processing (NLP) step is a building block to any cognitive system. Current state-of-the-art software stacks, like Stanford NLP or Ling-Pipe, claim 97% or higher accuracy for noun-phrase extraction. However, when we took a set of ~1,000 real world documents and fed it to both the systems, the overlap was shockingly low at ~70%. In other words while the two softwares think they are 97% right, they are only 70% right.

Accept it. This is a reality. Noun phrase extraction is a very basic step. Cognitive computing engineers have had to design systems where the input itself is at best 70% accurate. There are many ingenious solutions that add layers to make this number better, but a successful workflow based on a cognitive computer should still include human failsafes and oversight. Lets plan for errors and manage them well. Lets definitely not plan \$100m trades based on a signal from a cognitive computer.

Avoid third party human layer around cognitive products.

An IT application needs to run on its own. It does need human support for on-boarding, training, trouble-shooting, upgrades, etc. but it really needs to run on its own. Some cognitive vendors stretch the idea of support to create teams of their employees even for routine operations. These so called experts are the only ones that interact with the technology and with your data to

get you meaningful results. The aim is to deliver that 100% accuracy their salesperson promised you - there is no known cognitive computer better than a human brain.

This may be a great arrangement if you plan to engage cognitive applications for some one-time project. However, if this is going to be a core part of any of your business processes, please insist on direct access to the technology. Three reasons -

- First, if you agree to the service, the vendor is not selling you the technology. They are actually selling it to themselves and charging you more for something they would have happily provided even without that cognitive technology. You do not gain direct experience, which is important in these formative years.
- More importantly, you are letting the vendor make crucial business decisions for you and learn from your data, at your expense.
- Lastly, you can never push the system to its true potential unless you control it, which, as we have said before, multiple times, will be such a wasted opportunity.

Instead, assure the vendor that you understand the inherent inaccuracies of a cognitive system and have planned for human failsafes and oversight in your business process. Remember that the salesperson promised you 100% only because he assumed given your accomplishments that 99% will not be acceptable. You should insist that your team be trained to handle the system's quirks. You should also insist on single tenant deployments, which lets the vendor better customize the solution for you. No, in this world of cheap compute this is not blasphemy.

Walk away from the traditional RFP process.

Not only do the cognitive computing engineers have to deal with inaccurate inputs, everything the academic world promised them is turning out to be untrue. They were told to use TF-IDF for search 30 years ago, yet none of the successful search engines does that. They have been told about wonderful semantic algorithms, but these have $O(n^4)$ complexity and are impossible to scale. They are being asked to borrow from the world of distributed computing and Big Data, but actually the cognitive systems are not that easily distributable. (Write to us if you want to talk more about why.)

Left to themselves, these engineers have had to rethink basic principles. Not always a good idea. Well, alright: Always not good an idea. As a result there is a hardly any standard benchmark or test that is truly capable of objectively comparing different cognitive systems. Do not imagine one and force it on your vendors. Maybe it is better to identify the key business metrics that you care about, and let the vendors compete for that. Counter-intuitively, (and self-servingly) cost of the system is probably not the best metric to include in this list. Instead its probably better to take a consultative approach and help cognitive computing vendors design highly effective bespoke systems using their capabilities.

At the same time we think it is fair to ask those pesky, prying, impolite questions about the underlying technology. It is also prudent to think about:

- What is the truly cognitive part of their technology - Word Clouds and Sentiment

Analysis are not really cognitive!

- What happens if the throughput goes 10x, as your business grows?
- What happens if the nature of the content being processed changes drastically with changing needs of your business?
- How can the confidence interval around any recommendation be made transparent?

If someone tells you they have a TF-IDF system coupled with Latent Semantic Analysis running over a Hadoop cluster, you know something is wrong. (Write to us if you want to talk more about why.)

Of course, like any other IT deployment, cognitive technology will disappoint before it delights. Well - you already know that.

What is Next-Generation Enterprise Search?

Coseer's search solutions are transforming industries from healthcare to finance. Our point-and-shoot AI trains finds answers and insights with 95%+ accuracy within 4-12 weeks - all of this in 100% security. The reason? We founded Coseer on the principle that computers should take care of the boring stuff so that humans can focus on creativity and judgment. To that end, we've built enterprise search solutions to complete complex workflows just as humans would in a fraction of the time. Fortune 500 leaders are using Coseer to speed up and automate their most complex work.

We follow a tactical approach to enterprise search:

- We deliver 95-98% accurate solutions within 4-12 weeks.
- Our solutions deploy entirely behind your own firewall for 100% security, and every decision point is logged for full transparency.
- You add the finishing touches, but our point-and-shoot AI practically trains itself. No more huge training data sets or time wasted annotating and tagging.

Visit our [website](#) for in-depth case studies, ROI breakdowns per industry, and other insight.