

# Making Intelligent Virtual Assistants a Reality

---



## What is an intelligent virtual assistant?

---

The ability to learn is part of the basic definition of “intelligent,” and that is an essential quality in an intelligent virtual assistant. “Intelligent” also implies an entity that has independent, accurate knowledge. An assistant is someone who helps someone carry out a particular type of work. “Virtual” of course means that the intelligent assistant does not exist in the physical world but lives in cyberspace. Putting it all together, an intelligent virtual assistant has knowledge and can also learn, is a digital entity, and helps someone accomplish a task.

According to the CTO of the company that developed Siri, before Apple acquired it, the number of situations in which people ask a mobile device a question is relatively small. He identified seventeen different possibilities. Typically, people are searching for something; it might be tickets to an event, or, places to eat, or directions to a destination. Once the context is defined, it is easier to provide help, because voice recognition becomes simpler when the context is known. The same holds for an intelligent virtual assistant. When the job and task context is definable, it is more feasible to design an application that achieves the goal.

## Example of an Intelligent Assistant [ABLe]

---



ABLe is a virtual assistant that EIS developed for Allstate Insurance to help Allstate agents answer questions about business insurance policies that they were writing. At the time that ABLe was launched, business insurance had been introduced as a new line of business to Allstate agents, who had been writing mostly personal insurance policies. The company wanted to encourage the agents to cross sell from personal lines to business lines. For example, when a customer who was a contractor came in to renew their homeowners insurance, the agent could talk with the contractor about commercial auto insurance, or insuring their tools within a business owner's policy. EIS helped build a system to help fill out complex underwriting forms. The ability to vary its state in response to different situations and past experience is one of the trademarks of an intelligent system.

## Curation and Automation

---

The approach used for ABLe included continuously curated content where learning can take place through both manual and automated processes. The system requires core content that supports specific user tasks. Browse paths and search history are examined on an ongoing basis to identify successful and unsuccessful responses and then curating content, search configuration and terminology to create the best results. This includes updating and adding terms to

the taxonomy and adjusting query processing. The same metrics used to manually tune the system can then be applied to automated techniques.

## IVA IA (Intelligent Virtual Assistant Information Architecture)

---

In order to develop the information architecture (IA) for a search driven virtual assistant it is essential to define that problem in detail as well as all of the contexts in which that problem will be approached. In the case of ABIE, a number of use cases were available, along with domain models for the different types of terms, and for the various contexts of the agent and the information they were seeking. Content models were available for different types of information, including frequently asked questions, reference materials, procedures, and manuals that included everything the agents needed to know about business insurance products.

## Context is Key

---

The next step was to break that information into pieces and to contextualize it. As a result, when the agents hit the help key in the application that they were using to write a policy or a quote, or to issue a business insurance policy, the context was clear, their identity was known, and the right information was available. The 270 contexts in the workflow were defined, so that when a piece of granular information was delivered it just right for that context.

The solution at Allstate was based on a well-understood problem and used conventional technology. The task was clear, build a search-based application with structured content. This made the outcome manageable and successful. So much so, that when agents make long inquiries, the system is so responsive that agents may think a person is answering their questions. There may come a point when agents need to chat with a real person, and at that point they are handed over to a live support resource.

Over time, very specific inquiries are built into a question and answer system, similar to IBM's Watson. All of the information that was gathered in the implementation of ABLe, including chat logs, search logs, and other sources become training sets for the system.

These manual, curated and architected approaches may not seem like the magic that many virtual assistant vendors are claiming, however, that is the point – there is no magic. The foundation for these tools requires curation, well structured content and information architecture. They belong to a class of search-based application that also leverage advanced techniques around query processing and task contextualization. This forms a foundation for additional machine learning technologies which consume responses and live agent chat dialog to improve performance.

## Framework for a Virtual Assistant

---

The framework that builds on technologies that culminate in an intelligent virtual assistant is a cumulative model (see Figure 1). Many decisions need to be made around the application of technology, the application of content strategies, and the user experience interactions to address a specific set of cases or business problems. The model shows a way of classifying the space that encompasses these technologies, to help provide a common view and understanding of what constitutes an intelligent virtual assistant.

	BASIC SEARCH ENGINE	KNOWLEDGE PORTAL	VIRTUAL AGENT	INTELLIGENT ASSISTANT
KNOWLEDGE BASE	Any text Multiple sources	Multiple sources, separate Ontologies and schemas	Domain specific Highly curated sources	Dynamic. Info enrichment improves with interaction
SEARCH INTERACTION	Keyword or full text query	Full text query or Faceted Exploration	Query, explore facets Offers related info	Implicit query / Recommends based on users' history
INFORMATION ARCHITECTURE	None necessary, but Improves with metadata	Ontologies, clustering, classification	Ontologies, clustering, classification, NLP	Ontologies, clustering, classification. NLP, personalization
USER EXPERIENCE	Search box, documents list	Role-Based	Conversational	Conversational, personalized, contextual
ENABLING TECHNOLOGY	Search	Search, classification, databases	NLP, search, classification Process engines	NLP, search, classification Machine Learning

Figure 1. Search to Intelligent Assistant Continuum

## Basic Search

Basic search, shown at the left side of the figure, is a web crawler that goes through many text sources. The search interactions are typically keywords, so the interaction is through a search box. The information architecture is often the Wild West, with little explicit design. Information architecture, metadata, and well-formed content are all signals used in search, so every investment made in those areas improves the quality of the search.


## Knowledge Portal

A knowledge portal, shown in the next column has multiple sources and more knowledge about the user context, so schemas and ontologies can be more fully leveraged. Additional capabilities such as facets or browsing become possible to surface the content and to exploit the underlying information architecture. The knowledge portal would also be aware of roles, which provides context that helps to surface content that is more relevant to the users. Many employee portals and intranets exploit knowledge portal approaches.

## Virtual Agent

---

The virtual agent, in the next column to the right of knowledge portals on Figure 1, is very domain specific, task oriented, contextualized, with highly curated content sources for user tasks. Information architecture is a very structured set of knowledge representations, classification, with natural language algorithms to interpret term and phrase variations. Interactions are structured with styles tailored to content types and more granular contexts. The virtual agent should know enough about what the user is trying to do to present the information in actionable form, such as movie times and information to buy tickets for example.



## Intelligent Assistant

---

On the far right of Figure 1 is the Intelligent Assistant. This stage is characterized by dynamic interactions and learning from interactions to improve the knowledge base. The assistant might also make proactive recommendations that amount to responding to implicit queries, drawing from data in contextual cues. The intelligent assistant adds rules for an inference engine to derive responses that were not explicitly programmed. The interaction may be quite conversational. The user might get an answer and then refine the question, but is in a more natural situation versus a search.

The information architecture at this stage reflects very rich, well-defined content along with personalizations. It might draw from big data sources, aggregation of various feeds, and integrated with rich enterprise data sources. These are essential in providing the level of precision and proactive recommendation that characterize intelligent assistants.

## Humans Working with Computers

---

Both humans and computers have roles in developing intelligent virtual assistants. People excel at identifying the best use cases, defining the meaningful interactions with the data, and creating the initial context and examples; i.e., the so-called “seed data.”

Machines can leverage this seed data in a phase of supervised learning, and can auto-classify new data. Machines can also mine analytics and incorporate new knowledge so achieve unsupervised learning and improve their accuracy.

## Conflicting Results

---

The systems that support intelligent virtual assistants may work in collaboration with one another but sometimes they conflict. For example, if a user asks how long it will take to get to downtown Chicago from one of its suburbs, the system does a great job of getting contextual information, combining evidence about the user from processing the query, disambiguating the question and inferring that the user would be driving. But then search may begin reporting that there are five cheap flights to Chicago for \$48. This response totally ignores what the user was trying to do. Once multiple technologies are working together, it is important not to lose the user by having one system ignore information that was being successfully used by another.

## Machine learning, predictive analytics, and algorithms

---

Machine learning breaks down into a couple of simple concepts. One is algorithms, which are programs that represent steps to solving a particular problem. Some algorithms can learn from the data and make predictions based on the data. That is what machine learning is about. There are many examples of this kind of prediction. Spell check, for example, is a prediction that a user intends a certain word based on a pattern of letters and when it is not quite sure, the computer presents an alternative. Voice to text, handwritten character recognition by the postal service, translations, and autonomous vehicles all use learning algorithms.

Recommendation engines access information, leveraging different types of principles around search. Successful recommendations understand the user's needs based on the signals about identity or goals. With an intelligent agent, an application provides not just a

list of documents, which is what basic search would do, but a specific answer to a question. The algorithms make the process more powerful, although they still need to be taught. You have to teach them a certain amount. Content is now being optimized for virtual digital assistants. Marketers are wondering whether this technology will replace search engines, and the answer is yes. Eventually, personal agents will know so much about the user that they will start proactively finding information. Looking farther into the future, virtual assistants are likely to become the testbeds for tomorrow's personal robots.

## Understanding Foundational Content

---

Analysts agree that these advances will require a thoughtful view of content and proficiency in other facets of the technology, such as an expertly developed content architecture, domain model and knowledge base representation of the information. Intelligent virtual assistants need to be focused within a narrow context, and therefore benefit from clearly defined use cases and codified human expertise.

The technology is cutting edge however leverages foundational principles of content and structure. The first step is to define the problem to be solved, the information needed to solve that problem and the steps to the process.





## ABOUT US

---

### **EARLEY INFORMATION SCIENCE (EIS) HELPS ORGANIZATIONS HARNESS THE VALUE OF ONE OF THEIR MOST IMPORTANT ASSETS – INFORMATION.**

We improve enterprise information agility to drive digital commerce innovation, increase revenue, and enhance customer experience with a strong enterprise information architecture and management foundation. EIS helps make information more findable, valuable and useable — thereby helping organizations improve operational efficiency and effectiveness. EIS has over 20 years of experience in working with Fortune 1000 organizations globally across many industries, including manufacturing, retail, financial services, healthcare, life sciences, the public sector, and professional services.

Earley Information Science  
PO Box 292  
Carlisle, MA 01741  
P: 781-444-0287  
[www.earley.com](http://www.earley.com)