

Shane Greenstein:

This is professor Shane Greenstein with the HBS Digital Initiative from the 2019 Future Assembly at Harvard Business School. We are pleased to present Flash Talks, exploring technology in a multi-disciplinary world.

Perry Hewitt:

Good morning! It's great to be back at Harvard. I should admit that I'm only Marriott Bonvoy Silver, but in the last year I've stayed in both a penthouse in Paris, and a trailer in Dallas, in a Airbnb, and I would not advise the latter, just FYI.

Perry Hewitt:

So, today we're going to look at the known issue of normative defaults in interface design. We're going to consider the lessons we learned during this normative defaults are and are not applicable to the AI era, run through a couple of examples, and then some ideas for mitigation.

Perry Hewitt:

As the GIF shows, and I think only David and I remember this GIF, we've come a long way on the internet. Some of the people in this room have been part of shaping the past quarter century of growth on the internet from the era of million dollar web servers to codeless applications. But as a result of that rapid growth, we've had to make a lot of design and development decisions about digital interfaces that shape an enormous number of people's experiences.

Perry Hewitt:

We all make assumptions and they can certainly provide useful shorthand. But to believe assumptions are universally applicable is a means to being really wrong, really fast at scale. With a relatively smaller number of people creating interfaces during this rapid growth era, we endured a ton of baked-in assumptions. Despite the utopian hopes many of us held during the early years of the internet, the offline attitudes and behaviors, of course, shaped online experiences.

Perry Hewitt:

For example, cultural assumptions informed the normative default copy is my favorite sports team, and anyone who's a Gunner's fan should totally come to my breakout session. But the microcopy, e-commerce microcopy you choose, it makes you decide whether you're going to be an adult or a woman. I guess the two are mutually incompatible. Cultural assumptions also inform the normative defaults and account security questions. My favorite ones include:

Perry Hewitt:

What adventurous sport does your father like? Because only men really like sports.

Perry Hewitt:

Which company did you first buy shares in? We assume you, Mr. Moneybags, are a big time and longterm shareholder.

Perry Hewitt:

And where did you meet your spouse? We assume you are partnered.

Perry Hewitt:

But do normative defaults like this matter to users of the internet? Well, I guess they do affect those users who happen to be people with feelings. [inaudible 00:02:23] internet tendencies already lampooned this kind of angst and do need questionnaire. I mean, what is the color of the fur of the cat who'll be the first to eat you after you die, miserable alone at home, right? These kinds of security questions elicit emotions when they confront users with these default assumptions about who the respondents are. Their car ownership, their marital status, all those things we tried to portray so strongly to one another in disclosing our statuses.

Perry Hewitt:

Ideally we can find a way to create accounts that causes fewer existential crises in the future. So I picked on mostly language in the examples I've chosen. But there are, of course, broader components to consider when creating interfaces including visual design, interaction design and the presence of absence of certain elements.

Perry Hewitt:

A few of the ways to mitigate them:

Perry Hewitt:

You can diversify your product teams. You can steward live research practices that ensure that users remain top of mind rather than finding blocked into some hideous SharePoint server you can never quite access.

Perry Hewitt:

You can develop intentional and not accidental content strategy. I could really tell some tales about the number of critical large scale applications I've seen where the microcopy at the end of the application is left entirely to the UI developers.

Perry Hewitt:

The last one, to touch the application. Finally, having regular conversations about what makes design decisions matter. What makes them normative before issues arise and not after them?

Perry Hewitt:

So as we've seen play out in many settings, a few individuals' blind spots mean a massive world changing system has those same blind spots. Now we're ending the era of AI everywhere. Sure, it's not yet the robotic super Scarlett Johansson and her general artificial intelligence that mimics human behavior, which our society seems to both revere and fear in equal part.

Perry Hewitt:

It's more a narrow AI, which serves as an underpinning for a number of our day to day applications. If you work as I do in a marketing and business context, implementing AI in 2019 is not so much about these sexy Turing Award winning neural network theory pieces. It's deployed much more under the hood to make a better digital experience through platforms that enable things like marketing or automation or even internal document discovery. So everyone could see the issues, right? You could see them out loud with the normative of defaults in the first era.

Perry Hewitt:

Our challenge with AI defaults is they can be largely invisible to the end user. This is true of any personalized service on the web, of course. My airline flight pricing, very on topic with OTAs, is likely to be very different than yours based on our past purchase behavior. But AI's inferences may have a broad range of inputs that affect decision making and we need to explore whether those inputs are accurate.

Perry Hewitt:

Scott Belsky wrote a great piece about the tricky nature of these new defaults framed as disruptive interfaces. What I'm calling normative defaults in the AI era, he's saying, "They're just disruptive interfaces. They're drastically simpler, they're more accessible and they ultimately commoditize everything underneath." The interfaces are successful because they offer a better UX. It's a much more streamlined user experience than the systems they're replacing. And the battle to be that default is on.

Perry Hewitt:

So how much should the end user be informed about the AI that has gone into that decision? As these defaults become less visible to users, what is the new industry standard for communicating the presence of artificial intelligence and the reasons behind the choices you are seeing? How about the reasons behind the choices you are not seeing? What are a few of the risks of determining these due to defaults and how might we mitigate them?

Perry Hewitt:

Risk one from this Google trends chart, you can see that we've had a huge uptick in interest in fields like machine learning, and deep learning looks a lot like that. But less attention to the underpinning of training data, which is the core. A lack of solid training data represents a real risk for organizations seeking to create AI informed interfaces. To be effective training data sets must be sufficiently large and accurately labeled, and represented of the problem your product or service seeks to solve.

Perry Hewitt:

The data challenge is real, not only in availability and cleanliness, but in selection. A colleague recently described to me a challenge with a big box retailer in a South American country that was using a set of data to serve all its stores in its countries. The goal was for AI to automate and deliver better decisions and make it possible for the associates to have customer experience be a real differentiator on the ground. Unfortunately, the data set they focused on was way too broad, meaning that they included all the underperforming stores and resulted in a very costly flawed model.

Perry Hewitt:

In the social sector, I've looked at a few ways to highlight the training data, but in the social sector it's important to think about what's in your training data set. So this is an example from an experiment that was done at MIT, where they looked at these different companies that do facial recognition software. They all did poorly on darker people and they did worst of all on women of color. So it's very important to understand what's in there.

Perry Hewitt:

So, how do you mitigate these? You need to look about training datasets that are used to most accurate they can be. You need to think about how you scale and grow them where you don't have enough actual

data and you need to artificially grow them. You need to think about what's in your dataset. For example, 60% of location based data in the image net open database. So this is used in development for a lot. It comes from six countries in North American Europe. Understand what's in the data set and the accuracy of your labeling, which may vary by geography, will affect your defaults.

Perry Hewitt:

A second risk is, of course, releasing your AI to production full bore without thinking through potential outcomes. In 2016, where most of us will be familiar, AI was released, an AI chatbot was released by Microsoft. They said, "Oh, this is totally built using relevant public data. It's been modeled, cleaned and filtered. Don't worry about this." But it didn't do us much good against Twitter trolls, right? Your AI won't exist in a vacuum and unlike an interface, it will evolve with use.

Perry Hewitt:

How might [inaudible 00:08:05] have been averted? Well, how about a prelaunch checklist? A 2018 article talked a lot about the checklists in data science versus the lofty ethical oaths people make at the executive level and flaunt in press releases. Which makes me think of a high profile AI ethics board that was formed and disbanded this week. Broad public oaths have a number of weaknesses including their one shots. You take it once and then you're done. They can falsely inoculate you against doing unethical work. It makes you think like, "We can't be unethical. I totally signed that document. We're good."

Perry Hewitt:

AI needs a set of practical safeguards. More like the checklist popularized by Atul Gawande. There's a much longer checklist in the article, but I thought these were pretty good representative questions, particularly about shutting it down.

Perry Hewitt:

Risk three, the disconnect between the elite squad known as the special AI unit and the rest of the people who are working on in the organization. Reminds me of a little digital 1.0, when the web was surgically distanced from everything else going on in your business. A growing gap between the product function and how people in your organization understand it, can really present a risk. If your product is being built in a way that's antithetical to the values of your employees, the risk intensifies and Edelman Trust Barometer found a dollar costs associated with the gap in alignment between what your products and services do and your companies and your employees values.

Perry Hewitt:

So what might that look like? For example, building a hypothetical facial recognition model that charges women more than men for a vending machine espresso would cost you in credibility in dollars with your employees.

Perry Hewitt:

So when people think about preparing the workforce for AI, they think first about getting the data set right, the data science in place. But they think less about as the broader employee base, customer facing and beyond with a practical and educational need to understand how AI is doing. Employees need knowledge of and trust in the systems. I've seen this with digital marketing. When people don't know how to work alongside AI, they can even sabotage the efforts that are working on your organization.

That education may need to extend to the rank and file in industries you might not normally associate with AI. Right?

Perry Hewitt:

I'm sure many of you saw last week that McDonald's paid 300 million reportedly for dynamic yields. It's first application to say, "I will be the drive through window." So maybe my idea of charging sleep-deprived mothers more for espresso isn't so far off. Thus far, knowing your facial recognition is planned, but they will use it to artificially change... tweak drive through menus. For example, it might give McFlurries more prominent placement of the day.

Perry Hewitt:

Going back to the point about education, who needs to know about this baked-in default? The coder and the interaction designer? How about the marketer and the onsite cashier and floor manager? Does the customer need to know why items are on or off the menu on a typical day? Will they feel catered to or manipulated? A lot of people are feeling pretty manipulated these days, right? As a marketer, it's easy to get excited about the possibility for AI to foresee the consumer need, whether it's for traditional CPG software or the field I work in, which is nonprofit education development or a cultural product.

Perry Hewitt:

What we looked at, some of the genuine problems of AI fueled defaults, there are disruptive interface applications that could delight users with their time-saving and accurate defaults. We want to do it the right way. Social scientists have long known that algorithms do not just portray the world, they also change it.

Perry Hewitt:

So one of the ways to tackle this problem is to get closer to academia. There are many ways to approach this exchange between academia and industry. Greater collaboration and relieving datasets is a big one to me. Updating ethical checklist practices and getting to know one another in convenings. Just like the one we're having here today. Thank you.