

An Interactive Intelligent Advisor

Roger Schank, Ray Bareiss, Chris Riesbeck, Wendy Lehnert

The incredible amount of hype that Artificial Intelligence is currently receiving is mystifying to those of us who have spent our lives as AI researchers. “Novel written by AI wins contest in Japan” is the headline you see until later you find out it’s not quite true. Or “AI wins at Go” until later you find out that this was just a computationally intensive neural net which is quite different than true artificial intelligence. Or “Watson wins Jeopardy” until later you find out that Watson can barely understand any English sentences.

The main issue in AI has always been the same: There are those who think that any program that beats a chess or Go master, no matter how, is “intelligent,” and there are others who think that being able to do massive computation does not have very much to do with intelligence. The latter group, to which we belong, would like to build systems that really do display some intelligence. What would that look like? For one thing intelligent entities can explain what they did and why they did it. For another, intelligent entities can communicate through natural language and actually can at least act like they understand what they are being told. Intelligent four year olds don’t understand everything you say to them, but they do have *goals* of their own, *plans* to achieve those goals, *ways of explaining* their actions, and *ways of understanding* what you might want from them. Let’s take this as our basis for defining what a real AI system should be.

Now let’s ask the question: *What can we build now that would be both intelligent and useful?*

Whatever we build must have three distinct capabilities that correspond to three integral parts of human intelligence. These capabilities are:

1. Natural language comprehension
2. A model of the world
3. An appropriately indexed memory of experience and expertise.

Let’s talk about each in a bit more detail.

1. Natural Language Comprehension

AI researchers have worked on this problem for many years – because it is hard. People who have taken the problem seriously have come to the conclusion that natural language comprehension by an AI system is possible if the domain is highly constrained. Constrained domains (e.g., talking only about moving blocks or understanding texts about terrorism) enabled good progress more than 30 years ago (when AI was an honest field). On the other hand, trying to build an AI that you can say anything to leads to systems like Siri, Cortana,

Alexa, and Watson that really have no idea what is said to them but can make statistically-derived guesses that sometimes result in a useful response.

2. A Model of Its World

If the world we are talking about is highly constrained, it is possible for a system to learn and employ a model as a basis for understanding the domain's nuances and people's goal-directed behaviors within it. For example if we constrain the world to "air travel," it would be possible for a system to know a great deal about that domain and to use that knowledge to help a person to optimally plan a trip in line with their needs and preferences. The user is also a key part of the system's world. To optimally satisfy the user's needs, the system must also have a model of the user – his or her level of expertise in the domain, previous experiences, preferences, and so on.

3. An Appropriately Indexed Memory of Experience and Expertise

When we have a problem, we want help, preferably from an expert. Just like we sometimes needed our parents to give us advice, we need someone who knows about Iceland to tell us about where to stay and what to see if we are planning a trip. If we want to learn to program we need an expert to tell us when we make a mistake and to offer helpful hints. If we are trying to choose a school to attend, we would like to hear from people who have gone there and liked it, and others who have hated it. We would like to hear from both experts and people like us who can help us think about making better choices. Therefore, an intelligent system needs a memory of advice that is indexed in a way that enables the system to retrieve a piece of advice when it is likely to be relevant to what the user is doing.

With these capabilities in mind, *we propose to build Interactive Intelligent Advisor that could work in one domain and one domain only (although we could use the same underlying architecture to build Advisors in many domains).*

Here are some examples:

- *An Intelligent Interactive Mentor* that could teach you something in the context of a learn-by-doing course, say, a course to learn how to program.
- *An Intelligent Trip Planner* that could build a model of you and of travel options (as any good Travel Agent does), and based on that model, provide personalized, expert recommendations.
- *An Intelligent Home Buying Advisor* that would interact with a user to understand his or her wants, needs, and finances to provide personalized assistance in the home-buying process.
- *An Intelligent Car Shopping Advisor* that can help you to choose the make and model that's right for you, priced so you can afford it.
- *An Intelligent Coding Assistant* that can speed up the software development process while helping a programmer to avoid the mistakes he or she typically makes.

(The appendix to this document sketches several example interactions.)

All of these systems (and many more like them) would:

- observe a user as he or she works online – tracking both what the user is doing and parsing user-entered information
- proactively offer help, advice, and feedback as appropriate, drawing from a memory of expert and peer experience
- query the user when his or her preferences or actions are ambiguous
- answer questions, especially in response to system-provided advice
- expand its knowledge base by posing questions it cannot answer to external human experts and automatically adding their responses to its memory
- learn the user's specific preferences and tailor the interaction accordingly.

Such an *Interactive Intelligent Advisor* would be a “real” AI system that could learn who its user is and what they are doing, and respond appropriately in light of the user's needs, actions, and preferences.

All of these Advisors will be based on a general AI architecture comprising mechanisms for comprehending natural language, maintaining a world model, and accessing expert advice and experience. This architecture can be instantiated to create specific Advisors for a range of domains. Through use, each advisor learn more about its users and its domain, improving the quality of help and advice it provides.

We would welcome the opportunity to tell you more about our vision and our experience.