

How AI-Driven Search Could Bring Us Closer to the Intelligent Workplace

Seth Earley, CEO Earley Information Science

Tobias Goebel, Director Emerging Technologies

Artificial intelligence (AI) isn't a new magic concept.

AI has been kicking around for a while now, though many associate the concept with Alan Turing, who introduced what we know as the "[Turing Test](#)" in a 1950 paper, "Computing Machinery and Intelligence."

What is new is the ability to scale human to machine conversations – once the realm only of humans – where the machine provides relevant answers to specific questions at an individual context.

We see this scale in the mass market thanks to entrants from large consumer electronics brands such as Apple's Siri and Amazon's Alexa, but this capability has yet to reach the corporate market.

What Will Push AI into the Workplace

AI-enabled search promises to transform the way people interact with information and digital assets, driving new efficiencies and creating value from information that has been all but lost in the "digital junk drawers" that are our corporate information management systems. For AI to deliver on its promise, these junk drawers need preliminary organizing structures before the vision of conversational interactions can be realized.

In other words, *you can't have AI (artificial intelligence) without IA (Information Architecture)*.

So what goes on behind the scenes to make this "magic" happen? Knowledge engineering helps businesses realize the promise of AI on enterprise search, information access and knowledge management, and allows us to progress toward establishing an intelligent workplace.

Search as a Conversational Interface

Search is a conversation. When you type a query into any search engine, it returns results that are an approximation of the searcher's intent.

"Web queries are short, ambiguous and an approximation of the searcher's real information need," according to [search experts Ryen W. White, Joemon M. Jose and Ian Ruthven](#). Therefore, disambiguation is typically part of the search interaction. (If you say, "Tell me about AI," I will respond with a clarifying question "What would you like to know about AI?" to narrow the topic.)

Chat is a natural extension of the search interaction. Chat feels more natural with the advent of voice recognition and natural language inquiries, such as when we ask our smartphones where the closest gas station or sushi restaurant is located. It is the transformation of queries into natural language.

Early experiments in natural language queries have grappled with the challenge of ambiguity. Words have multiple meanings (i.e., "what is the stock price..." can result in finding the price of "a unit of financial value," "the handle of a gun," "a base ingredient of soup," etc.).

The marketing for Amazon's Echo promises users can search for a large scope of topics via the chat interface, Alexa. This broad approach does not align to workplace needs, where language and terminology is specialized. Generalized needs cannot be implied from the variations in users' requests.

Corporate information is more nuanced and complex than gas stations and restaurant locations.

Conversation Gets to the Heart of Search Faster

Traditional search is limited because it does not rapidly filter the query to its true essence. People use search terms that fail to deliver the richer meaning that a conversation can elicit.

But when people talk, they disambiguate. They have a conversation which gets to the heart of the query faster. This process occurs through a series of refinements that come as a result of the give and take that questions and answers provide. For example:

CFO: "How can I help you?"

User: "I need the quarterly report numbers for net revenue"

CFO: "For what division"?

User: "Consulting services"

CFO: "For which Fiscal quarter?"

User: "Q3 and Q4 2016"

CFO: "Here is a link to the 2016 annual report. Page 16 has Consulting Services revenue by quarter. Does that meet your need?"

User: "Yes, thanks."

The conversation zeroes in on the specific need through this refinement process, and saves the user the trouble of looking through a long list of search results and opening individual links repeatedly until they find the information they need.

An iterative exploration of a topic through conversation provides a higher chance of achieving a relevant answer, as opposed to receiving an entire set of possible answers at once, as in an FAQ setting or a Google search result.

This is what AI brings to enterprise search. The iterative process results in higher efficiency (time-to-information), increased accuracy (retrieval of the right information), and a greater utilization of corporate information assets (information is put to work).

What AI-Driven Search Looks Like in Practice

AI-driven search, in the form of digital assistants and chatbots, builds on activities we already carry out today. For example, a popular corporate travel administration portal includes a chat box where users can enter free-form text describing their travel needs:

"I need to fly from Boston to Chicago on March 20, and return on March 24."

The system parses the request, conducts the search and returns the same flight choice result screen that is produced using the standard form-based interface. The difference is the user made the request in conversational language, avoiding the confusion that can occur when switching modes to accommodate the browser-based flight search form fields.

Here is a more mundane yet very useful example of AI's potential impact as a new search interface that you can try today:

Ask Siri, "Show me emails received from Joe Smith since Monday."

Assuming that a user has been performing tasks on a Siri-enabled device, this capability can transform a very common business action – searching for the latest version of a slide deck or proposal – by searching for the colleague who most recently sent it to you.

Test this form of search with your smart phone and see what you get – or your Mac, where the latest version of the OS now supports Siri as well.

Imagine what else we could do as these assistants and chatbots are refined.

Then There's the Rub

While this simple example represents the potential for a tangible result, there's always "the rub" – the very real set of challenges and hard work that brings the potential from hype-driven promise to business reality.

In the example above, the data is structured. Emails have a send date and a field containing the sender. Underlying the promise of AI-driven search is the requirement that businesses organize their information in a way that enables machines to parse queries in relationship to the structure. This is where knowledge management enters the equation.

Tomorrow we will lay out the specific steps of this foundational work, providing a detailed explanation of the three major steps required to organize information so intelligent applications can consume it.

What it Takes to Deliver Successful AI-Driven Search

Artificial intelligence (AI) promises to deliver enterprises higher efficiency, increased accuracy and [greater utilization of corporate information assets](#).

But these promises can only come true if the AI is built on a solid information architecture.

While many would like to believe that AI is combination of magic and pixie dust, for human-to-machine conversations to become reliable, a significant amount of foundational effort must take place. And this process begins with the basics of knowledge management.

3 Steps to Establishing an AI Foundation

1. Develop a Framework

Domain Models and Schemas

Developing a framework starts with the concept of a "domain model."

A domain model represents the concepts and terminology of a particular industry or specialized area of knowledge. For example, a domain model for an insurance company might include products, services, risks, regions, topics, processes, audiences, content types, customer types, etc.

The domain model represents common terminology, structure and concepts for most business processes within the target domain. At a deeper level of detail, the domain model is converted into "schemas," which more precisely describe the process, data and context for the area of interest. The schemas become your metadata structures.

The business world contains thousands of schemas and contexts, and you cannot resolve all of them once. Begin the framework process by starting with one, something critical to the business. By narrowing the focus, it will allow you to measure the business impact more readily.

Writing a schema creates a structure which supports machine training. Machine learning models function more accurately when told what is important to your organization.

Search also benefits from domain models and schemas by limiting and contextualizing possible results to make them more meaningful to the user and relevant to the domain.

Digging Deep into Ontologies

Start by building a logical map of information related to a specific business process or need.

Take the proposal creation process of an enterprise as an example. Proposals in this setting contain many information elements. The challenge is finding the right mix of information which relates to the potential client's requirements and to deliver the completed proposal on time.

Businesses waste hours searching for the various content components, whether product configurations, specifications, information about people such as bios and photos, and other components such as case studies, project plans and cost elements.

Data structures, domain models and schemas (with taxonomies and controlled vocabularies) comprise the ontology – defined as the elements described and all of the relationships between them.

In the insurance domain, the relationship between the vocabulary of risks and that of regions can be called “risks in a region.” These ontological relationships capture knowledge about the organization, its processes and relationships in the real world and become the framework on which knowledge is organized.

Knowledge Engineering, at Scale

Machine learning algorithms can place knowledge elements and artifacts into a knowledge base, and then surface the content users need in the context of their goal.

A salesperson creating a proposal can retrieve the information needed to construct it because the process of completing a proposal is part of the framework. The schema structures the information for that process.

If this sounds like old fashioned knowledge engineering, it is. The difference is how the tools scale across content and data sources as well as interpret human intent.

Many of the elements humans need to learn how to complete a task are the same ones required by an AI-driven search application.

In fact, “training content” remains the biggest obstacle to enabling AI programs. IBM Watson's creators noted that – as one would expect – feeding the AI the wrong information sources degraded performance.

Over time, the machine learning that AI offers will move search to a proactive experience. Machines can watch/listen to what humans are trying to do, and refer to repositories of past search actions, improving its ability to return the most relevant results.

2. Build a Knowledge Management Foundation

Serving Up Search Results, Within Context

AI-enhanced search depends on a knowledge management foundation. Even when the algorithms perform without external ontologies, someone made these decisions and embedded the architecture in the code.

Knowledge engineering will play an even more important role in the success of emergent technologies such as bots and conversational interfaces.

AI tools do not remove the need for domain models, schemas, content architecture, ontologies and other design elements. To the contrary, context must be defined and retained via knowledge engineering approaches.

Knowledge engineering means you are building mechanisms that deliver search results containing the right content to support your employees' processes.

Taking this a step further, using component authoring enables content to be re-used at even finer levels of granularity and in more contexts. AI interprets human intent and contextualizes and personalizes results at a scale impossible with knowledge engineering alone.

Google Sets the Bar High

These precepts aren't new.

Google has been enhancing its contextualized search for a while, showing steady improvements over time. Traditionally, Google searches resulted in a long list of web pages which contained the answers buried inside.

Well-worded searches sometimes resulted in direct hits, where the answers were part of the URL's page title. This "trained" the world on how to tag web content for easier retrieval (and is why marketers fear the dreaded Google algorithm!).

You may have noticed, however, that Google has started giving direct answers when it can (e.g. search "How old is Drew Barrymore?" – the answer appears in a bio box astride the list of search results).

These responses use the Google "knowledge graph," a structure of related concepts and attributes, which is a form of schema related by an ontology. But progress has been slow. Even while possessing the largest database of queries known to man, Google still struggles to consistently offer this level of response.

On a positive note, you don't have to solve what Google is trying to solve. Your business's area of expertise is a tiny subset of the world's knowledge that Google tries to master.

Tying Information Architecture to Business Context

Unfortunately, when it comes to enterprise search, Google doesn't reach inside the firewall.

Much of corporate knowledge is stored safely within protected networks that commercial search engines cannot reach. Further, this information is often unstructured and lack the links that facilitate ranking and retrieval algorithms.

Corporate information often exists in untagged documents and files, outside of a meaningful metadata structure. In order for a machine to understand this unstructured content requires it first be classified and mapped, as well as placed within a context of its use in the business.

Many classification attempts fail because the information architecture does not account for the potential uses of content and information. Instead, businesses organize unstructured information in haphazard ways – either through an accidental architecture or by way of personal, idiosyncratic and inconsistent organizing approaches.

Haphazard approaches do not support user needs within their business processes. AI-enabled search is contextual by nature. Search is always tied to a step within a business process. Thinking about search this way changes how you approach information architecture design – basing it on the context of the business process.

For example, a pharmaceutical company is constantly in the process of complying with FDA regulations in order to advance its products to market. All of the related information that goes into this process can be architected, indexed, and classified to enable much faster access when searched. Advanced knowledge management

professionals are attacking this problem today, resulting in thousands of hours of labor reduction, and a faster time-to-market.

3. Train Machines Like a Human

If your company is venturing into AI-enabled search, don't start with machine learning.

Machines first need to be trained on how to learn – whether through schemas or through classifying and organizing information to a set of common purposes within a standard table-driven database. Training sets need to be the gold standard of content and represent clean data in order for them to be useful.

Depending on the purpose, the training sets may require varying levels of structure, but the data must always be of superior quality. Remember what Watson's creators taught us: Putting garbage into an AI system results in garbage AI results.

Consider how humans learn and apply the same approach to AI.

Humans do not learn about the entire business at once. A talented knowledge worker starts by learning about their part of the business. In the same way, a learning “AI machine” is far more successful when tasked with a focused set of business processes and clear objectives. Most of today's successful AI solutions are narrowly focused.

Feeding the system knowledge about the user, the task and the related content is “the learning” that must occur. It provides context. This process starts one department or one function at a time. Departmental and process-specific AI search applications can then cross over within industries, then cross to adjacent industries, and finally to “business” in general.

But the industry is not there yet.

Setting the right expectations is difficult. AI is not magic, but its value amplifies when you take the machine learning and enabling it in new environments. Each iteration results in refinement for all.

AI-Driven Enterprise Search Is Closer Than You Think

One of the biggest challenges of bringing artificial intelligence (AI) into enterprise search is simply getting it in the door.

The people building information management solutions have little interest in conversational applications. The inertia of ongoing projects (and of the new AI snake oil salespeople who promise the world) will siphon all available resources away from approaches requiring a [strong knowledge architecture foundation](#).

Because we have few examples of big successes in this area, executives can break out in a sweat when asked to back something lacking short-term ROI. At the same time, executives are justifiably concerned about being left behind.

Where's the ROI?

To move ahead, executives need to know the limitations and when they are engaging in pure research versus experimenting with innovative approaches and business models.

You will not realize the [benefits from AI-enabled search](#) immediately. Not only do the machines need to be trained, but the humans do as well.

It's not just about building a search bot, it's about learning how to interact with these systems. This learning process is a form of discovery, socialization and communication: users need to learn what they can do and this is not as intuitive as the Amazon Echo TV ads would have us believe.

Corporate information is complex, and knowledge work is nuanced. Work is hard, which is why they call it "work." The tools require some degree of training to make the most of them.

ROI metrics can be "hard" and measurable or "soft" and less tangible.

The former could include a significant impact on "time to access information" or "time to complete a process," given accurate baselines and metrics. However, companies rarely measure things like "time-to-information" or "employee efficiency" unless that efficiency is carefully tied to a specific process.

Knowledge workers' gains and losses related to time spent looking for information must be tied to adjacent performance measures. Time-to-market is one measure that ties to business efficiency – again, the key is linkage.

Some organizations see correlations between information system efficiencies and measures such as "rate of voluntary terminations," and "employee satisfaction." While these measures do not always come with hard dollars attached (though high turnover is costly to any business), most companies realize beating the competition to market requires retaining talent and corporate knowledge in order to create a competitive advantage.

Where's the Talent?

AI-enabled search applications requires new thinking and new teams. Your employees who are currently using traditional methods to solve the same business challenges are ideal candidates for this team, the trick is getting them to think in new ways around natural language conversation responses to queries.

Some of the talent required may come from outside the traditional technology department roles. People with backgrounds in the liberal arts, creative writers, authors, theater/performing arts all help here – these people intuitively know the conversational language that occurs across many walks of life. User experience designers interested in a new user interface paradigm – the Conversational User Interface (CUI) – can also be important contributors.

Crafting scripts and responses in conversational language is a new skill set. Is there a "skills gap" to overcome? Demand is higher than the supply of the conversational language programming skillset. Amazon, Microsoft, Google are all hiring storyboard writers right and left, from Pixar and similar companies.

Today's AI Is Tomorrow's 'Search'

We live in an "interrupt world." When we can't find the information we are looking for, we interrupt our colleagues in the hopes that someone knows where to find our data.

This not only interrupts the requester, but also those targeted with the requests, and the impact may domino through several layers of the organization. Large enterprises suffer the most, battling through scattered information repositories and dispersed resources.

AI-enabled search promises to put a dent in this problem. The hype engine promises more than that, but for the purposes of this discussion, let's be realistic and be happy with a dent.

Knowledge is the currency of business today. We've all heard the mantra: "Knowledge is power." But if you don't build knowledge bases related to the "secret sauce" for how you do business, you will not lead ... you will lose.

Deciding to rebuild an enterprise search paradigm means taking resources away from how things are done today. It takes strong leadership – maybe even a change agent hired to disrupt – to realize the status quo isn't working and act accordingly.

Finally, realize this: It's AI when it involves knowledge or information previously thought to be only in the realm of humans. But as soon as it's automated, we stop calling it AI. Spell check, grammar checks, then style suggestions were all once considered AI. Now we take them for granted as part of word processing.

So maybe you'll now ask Siri different types of questions or make travel arrangements conversationally. This is today's AI, but soon it will just be regular old mundane "search."