PATRICK HEBRON

An Interview

A software developer, designer, teacher, and author, Patrick Hebron leads the Machine Intelligence Design team at Adobe. His work focuses on the emerging intersections between machine learning, design tools, programming languages, and operating systems.

What kind of relationship will unfold between humans and machines?

The path forward is for us to celebrate difference. The world is full of human thinkers. If we want human thinking, we should probably go to humans for it. There are a lot of them. I don't see the sense in trying to replicate the human perceptual system. One of the interesting things about artificial intelligence is the prospect of having another intelligent species on this planet. This is of real value because it holds a mirror up to the nature of our intelligence and to intelligence in general. It also brings a new form of intelligence into the world, which can be constructed as a harmonious counterpart to our intelligence.

DeepMind's work on gameplaying has been particularly inspirational to me. In one of its games against the human champion, Lee Sedol, AlphaGo played a now famous move that the commentators initially thought was a mistake. Ultimately they realized that this was a brilliant move—a move that defied three thousand years of human strategic wisdom. It discovered this novel strategy through its own intuition rather than adhering to someone else's preconceived notions of how to play the game. There's a real value in that. After the match against Lee Sedol, DeepMind built a new version of its system that only learned from its own gameplay and had no contact whatsoever with any human-played games. This new version was substantially better than the previous. It's amazing what a fresh perspective can do!

Though Go is just a game, we can start to leverage this kind of strength in other things too. To be clear, this strength is not a matter of raw intellectual horsepower. It's more a matter of how we can commit the machine's intellect to a particular domain or task. We are temporally confined by our lifespans, but can build ideas across numerous lifespans by communicating with one another, by writing ideas down and sharing them. Machines also have limited lifespans, of course. But they can approach communication somewhat differently than us. When humans communicate ideas to one another, they start from some sort of mental representation and distill it into concrete language, which is then interpreted by the listener back into a new mental representation. A lot can be lost in translation. A lot of the nuance can be lost to the low fidelity of language.

"How can we bridge the perceptual gap? Humans and machines are both perceptual entities. But their mode of perception is fairly different from one another, regardless of raw horsepower."—Patrick Hebron, Adobe

With machines, we can set up thousands of learners to interact with a given task or material and discover a wide variety of unique approaches. Ultimately though, we can take each of those machines' mental representations and merge them together without having to pass through a low-fidelity language. So a thousand separate Go players don't need to result in a thousand separate Go players, they can result in one extremely intelligent one—a player that has benefited from the integration of many different strategic approaches. Of course our approach to knowledge transfer also has its strengths. The loss of fidelity through language may be a good thing in some ways—a kind of one step backwards, two steps forward kind of progress that leads to a constant pruning and reorienting of human knowledge over the ages. In any case, having different sorts of intellectual strength in the world is a good thing. As we bring these technologies into our lives, we need to think about how we situate them in relation to what we each are good at.

When it comes to design, there are always going to be factors that will have to come from humans. Autodesk has done some interesting work on constraint satisfaction. In this framework rather than designing the specifics of a bicycle, the human instead appoints particular attributes they would like the system to include—for example, they want a bike to be both fast and lightweight. The machine then tries many different possibilities, running each through a physics simulation to determine which best expresses the desired combination of attributes. This kind of system can be quite successful in meeting its stated goals. But, of course, in bicycle design, there are many important considerations beyond performance characteristics. Is the design a good fit for the human physiology? Will a human know how to sit on it? If the bicycle is too different from anything that the user has encountered before, then the bike may not be very useful. Human-machine symbiosis means balancing the machine's ability to open up possibility spaces with the human designer's ability to mitigate that newness against historical legacy and issues of familiarity.

In my work at Adobe, the focus is somewhat different because objective optimization criteria, such as a bike's speed and weight, are generally less applicable in the domain of graphic design and image editing. In this context, the machine can also help to make the possibility space more accessible to and navigable by users. But it is clear from the start that a human must guide the process towards an outcome that will be deemed pleasing or useful by a human audience.

How will machine learning influence design practice in the future?

People have the initial reaction of, "Is this going to replace designers?" In some respects and to some extent, the answer is likely to be yes. Any massive paradigm shift is going to have uncomfortable growing pains. There's no denying that. But from another angle, this paradigm shift is not unique in its disruption to the organization and economies of how work is performed. For example, in architecture, the Florence Cathedral took about one hundred and forty years to go from initial conception to project completion. A much more complicated and recent building, the Burj Khalifa, took about five years. One of the key differences between those two projects is the advent of computer-aided design, or CAD, tools. These tools make it possible to conceive of systems that are too grand and complex for any one individual to keep all of their big picture goals and specific details in mind simultaneously. The importance of being able to scaffold complexity cannot be understated. If you look at the early versions of Adobe's Photoshop, many of the features are direct translations of concepts that existed in predigital tools—pens, scissors, and the like. Initially, designers approached these tools with the same tasks and the same aesthetic goals as they had before. Quickly, though, they started to see opportunities that were not possible from the capacities of a darkroom. They began to have ideas that would have required so many steps or such precision that those ideas would simply have not come to mind in the earlier context of paper and photographic emulsion. With digital tools, it became easier to scaffold certain types of processes and therefore to scaffold increasingly complex aesthetic goals. We are now entering another dramatic leap forward in this respect. Machine intelligence will enable creatives to do even more and to think even bigger.

Machine intelligence will also enable people to interact with design in a way that requires a lot less tool learning. Rather than you having to speak the language of the tool, you will be able to express your ideas in the form that you hold them. If you want to make the sky brighter, for example, rather than saying, "I need to select all of the pixels that represent sky and then go to this menu and drop down and go to the brightness slider," we can instead issue a semantically formulated command to simply make the sky brighter. That opens the door to more people; it democratizes the design process and leads to a greater volume of design work being created by a wider range of people. From an accessibility point of view, we can start to think about the machine understanding lots of different modes of input. Different people think in different ways—speech may be preferable for one person while demonstrative gestures may be preferable for another. There is a very real path to opening up new opportunities for human designers through these technologies. I am very excited about that and glad to be a part of it.