

Conversation Agent AI

Blog 3 of 3

Andrew Lea, FBCS, FRSA,
Visiting Research Fellow at the University of Brighton
Head of AI at HelloDone
andrew.lea@hellodone.ai

In previous blogs we looked firstly at the benefits of Conversation Agents, and then the issues they must address and how they could be architected, but the astute and erudite amongst our readers will have noticed that we dodged the actual question of how to make a Conversation Agent understand what it was reading, take sensible actions, and generate engaging replies.

In short, how do we make Conversation Agents actually intelligent? Enter AI. We aren't going to give you the full recipe (after all, we want you to procure the highly cost-effective Cicero!), but we will tell you the ingredients. Well, mostly.

A bit about Artificial Intelligence

Unless you live on Mars, you won't have escaped noticing that "Artificial Intelligence" is all the marketing rage at the moment, and may even think that "Machine Learning" or the distinctly non-eponymous "Deep Learning" is all there is to AI.

Leaving aside both the murky depths of **Fake Marketing AI** and the esoteric heights of **Real AI** which hopes to make computers "think like people" with all that entails, not the least being self-awareness; we can view practical **Applied AI** as the "intelligent acquisition, manipulation, and exploitation of knowledge". Machine learning can be one source of that knowledge, but so is the systematic acquisition and codification of high grade knowledge into a knowledge base. You will be unsurprised to learn we use both.

Another view of Applied AI is using a computer to accomplish that which, *were* it to be done by a computer, would require intelligence.

Engineering AI into Conversation Agents

Most practical applications of AI, such as a conversation agent, use multiple AI techniques. After all, that is what we do as people. The requirement for AI is conceptually simple: we need the agent to be intelligent, as if the user were conversing with a person.

As AI was developed, it became apparent that the tasks we thought would be easy, because we find them easy - understanding language, recognising objects, running and jumping - turned out to be hard. The reason is simple: the brain devotes a large proportion of massively parallel processing to handling these tasks in our sub-conscious. We are simply unaware that we are doing them at all. Conversely the tasks we thought would be hard, because we find them hard - playing chess, solving mathematical problems, thinking through complex problems - turned out to be easy, or at least no harder.

The other elements of human intelligence, especially self-awareness and sentience are almost completely beyond us¹.

¹ In fact so much so that they are now taboo subjects, and very few of us, with the author being one of the exceptions, are trying to develop self-aware AI. Some people think we shouldn't, because AI may turn out to be an existential threat. The good news, however, is that AI won't be able to finish off the human race, because climate change will do so first.

Our Conversation Agent has some straight forward tasks - logging errors, talking to external systems, reformatting output - which are merely “work”, requiring nothing more than lots of high quality software engineering.

Many of its tasks, though, are more interesting (and which cultured human being isn't deeply fascinated by AI?) since by its very nature the task of behaving like an intelligent person requires the mimicking of human thought - applied artificial intelligence.

A Cornucopia of AI

Here are some of the AI components and technologies we have engineered, in once case subscribed to as a service, and in another rejected:-

Voice to Text - ie listening to sound and extracting the text of any utterance - is now commodotised. It is, however, extremely hard and many millions were invested in cracking it. Although modern smart-phones are easily powerful enough to run voice-to-text, the major vendors still prefer to convert the sound to text on their computers “in the cloud”². One can only speculate as to their motivations.

Artificial Neural Nets for concept extraction. Artificial Neural Nets are inspired by the concept of the neurones that are linked together in our wetware brains, and are nothing more than approximation functions³. They are, however, enormously powerful approximation functions which can also generalise, ie find the pattern in messy examples. Essentially we train Neutral Nets with example text and the concepts they contain, so that they can recognise the same concepts when they occur in other text, such as a message from a user. This is an example of *supervised learning*.

Semantic Representation. It is all very well saying “extract the concepts”, but how do we represent those concepts? There are of course several possibilities, including semantic “graphs” (dot-to-dot diagrams where the dots are concepts which are linked to others), sets, fuzzy sets (where a concept may be a slightly present), and so forth. In our world, the semantic representation is based on the concepts of:

- “intents” - what is trying to be achieved
- “frames and slots” - small snippets of bundled facts with known structures
-

We use this as the “lingua franca” within Cicero, so that one component (for example the intent recogniser) may talk to others (such as handlers or the FAQ system), who may indeed reply.

It is possible to train a computer in the **Machine Translation** of human languages, such as English to French⁴. In theory we could have our entire system run in, say, English, and translate incoming language into English, process it in English, and translate the output back. We don't do this for four reasons:-

1. Our system already has a more tractable representation than English, so its not necessary.
2. We can train the Neural Net based intent recogniser to work in other languages too.

² There is no cloud. There is only other people's servers.

³ Please do *not* mention this when speaking with a Neural Net practitioner, as it tends to make them sad.

⁴ Your author once programmed a system to learn, by reading examples, how to translate from one language to another. Although the program took weeks to write, depressingly it only took the system two minutes to learn how to translate English to French.

3. Translating English to another may yield intelligible text, but it *will* miss the cultural aspects of talking in another language.
4. We have a more powerful method of generating language than translation...

Our **Natural Language Generator** can take details, and a knowledge base with a generative grammars and generate maximally informative, varied, even interchangeable personality, and culturally appropriate text. It is based on the system that would have been on the Mars lander that crashed.

Naturally people write with spelling mistakes, typos, abbreviations, and even emojis, all of which Conversation Agents must be able to understand. Internally, then, we use various **Fuzzy Matching** techniques, including edit-distance, n-grams, dictionary methods, soundex-like transformations, and so-forth. As with most AI techniques, the appropriate fuzzy-matching technique depends upon the context. One last wrinkle: fuzzy matching is improved with domain knowledge, which we often use to increase recognition rates.

On many occasions we need to deduce what to do next: enter **Inferencing**. Inferencing is a bit like programming - if this then that is implied - but where the underlying system (or inference engine) determines which rule to apply next. This is a very effective approach where the sequence of events is unknown, which is certainly the case when speaking with people who may express the same set of ideas in many orders.

Reasoning with *Uncertainty*: **Fuzzy Logic**. Often we are faced with the situation where a particular fact may be true (this is a returning customer), false (this is a new customer), or unknown (may or may not be a returning customer). To none the less reason on the situation, perhaps using inferring rules, we use a type of belief-based logic akin to Bayesian logic.

We also use **Unsupervised Learning** and **Clustering** for purposes of continuous improvement.

There are other techniques we are developing, but which we do not wish to disclose at this point.

Make and Go Mad or Buy and Relax

So if you want to build a robust real-world capable Conversation Agent, these blogs will have give you your to-do list.

Less stressful and more successful approaches include:-

- Getting HelloDone to integrate their service into yours.
- Using the HelloDone framework to coral your multiple corporate ChatBots into order.

If you would like help, drop us an email, and either us or Cicero will get back to you.