

Wanted - Smart Humans

Posted At : February 28, 2017 1:38 PM | Posted By : Admin

Related Categories: Inspiration, General Info

Remember PDAs (Personal Digital Assistants) - All that technology and a lot more are in our smartphones and smartwatches. It's likely that those devices will become the basis for AIA (Artificial Intelligence Assistants).

If you canvas what is being covered by the media and statements made by technology vendors, you might come away with the notion that we're all just years—or even months—away from being able to have our own personal artificial intelligence (AI) assistants, and businesses will be able to enjoy off-the-shelf AIs that they can train to perform various complex functions.

It is seductive to imagine downloading open source AI software like Google's Tensorflow or Microsoft's CNTK and training it to perform trend analysis on sales over the last 2 years to identify your best customers. Or, use it to pore over customer service interactions and online forums to understand customer satisfaction.

The reality is that AI is still heavily-reliant upon smart, willing and trained humans in order for AI to behave in a manner that we would expect. Humans are needed to scope the problems, identify relevant examples and verify the results. Without humans as a guide, current AI is no more capable than a computer without software.

All are not equal

AI is an umbrella term that means different things to different people. AI is not a single technology. Deciding upon which to use when can be complex. Without proper scoping of a particular problem, AI cannot deliver relevant output.


The ultimate promise of AI and learning machines is that they will “live” in our world and develop general intelligence in the same ways humans do. This conjures images of robots that can perform a variety of tasks from mundane house cleaning to more-sophisticated activities like creating and balancing a budget for the family. This is known as “artificial general intelligence.” The reality is (at least for the near future) that just as we have developed specialization for our own jobs, AI also needs to be specialized. We call this “applied AI” where the machine is taught a very specific problem-solution domain.

Scoping the problem, selecting the technology

When it comes to tackling a particular problem, there are many different machine learning techniques. While “deep learning” neural networks are the most popular, there are always specific machine learning technologies suitable for distinct tasks. Most of us would not understand the technical differences (let alone the strengths and weakness of) convoluted neural networks versus its recurrent variant or if support vector machines perform better at classification than deep learning alternatives.

Understanding which AI to use requires an intimate knowledge of the mathematical underpinnings of each program. And sometimes, it makes sense to combine several different technologies to produce the best results. Getting from A to Z is even more complex since the user will have to understand how to properly synthesize results from several different machine learning programs. This “voting” technique can be very complex and based upon independent variables of each machine learning program.

Garbage in, garbage out (GIGO)



As if selecting the proper technology wasn't difficult enough, just as difficult a task is scoping the real problem, inclusive of feeding the right amount of relevant data. Most of us have heard the axiom "garbage in, garbage out" (GIGO). And this is true for learning machines, and many times it can be impossible to understand if the inputs are garbage.

For instance, let's take a deep learning convolutional deep neural network (CNN) to identify the most profitable customers of a company. Even with the defined goal of "find the most profitable customers of 2016", the AI wouldn't have any idea of where to start with regards to data. Pointing the AI at the company's accounting database is highly unlikely to result in any output, good or bad. The AI needs to be trained to complete a specific task using relevant and unbiased data. Relevant data also cannot be missing or incorrectly weighted. What if key data required to make the calculation of profitability, such as cost of sales is not available? The output would obviously be erroneous.

Training for inference-building

The other part of training is that AI must also be provided various outcomes so that it can develop inferences as to what scenarios produce relevant results. This is typically referred to as "ground truth" data that is critical for machine learning software so the output is as precise as possible.

In the profitability example, a trained data scientist would research the problem domain, identify all relevant inputs for customer data and then assign profitability scores for each as outcomes. This data is then submitted and the software analyzes each along with the outcomes to develop inferences. The scientist must then examine the outcomes and determine if adjustments need to be made; either to inputs, the model or both. Input may be missing key causal data. This process goes on until the software outputs the results that are deemed acceptable.

In the future

At some point in the future, trainable machines will offer end users the ability to refine models or even train new ones, but that day is still a ways off. AI software will have to undergo more refinement to get to the level of simplicity that can support less-technical people using it. For today's reality, AI is best focused on constrained, specific tasks after being trained by experienced data scientists who can properly teach and test it.