



Living with Digital Employees

Ready or not, various forms of artificial intelligence (AI, mostly machine learning variants) are coming your way. In our research, the number of companies offering some sort of AI for HR is growing at a rate of 500% year over year. In May, 2017, there were 40. We forecast 200 by the end of May 2018.

The growth rate is not surprising. At the simplest, AI is the next level of abstraction in software. All providers will be delivering some sort of predictive/learning/recommendation service as a part of their offerings within four or five years.

The large providers of cloud processing/storage services (Amazon, Oracle, Google, etc.) are in an arms race to deliver usable AI tools at no charge to their clients. Their motivation is simple. Artificial intelligence uses massive quantities of processing and storage. Growing the capabilities of their clients is good for business.

Meanwhile, there is a shortage of technologists who actually understand the implications and ramifications. A search of LinkedIn for data scientists shows slightly more than 22,000 Ph.Ds. Even if that number is off by 100%, there are nowhere near enough competent practitioners.

But, that doesn't stop the marketing machines from deploying high-intensity hype campaigns. Since this is HR and all sales in our industry boil down to cost savings schemes, the claims have to do with efficiency rather than effectiveness. As a reminder, efficiency means getting things done quickly while effectiveness means getting the right things done.

The argument that you are most likely to hear is, "Machines have a higher rate of accuracy than their human counterparts." Sometimes, it's framed as, "Machines perform

measurably better than people." Another version goes, "We let humans learn their jobs, why shouldn't we do the same for machines?"

In each case, the defender of machine-led decision-making is making the assumption that humans and machines make the same types of errors. They're saying that any old 80% is the same as any other old 80%. In the case of machines, it's 80% likely that machines can repeat prior performance. In the case of humans, it's a question of whether they can make the right decision 80% of the time.

If this were a Pareto analysis¹, we'd be wondering whether the 80% is the 80% of the most important results or the 80% that is the most trivial. All we really know at this point is that machines can repeat history with 80% accuracy. The measure itself is a thing of nonsense rivaling Alice in Wonderland's most tortured logic.

"I know what you're thinking about," said Tweedledum, "but it isn't so, no how."

"Contrariwise," continued Tweedledee, "If it was so, it might be; and if it were so, it would be; but as it isn't, it ain't. That's logic."

"I was thinking," Alice said politely, "which is the best way out of this wood; it's getting so dark. Would you tell me, please?"

- *Through the Looking Glass*

This is not to say that all AI offerings (or even most) are flagrant examples of hype-infested vaporware. They are not. We are seeing tremendous strides being made by the majority of providers. The companies offering intelligent software almost all look like laboratories or innovation centers. The workers diligently spend their time

Harnessing these powerful new technologies requires a different kind of thinking. This is not more of the same old software.

attempting to build models that their offerings can use to make our lives easier.

But, this is new territory. Harnessing these powerful new technologies requires a different kind of thinking. This is not more of the same old software. The products and services offered under the rubric of AI make decisions and recommendations that used to be the province of human beings. We don't really know if they are any good at it. We don't really understand the consequences of using the tools.

We are lurching into the future.

In a well-reasoned paper, *Will Robots Steal Our Jobs*,² PwC lumps AI with other forms of business automation. They forecast the eradication of about 30% of jobs by the mid-2030s. The change happens in waves as we learn to more fully embrace the emerging tech. The report alludes to job growth caused by various forms of automation, but doesn't spend much time imagining that growth.

In what seems like an afterthought, the PwC paper mentions "learning mishaps." A learning mishap occurs when the machine doesn't have all of the information but gives a recommendation or makes a decision nonetheless. As we develop skills at using machine inputs in our decision-making, we will encounter learning mishaps.

The way that most organizational life works is a cycle. The organization strives for routine process. Once they are in place, things run smoothly until the chaos happens. The chaos is any of a thousand events from a merger or reorganization to an economic downturn to a new product launch. Each chaotic event interrupts the routine. The organization works hard to re-establish a new routine that includes whatever was learned in the chaos.

That's the cycle of organizational life. Routine precedes chaos which precedes routine. As an organization, it is the conversion of chaos into routine that propels us forward (or, maybe it's the other way around, or maybe it's both).

In each transition between routine, chaos and back again, machines must be retrained. While they are being retrained, the quality of their decisions declines. This decision variability is impossible for the machine to understand. The humans who supervise the machine must be the ones to understand and evaluate the quality of the machine's work.

Since the machine can't understand its own limitations (they are much like young humans in this regard), their supervisors and managers must do it for them. In our research, we call this the "latency problem." The difference between what the machine understands and reality is a high-risk source of discontinuities and "learning mishaps."

The machine's ability to perform predictably is the result of a combination of factors:

- **Quality of the underlying model** – For the most part, data models are judged to be acceptable if they can predict a percentage of last year's results. That does not mean that they have a clear picture of reality (and that's where some of the risk lies).
- **Training effectiveness** – The difference between 80% and something better comes from allowing the system to learn based on real-time data. This is a pure volume game. The more instances in the training set, the better the model will perform.
- **Latency** – The machine sees the world through the narrow lens of things that can be measured. When circumstances change, it must retrain. Latency is the difference between the real world and the machine's picture. At its worst, latency means things are garbled. At best, it's like the difference between a color photo and a black and white version.
- **Task complexity** – Simpler tasks involving a repeatable decision tree and yes or no answers work best. Complicated real-time judgments in rapidly changing circumstances are much more difficult.
- **Data quality** – Self-reported data manually entered by humans is the worst. Data that comes directly from measurements that a machine makes is the best. There is every reason to believe that personality assessments will improve in quality because they are increasingly relying on machine-generated data.
- **Feedback loop** – When the machine errs or cannot perform (resulting in a blank screen or a call to a human for assistance), there must be a way to simultaneously answer the question and pass the feedback about the defect to the right place.

-
- **Self-diagnostics** – This is the trickiest. The machine should be able to tell you when it suspects that its answer is incorrect.

With each new installation of an intelligent function, you need a “manager” who:

- Understands the whole organization well enough to troubleshoot multiple issues simultaneously,
- Understands the data sources well enough to evaluate their quality, and
- Understands how to “fix” the errant AI.

Implementing an AI project will involve training and retraining the tool itself,

the users, the managers of the users, the managers of the tool in usage and development, and the overall management team. It’s a complex undertaking and the foundation of an argument that you can’t implement AI without an overarching transformation process.

The emerging generation of software does not behave like its predecessors. Where the measure of adoption used to be how much time one spent working in the software, these new tools directly impact job performance. The job will no longer involve working in the software. Instead, we are headed back to doing our jobs.

Endnotes

¹ https://en.wikipedia.org/wiki/Pareto_analysis

² <https://www.pwc.co.uk/economic-services/ukey/pwcukey-section-4-automation-march-2017-v2.pdf>

About the Author

John Sumser is the principal analyst and editor-in-chief of the HRExaminer. He is an independent analyst covering the entirety of the HR technology ecosystem from Payroll and Benefits to Recruiting. He has a particular focus on ethics and practices in predictive analytics. Sumser routinely advises Human Resources, Recruiting Departments, and Talent Management Vendor teams with product analysis, market segmentation, positioning, strategy, and branding guidance. He’s been published and quoted in every imaginable outlet from *The New York Times* to HR industry trade magazines. He is currently experimenting with the use of Facebook as a forum for industry dialog. Follow him there. You can also reach him at john@hrexaminer.com.