JS

FORMS

Haptics

AUDIO

VIDEO

INTERACTION TECHNIQUES

MULTIMEDIA

Chatbots

AGENTS

BIG
DATA

HILL
CLIMBING

NEURAL
NETWORKS

MISCELLANEOUS

ARTIFICIAL INTELLIGENCE

VIRTUAL REALITY
TREE MAPS
HEAT MAPS

VISUALIZATIONS

CONCEPTUAL MODELS

ACTIONS

OBJECTS

CLASSES

DEFAULTS

Intelligence on Tap:

Artificial Intelligence as a New Design Material

Lars Erik Holmquist, Northumbria University

Insights

- Through a combination of factors, AI has recently made significant progress and is now integrated in many successful products.
- In the future, AI will become available as a resource to use by non-experts—intelligence on tap.
- Interaction designers need to consider AI as a new design material, with its own unique opportunities and limitations.

There has been a revolution, but it snuck up on us so gradually that you'd be forgiven if you missed it. It's called artificial intelligence, and it will have a profound impact on how we design digital products in the near future.

This has been something of an unexpected comeback. In the very early days of computing, many expected that machines would soon be able to complement or even surpass humans in tasks requiring intelligence. But while well-defined undertakings, such as playing chess, have proven to be solvable by using strict rules, more fuzzy problems, such as recognizing a cat in a photo, have turned out to be

much more elusive. And so for decades, the idea of artificial intelligence has been considered mostly an unkept promise. While applications of machine learning have been increasingly useful when it comes to processing big-data collections at major Internet companies, the consensus has been that for most practical applications, human intelligence simply cannot be replaced.

But recently, artificial intelligence, or AI for short, has actually begun to deliver. New or revitalized techniques have started to equal or even surpass humans in tasks previously thought out of reach, from speech recognition to playing complex games. The rate

of this AI resurgence has taken a back even leaders of the industries being affected the most. Google co-founder Sergey Brin said in a recent interview that he has been surprised by the recent surge in practical applications for artificial intelligence [1]. Approaches such as neural networks and deep learning, coupled with access to massive amounts of data and new computational hardware, have led to significantly better results than traditional methods in areas such as image recognition, machine translation, and speech synthesis.

In one of the more spectacular examples, Google's Deep Mind software was able to beat a grandmaster of the ancient Chinese game Go in March 2016 [2]. Computers had already proven they could beat world masters in chess, but Go was thought to be out of reach because it contains exponentially more possible move combinations—far beyond what can be stored in any computer. In order to achieve this feat, rather than working from a list of possible moves, the software instead taught itself to play the game. First, it got a solid foundation by training on millions of existing human Go games. But that was not enough. To improve its game, the network then played many more matches *against itself*. In this way, it ended up with a vocabulary of moves that consisted of both human and self-taught strategies. The result was a game-playing software that played from its own experience, not from any strict set of rules.

If playing an old Chinese game sounds too esoteric, consider this: A neural network derived from the same basic techniques was trained to control the cooling processes in Google's data centers, in a way similar to how it learned Go. This time it was no game; it had dramatic financial consequences. The company claims that through the smarter control provided by this software, it has been able to save *several hundred million dollars* in electricity per year—thus by itself paying for Google's

acquisition of the AI startup company whose research laid the foundation for the system. [3]

And that is just the beginning. In the past year, the collection of AI techniques called *deep learning* have contributed to significant advances in a whole range of areas, including speech synthesis, speech recognition, machine translation, image recognition, and image compression [4]. And although the results are still largely coming in areas dominated by big data and big Internet companies, it is clear that AI will soon have implications for a whole range of new products. It will eventually make it possible to inject a little bit of intelligence into even the most mundane product, whether a toaster or a car. By extension, this will fundamentally affect HCI research and the practice of interaction design.

But before we go on, let's try to unpack the recent developments that have surprised even people like Google's co-founder.

AI COMES OF AGE

Algorithms inspired by how the brain works, so-called *artificial neural networks*, date back to the 1960s; most computer science students still encounter them in introductory AI classes. These networks are formed by connections of artificial “neurons,” which are basically just weighted links between nodes in a graph. The actual network itself does not have any inherent meaning or knowledge. But by subjecting the network to stimuli and reinforcing the links that are used when it makes the correct choices, it is possible to train the network to make choices. For instance, by subjecting a network to a sequence of pictures with simple geometrical shapes but reinforcing it only when the network selects those that contain a circle, it would be possible to teach it to pick only images that depict a circle.

The main technology leading the current AI resurgence is neural networks. However, for a long time it

looked like neural networks would be limited to simple problems with little practical use. This is because, first of all, to do anything beyond the most trivial tasks, the number of nodes and connections in such a network would have to be very large. This means that it would take a long time to train it, and even when it was fully trained, the time it would take to reply to a query would be too long for any time-critical applications such as automatic translation. Second, in order to learn anything meaningful, the network would also need huge amounts of training data. Such data would need to be in machine-readable form. The data would also have to be coded, meaning it would already contain the answer to the question the network was being designed to answer. For instance, for our fledgling network to learn to recognize circles in pictures, we would have to subject it to a large number of pictures that contained circles and were correctly labeled as such, as well as pictures that did *not* contain circles, so that it would eventually learn the difference.

But recently, these barriers have all but disappeared. When it comes to size and speed, Moore's law has been helpful, but not sufficient, in reducing the cost of storage and processing time. Instead, a much bigger breakthrough came from an unexpected source: computer games. In 2012, researchers at the University of Toronto showed that the specialized chips that are used to generate fast high-resolution graphics in PCs, so-called graphics processing units or GPUs, just happen to be perfectly provisioned for processing neural networks [5]. This is because they are designed to process massively parallel tasks at a very high speed. In other words, the right skill for drawing realistic zombies on a teenager's video-game screen also turns out to be exactly what is needed for running a neural network! Thus, almost by accident, neural network researchers were handed fast and inexpensive hardware on which to run their experiments, something that is now revolutionizing the entire chip industry [5]. This in turn allowed for new and more effective techniques such as deep neural networks (the layering of several levels of networks) and unsupervised learning (which does away with explicit labels and presents the network with only rough clusters of data). Together,

The right skill for drawing realistic zombies on a teenager's video-game screen also turns out to be exactly what is needed for running a neural network!

these advances contributed to results like the Go game victory.

And when it comes to data itself, there's a veritable mother lode. Facebook, Google, Amazon, and the other Internet giants have already been patiently Hoovering up every scrap of input generated by their users for decades. They now have access to billions upon billions of photos, emails, videos, and chat messages, not to mention mouse clicks and finger taps on everything from inspirational articles about yoga to diaper advertisements. This manic data collection is also reaching its tendrils out into the real world, for instance through mobile phones, taking in things like the user's geographical location (through GPS) or their physical activity (through motion sensors). And if you hadn't noticed, neural networks are already listening to what you are saying! Companies like Apple and Microsoft are storing every command given to their respective voice assistants for future use, in order to better train their recognition software. In this case, Siri, Cortana, and of course Amazon's Alexa and their ilk, are serving not just as helpful assistants but also as Trojan horses to gather unheard amounts of voice utterances and associated behaviors to feed the neural networks of the future. As if this wasn't enough, emerging technologies such as drones and self-driving cars will soon add ever bigger piles to this data stash.

Of course, this data gold rush has consequences that can be troubling. Most obviously, consider the fact that all this data is in the hand of private companies. They now have literally unlimited access to everything generated by our private and public digital lives but are not governed by any of the rules for transparency or privacy that pertain to public organizations. This leads to another, less obvious, consequence, which is that many of the best minds in the field will no longer be found at universities, where they can freely share their knowledge. Instead, they are being aggressively recruited by well-funded companies, where they not only get better salaries (and free food to boot) but, more important, much more challenging problems to work on. This is because the big data that is necessary to provide truly groundbreaking research resides at these companies, where it is also increasingly well protected, since it constitutes the very

essence of the companies' value on the stock market. While once upon a time Flickr set its user agreement to the altruistic Creative Commons license by default, meaning that images could be freely used for noncommercial purposes and released as large training sets for the benefit of science [6], current services guard their content much more jealously. For instance, Instagram pictures, while free to browse, are bound by agreements that prohibit any application of computer vision, making them in effect inaccessible for any machine-learning approaches.

On the other hand, there are encouraging signs that the tools of this new and efficient AI will become more accessible, often when universities and industry work in concert. Open source software such as Tensorflow is already letting users adapt and train neural networks for new purposes [7]. These services are still far from plug-and-play; they require extensive handholding from experts to achieve any useful results. But they point to a future where neural networks are packaged in such a way that non-experts can use them through well-defined interface mechanisms. Most likely, due to size and speed limitations, this will happen not on individual devices but on remote servers. Thus, just like other data- and processing-intensive tasks such as cloud storage and Web hosting, AI will transform into a service.

And with commercially available AI services bound to arise, it will gradually become easier to obtain and train an artificial intelligence to do your bidding. This means that in the near future, designers will no longer have to be experts in neural networking to use AI, just as they do not need to know the ins and outs of TCP/IP or even HTML to design Web pages. The same services will be available when designing physical artifacts, too, to complement other elements such as sensors and actuators. When this happens, AI will be thought of not as an exotic and complicated technology that can be used only by gurus with Ph.Ds in machine learning, but rather as a resource you can plug into any new product when you need it. Think of it as intelligence on tap.

A NEW DESIGN MATERIAL

So what exactly does this intelligence on tap mean for interaction design? First and foremost, it means that *intelligence*

is becoming a new design material. As we know, the options of a designer are to a large extent defined by the materials they have to work with. For instance, a graphic designer working in the medium of print must be familiar with paper sizes and coating types, as well as color blends, printing presses, and other means of achieving their desired results. A product designer would need to be aware of the physical characteristics of materials such as plastic, wood, and metal, as well as how these fit together mechanically, in order to design an aesthetically as well as functionally pleasing experience. As AI becomes a more and more vital part of everyday products, designers will have to figure out how to work with intelligence as a new material, with its own specific quirks and opportunities. This will not be easy, as intelligence on tap could mean a radical departure from previous design practices, as when going from paper to screen in the early days of the Web.

For anyone developing products that contain AI (including but not in any way limited to designers), it will be necessary to form a clear understanding of what AI *can and cannot do*. Again, this does not mean that everyone has to become a neural networking guru, but it is necessary to understand the underpinning principles of AI. In particular, this means that if someone tries to design a product without a firm understanding of the limitations of AI, the result will most certainly be failure.

Here, the most important limitation to consider is the fact that AI still cannot form an actual understanding of the world. While neural networks can indeed work better than humans on problems that involve large amounts of data, and can seemingly reply in intelligent ways to many queries, they still cannot understand a basic sentence in natural language. This has particular relevance to some of the most hyped AI applications, such as natural-language dialogue systems, aka chatbots. As overly enthusiastic product designers have already discovered, it is currently far beyond the reach of any neural network to carry out an intelligent conversation. For instance, Facebook's recent experiments in chatbots ended in something of a fiasco after it turned out it could correctly fulfill only about 30 percent of the requests [8].

There is an important lesson to be learned there. Replacing human-to-human interaction in realistic situations

is exactly something that AI *cannot* do yet. This is the kind of problem that requires a real understanding of the world and the intentions of the conversation partner—something that today’s neural networks are simply incapable of. Furthermore, it is well known from research that dialogue systems are more efficient when users do not expect the bot to have full, human-like intelligence [9]. Thus, by trying to apply human standards to an automatic system, the constructors of the Facebook chatbot literally set it up for failure and made users even more disappointed and frustrated.

Instead, artificially intelligent systems should concentrate on things that humans cannot do but that AI can do well. In large part, this involves sifting through immense amounts of data and finding patterns. One area where AI is making great progress is image search, in which large amounts of data and new neural-network techniques have produced remarkable results, such as actually being able to find pictures that contain cats. Other areas where AI does well, as long as there is enough data, is matching one dataset against another, for instance in machine translation. It can also be used to extrapolate from existing data and make decisions based on that, as with Google’s server-cooling system. But this also means that AI systems are highly dependent on the data they have access to. If the data is lacking in quality or quantity, this will greatly increase the risk of the system making poor decisions.

Thus, anyone constructing an AI-based system needs to tread lightly, manage expectations, and be careful not to overreach when it comes to AI’s capabilities. But apart from understanding the overall potential of AI, I believe there are a number of interdependent challenges that pertain more specifically to interaction design. These have to do with how designers can take the behavior of systems that rely on artificial intelligence and make

it understandable for the end user. They include:

- Designing for transparency
- Designing for opacity
- Designing for unpredictability
- Designing for learning
- Designing for evolution
- Designing for shared control.

The first challenge means that it is necessary to let the user understand how artificial intelligence is actually affecting the interaction. It must be clear to the user that a system is actually making its own decisions based on incoming information, rather than working from a fixed set of rules. This might require the rethinking of fundamental UI components. For instance, there are interaction cases when users might want to override the intelligence, and others when they might want to cede control. For a device, this could mean that rather than just an on/off button, a device might need an “it depends” button that lets the device decide whether to turn on or off. Similarly, there will also be a need for interface elements that communicate when a system has made a decision, what that decision was based on, and even a mechanism to revert or undo the decision if the user does not agree with it. There could also be a need to communicate more complex concepts and plans to an AI, which might require more flexible interfaces such as natural language. In summary, designing for AI might entail a lot more fuzzy, open-ended user interfaces than we are used to.

The second, somewhat contradictory, challenge has to do with the fact that it is no longer possible to explain exactly why or how an AI does what it does—they are opaque. The way that neural networks are constructed means that their inner workings are hidden even from the person who programmed and trained them. For example, Google’s engineers recently made the discovery that a neural network trained for machine translation had created its own intermediary format [10]. This made

it possible for it to translate between language pairs on which it had not been trained; for instance, if it had done Japanese to English, and English to Korean, it could also in principle translate between Japanese and Korean. The point here is that this capability was not designed into the system, but rather evolved by itself. How can designers communicate to the user that there are things inside the product whose workings nobody can quite explain? And how does this affect qualities like trust and confidence in the system?

This leads to the third challenge: unpredictability. No matter how well trained a neural network is, it is still to some extent drawing its own conclusions from given data. This is not necessarily a bad thing. For instance, the Go-playing network we mentioned in the beginning had honed its game not just on humans but also in matches against itself, where it devised its own strategies. This led it to make some surprising moves that no human player would make. While some of the choices it made were inexplicable, they were also part of a winning strategy, and despite deviating from the human playbook, in the end the system was able to beat the human opponent. Designers thus must be prepared for and design for systems that behave in unanticipated ways, which can be jarring even when it leads to them solving the problem better than a human would. How can interaction design minimize the damage and maximize the benefits that arise from this unpredictability?

The fourth challenge has to do with improving the AI through constant learning. Ideally, a neural network should never stop learning; it should use all available new input to improve its basic algorithms and make the system even better. However, this cannot be a chore for the user. If the user has to explicitly train the system, that will most likely become a hindrance to efficient use. There are already clever ways of having humans solve problems to aid AI learning, such as the “captchas” that separate humans from bots on the Internet by having them do simple image-recognition tasks. Another example is recommender systems on sites such as Netflix that encourage users to rate the content they have viewed, thereby improving recommendations. But ultimately, the learning has to be

Anyone constructing an AI-based system needs to tread lightly, manage expectations, and be careful not to overreach.

built into the interaction itself and completely unobtrusive, so it does not feel like the user is doubling as the AI's training wheels.

The fifth challenge has to do with how these systems will continue to evolve over time. As AI products solve problems in collaboration with their users, they should keep improving. But this could be jarring if the system's behavior starts to get better than it was originally. In fact, we often build behaviors around flaws like squeaky doors or loose tiles in a staircase. If these flaws suddenly disappear without warning, it might be even more disorienting than when they first appeared. Say you have bought an intelligent coffee brewer that is supposed to prepare coffee at the right time and temperature to help you get up in the morning. You set it for a certain time, but you have a hard time getting up, so the coffee is always a little cold. And that's OK; you need your sleep. But imagine then that the brewer observes how you are always late getting up in the morning, and one day it proactively decides to delay the brewing of your coffee by 10 minutes to better fit your schedule. The result is that you scald your mouth—and probably throw the coffee maker out the window! As systems evolve and make new decisions, it will be necessary to communicate this to the user so that they know what to expect, and can benefit while avoiding unpleasant surprises.

The final challenge is one that springs from all the others. It involves how artificially intelligent systems can be designed to allow the sharing of control with the user. This will not be an either/or situation, where one or the other has full control. In systems built on proactive intelligence, there will have to be provisions for a truly mutual responsibility. The interface must give the user access to clear controls, as well as indications as to how the power is distributed in any given moment. This includes how much autonomy a system receives to make its own decisions and how much it is under the control of the user. It also includes how much it is allowed to evolve new functionality, how it collects and evaluates data, how it is to handle unexpected situations, and so on. Again,

some of this may be too complex to be fully negotiated by a visual or tangible interface, which may lead to the need for speech or other more nuanced modes of communication. But designing the interaction of an AI system so that it can work truly in concert with the user will be one of the key measures of success.

There will be many other challenges as well—what I've discussed here has just scratched the surface. We did not even get into ethics, which will have a huge impact. Who is responsible if an AI system causes damage or even the loss of life? This could happen if the system made an error or was inaccurately controlled by the user, perhaps due to some flaw in the interface design. This is not a science fiction question; it is already pressingly important for companies developing self-driving vehicles. And who gets sued for libel if an AI runs amok because it is absorbing data without questioning it, like the Microsoft chatbot that became racist by reading Twitter comments [11]? Another issue is who owns and takes responsibility for material that an AI produces? Ownership was much easier before autonomous systems, because the creation of content was the result of a conscious creative act. Now if an autonomous security robot, or perhaps an outdoor drone, manages to take compromising photographs, who gets to control the results—the subject, the owner of the device, or (most likely) the company that stores the images on its servers?

Full-fledged intelligence on tap might take a long time to arrive, but I have no doubt that it will. And while enthusiasm for AI in its many forms is very high right now (Gartner's hype cycle for 2016 has machine learning at the very top [12]) and is sure to hit many snags along the way, there is no doubt that the technology is going to fundamentally change interaction design. The sooner designers start to think about intelligence as a design material, the better prepared they will be for the coming shift in how digital systems will work, and in particular how AI can function in concert with their users. Hopefully, this article has provided some first steps toward understanding the future of AI as a new design material.

ENDNOTES

1. Kharpal, A. Google co-founder Sergey Brin says he's 'surprised' by pace of A.I. and uses a story of a cat to explain it. *CNBC.com*. Jan. 19, 2017; <http://www.cnbc.com/2017/01/19/google-co-founder-sergey-brin-said-he-is-surprised-by-pace-of-ai.html>
2. Metz, C. Google's AI wins fifth and final game against Go genius Lee Sedol. *Wired*. Mar. 3, 2016; <https://www.wired.com/2016/03/googles-ai-wins-fifth-final-game-go-genius-lee-sedol/>
3. Clark, J. Google cuts its giant electricity bill with deep mind-powered AI. *Bloomberg Technology*. Jul. 19, 2016; <https://www.bloomberg.com/news/articles/2016-07-19/google-cuts-its-giant-electricity-bill-with-deepmind-powered-ai>
4. Metz, C. 2016: The year that deep learning took over the Internet. *Wired*. Dec. 26, 2016; <https://www.wired.com/2016/12/2016-year-deep-learning-took-internet/>
5. Greengard, S. GPUs reshape computing. *Communications of the ACM* 59, 9 (2016), 14–16.
6. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., and Li, L. YFCC100M: The new data in multimedia research. *Communications of the ACM* 59, 2 (2016), 64–73.
7. Abadi, M. et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467* (2016).
8. Facebook scales back AI flagship after chatbots hit 70% f-AI-lure rate. *The Register*. Mar. 22, 2017; https://www.theregister.co.uk/2017/02/22/facebook_ai_fail/
9. Cassell, J. *Embodied Conversational Agents*. MIT Press, 2000.
10. Wong, S. Google Translate AI invents its own language to translate with. *New Scientist*. Nov. 30, 2016; <https://www.newscientist.com/article/2114748-google-translate-ai-invents-its-own-language-to-translate-with/>
11. Vincent, J. Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day. *The Verge*. Mar. 24, 2016; <http://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>
12. Gartner 2016 Hype Cycle. Aug. 16, 2016; <http://www.gartner.com/newsroom/id/3412017>

Lars Erik Holmquist is professor of innovation at Northumbria University, U.K. Previously, he did research in interaction design and ubiquitous computing in Sweden, Silicon Valley, and Japan. His first book, *Grounded Innovation: Strategies for Creating Digital Products*, was published in 2012. He just finished his second, a science fiction novel set in Silicon Valley.

→ lars.holmquist@northumbria.ac.uk