

Artificial intelligence

Artificial intelligence has been a goal since the beginning of computers. Why is it that after over 50 years of work there are still no systems out there that can be said to be intelligent? Computer power has increased by massive amounts, millions has been spent on research but still no major results. This article talks about the current batch of chatbots and where AI should be going.

I should state at this point that I believe intelligence is a set of processes in the brain that is perfectly possible to emulate in software so that a true artificial mind can be produced. I have been interested in Artificial intelligence almost since I first got into computers over 30 years ago. The promise of making computers smart, being able to talk to them seemed like a tantalizingly close dream. How hard could it be? The fame and fortune that would follow the person who first worked it out would be epic. I was hooked.

30 years later and I am still interested but the dream is just as far away now as it was then. What happened?

In a nutshell, nothing much.

Researches worked on the problem, systems were built but no real results, no systems then or now produce any sign of real intelligence. There are a number of systems (so called chatbots) which are designed to look like intelligence. I guess the most famous is the chatbot called ELIZA which was designed to act like a computer psychologist. It was simple and had a stock set of answers to certain keywords typed by the user. Obviously the results were patchy, it fooled some people but very quickly showed up as just a toy. There have been a number of systems since then that have tried to continue this line of work. I have designed a few myself, they scan the user input to try to get a sense of what the user is saying and then use a set of prepared answers to give the impression of a real person giving real answers. These systems have real limitations and are still no where near what is needed.

There are a number of other systems based on neural nets that are designed to emulate how the human brain works (a series of cells interlinked). Since they are designed to work in similar ways to the way the brain works it is a valid approach. Again after decades of work not much of any use has been produced. This approach is similar to trying to work out how a computer works by studying a transistor. It may get some results eventually but it is going to be very difficult.

So why hasn't there been progression over the years like you get in fields of, study? The answer may seem obvious but it is too complex. If you have billions of cells in your brain each connected to tens of thousands of others, it become impossible to figure out what is going on. Add to this the difficulty in direct experiment on living brains and it's not hard to see why we are still getting nowhere.

So what do we do? How can we progress after all this time. It is my assumption that if you mimic the way people behave then we can build intelligent systems, we should copy the way our minds work on a system level not on the level of the neuron. If we study what people do, what drives them then we will have an understanding of what is needed to make a system look smart.

So where do we start?

Computers know nothing about the real world. People take years to develop an understanding of the world; this "assumed knowledge" is taken for granted but is a vital part of what makes us what we are. Even a child knows a bird can fly, flames are hot and knives are sharp and can hurt. Computer systems know nothing about any of this, they have no assumed knowledge.

If a computer system is to be intelligent it must have assumed knowledge, it must either have ways to accumulate this knowledge (camera's and microphones) or we must give the knowledge to the computer. Chatbots or any AI system must have "assumed knowledge". Ideally a system should learn this assumed knowledge but the technology to do this is another whole subject on its own. So to produce a system in a reasonable time we must give the knowledge to the system at the start.

Knowledge is not everything. Knowing stuff does not make you intelligent, using those facts to your advantage is a sign of intelligence. Assessing advantage is the next key to intelligence. Even simple creatures discover techniques and new methods to get food, fight enemies and stay warm, they assess what happens and form opinions on those events designed to give advantage. Any AI system must be able to assess advantage to it and therefore must have goals. You cannot access an advantage if you don't know what you want/need.

Mood.

All intelligent systems have moods (emotion), sometimes they are angry, happy, sad, violent, and passive. These mood varieties are designed to aid the individual to survive but they also have a direct effect on intelligence, the mood you are in directly effects the way you think, work things out and assess advantage.

In summary all intelligent systems must have:

- Priority (Goals) - Advantage Assessment - Moods (Emotion) + - Assumed Knowledge

I will shorten this list to p.a.m. + a.k

In the next chapter, I will discuss each of these elements in more detail and how we might go about building this into a software program.

As a recap, I believe an intelligent system must have:

Priority (Goals). Advantage Assessment. Moods (Emotion). Assumed Knowledge.

If we can implement these features into a software program then the intelligence it produces would be a real step in the right direction.

Assumed knowledge. What is the capital of France? Where do birds live? What is cold?

These may sound like silly questions but they are of real importance. Any AI system MUST KNOW the answers to these questions, but how is this to be done, what does "know" even mean?

A dictionary definition of know. To perceive directly; grasp in the mind with clarity or certainty. To regard as true beyond doubt. To have a practical understanding of, as through experience; be skilled in To have fixed in the mind. To have experience of. To perceive as familiar; recognize. To be acquainted with. To be able to distinguish; recognize as distinct. To discern the character or nature of. To possess knowledge, understanding, or information. To be cognizant or aware.

It can be seen that "know" means "store information in the brain". It may be a smell, an image or a feeling, but they are all just storage of sensory input.

This brings me to an interesting point. As we have seen knowledge is an important factor in intelligence, therefore gathering that knowledge is essential: Any AI system must have senses, it must have a way of gathering information from the outside world.

This means that eyes, ears and smell etc. are an essential element in intelligence. It is not surprising that AI has proven so difficult; to do it properly the system requires senses. What can we do about this? Obviously we can give our systems eyes and ears but this is a complex problem in its own right and could takes years just to get something working. So how can we give our system a means of gathering data? There are two choices here; we either give it information or we allow our system to interact with other "systems". By systems I mean either people or the internet. In the past the only choices were to give the test systems data or allow interaction with other people via a terminal. Now we have a new and much better option, we can allow our system access to the internet directly; it can then scan web content and gather information from there. This is a vastly superior method as the internet has video, sound as well as text data and would allow a system to get data from multiple sources and compare for accuracy.

In the first instance we give a system a set of data from dictionaries; encyclopedias etc and then allow it to explore the internet at its leisure.

Since we won't have proper senses to begin with I suggest that we just give our system a comprehensive set of dictionaries and encyclopedia's etc to give it a starting point. This data will have to be organized so that it can be searched quickly, allow easy editing and additions and be multimedia so that we can store pictures, sounds and video. Obviously we need a full text indexed database. Actually we need more than this in the longer term since we will want to search for pictures based on picture input and sound based on sound fragments but a full text index is a start. How this data is stored and retrieved will be talked about in more detail in a later article.

Advantage and Priority

In the morning why do you get up? You have to be at work by 9am or you have to be at school, meet a friend etc. From the minute you wake up to the minute you fall asleep you have a list of things to do in your head. A list of people to see, places to be. These are your priorities and they change daily. On a basic level these things drive you on, they give you a purpose. What is really happening is that you are seeking advantage. You wake up and get in the shower, being clean is better for you and everyone else. You clean your teeth (so they look good and don't fall out), you brush your hair (to look good). You dress (to look good and for modesty) and then head down stairs for breakfast (you must eat to live). All of these actions are taken to achieve advantage to you. Everything you do is based on advantage. Looking for a new car, house job. Even when doing things for other people, it is still all about advantage to you. Those actions make you feel better about yourself and tell others you are nice so they will be inclined to help you if you ask.

My point in all this is simple. Any Artificial Intelligence must have priorities, pursue advantage and be able to assess advantage by comparing the world against its own priority lists. Obviously avoiding disadvantage like hunger, pain getting fired from your job is the opposite end of the same system.

Advantage is therefore defined as any result that achieves or approaches one or more of the system's priorities. A priority is defined as obtaining any object or state the system has decided it wants.

How is this to be done.

There must be an ordered set of priorities and a feedback loop where actions are taken, results are observed and tests performed to see if a priority has been reached.

Mood.

As stated before, emotions are an essential part of us all. They encourage us to do things that give us pleasure, avoid things we don't like. Be afraid of dangerous things that will harm us and get bored and venture out to discover new lands. Emotions are designed to aid survival. Any Artificial system should have basic emotional states built in. This will make the system more lifelike and aid in its ability to gain new experiences.

When you think about it, emotion modifies our behaviour. Sometimes on a subconscious level sometimes not. Some emotional states are based on outside events (A Tiger appears inducing terror). Whilst some are more cyclical. It would be a simple thing to write software that increases boredom behaviour the more something is repeated. To increase excitable behaviour when a priority is achieved. Display fear behaviour when a threat shows up.

My point is that emotion is essential to survival and simple behaviour modification routines can be added to achieve it.

So far I have talked about the A.I system in terms of what we should have it do and how it should behave, actually building something that can do all this is another matter. Since audio and visual elements are too complex to deal with yet I will stick to building a software system that observes and communicates with the outside world using just the English language. This will limit its capabilities but vastly simplifies the steps involved to make it work.

Look out for the next article that will outline how we can achieve the above in software.

Daniel Burke
www.danielburke.com
May 2nd 2011