“information” vs. “interaction”: a case study of affective computing

Noura Howell

I wanted to introduce affective computing because it’s one field where machine learning methods are used around something very subjective - our emotions.

With affective computing, there are two different approaches to emotion - an “information”-centric approach which emphasizes symbolic representation and an “interaction”-centric approach which focuses more on interaction We’ll get into what those are later on and see how they are at play in affective computing.

But then from there we’ll go more general and look at how these two different approaches play out in other fields, and think about implications for validity and measurement in data science.

roadmap

1. affective computing
2. “information” vs. “interaction”
3. semantic validity & measurement
4. other examples of this divide
5. alternative / hybrid models?

So, first I’ll talk about affective computing and some example applications.
Then, I’ll introduce more specifically what these information and interaction centric approaches are.
Then we’ll tie it back in to semantic validity and measurement,
look briefly at a couple other places where this kind of framing is at play,
and then open it up for discussion of how to design models with this in mind.

affective computing

- studying emotions scientifically
- designing interactive computing systems

Rosalind W. Picard. 1997. *Affective Computing*. The MIT Press.

So, what is affective computing? It draws a lot from scientific psychological studies of emotion, with the goal of designing interactive computing systems that that take into account the users' emotions. Let's look at some examples.

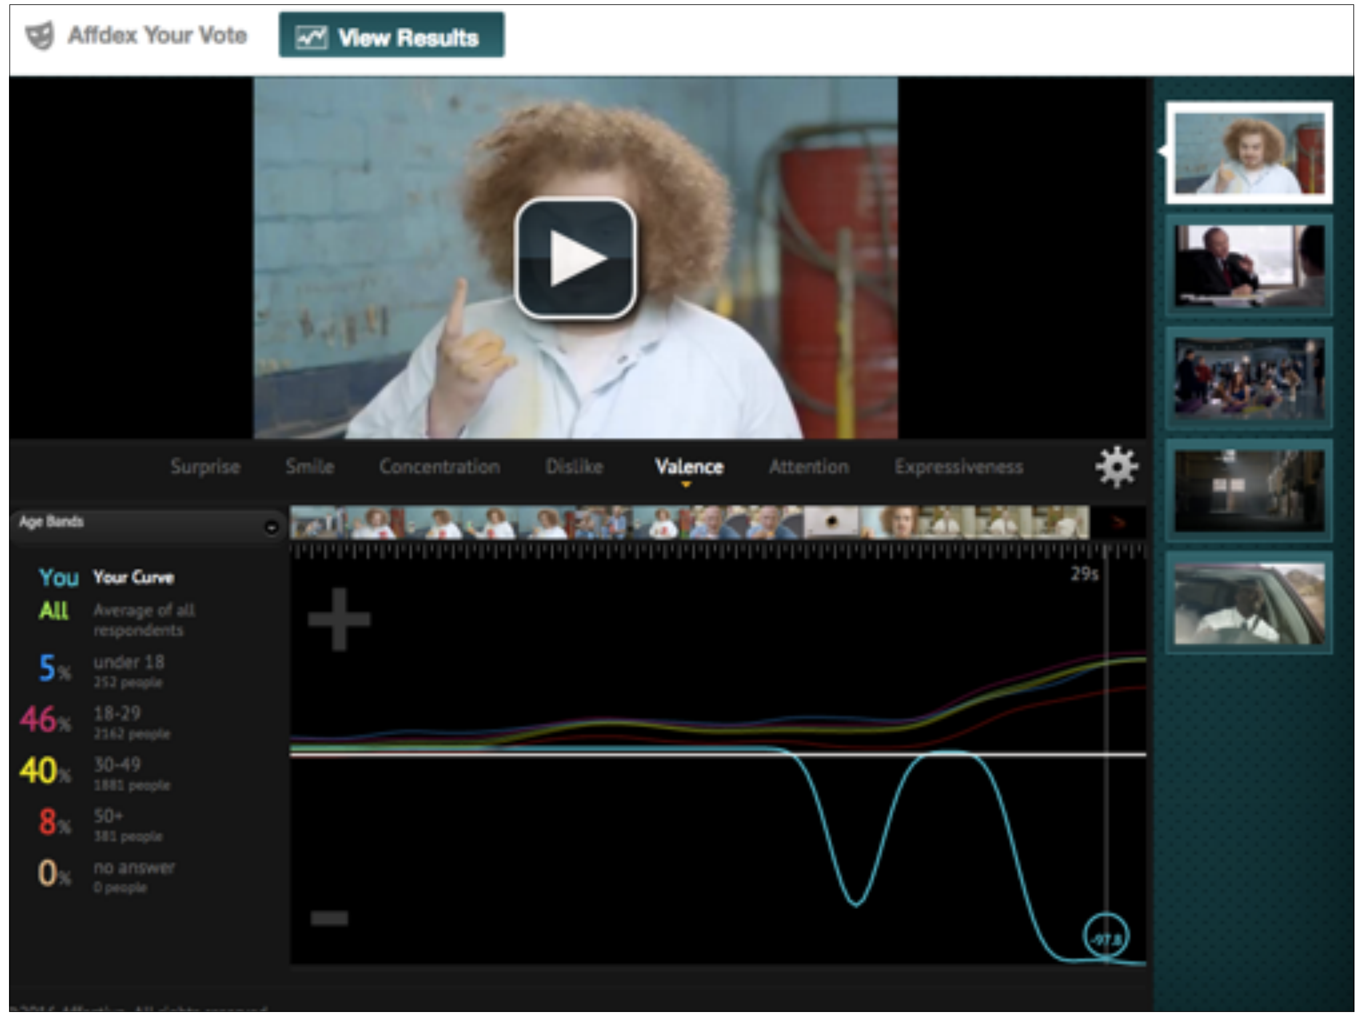

Affectiva is a company that spun out of research at the MIT Media lab, founded by Rana el Kaliouby. It takes in video of a person's face and detects different parts of their facial expression, such as the shape of their mouth or their eyes, and uses machine learning to map combinations of those to different aspects related to emotion such as valence (positive or negative feelings), surprise, smiling, etc.

I went onto their website and tried out their demo, where they tracked my face while I watched an ad. You can see my "valence" graph in the bright blue. Apparently I really didn't like this ad because the valence went super negative. Evidently most people had a more positive valence toward the end of the ad.

So, you can see how advertisers would love this kind of information to help better sell to consumers. There are a lot of other possible applications too. I think Kaliouby initially wanted to help children with autism understand how others around them were feeling to help them socialize. You could also imagine online courses that offer special help when a student appears to be frustrated.

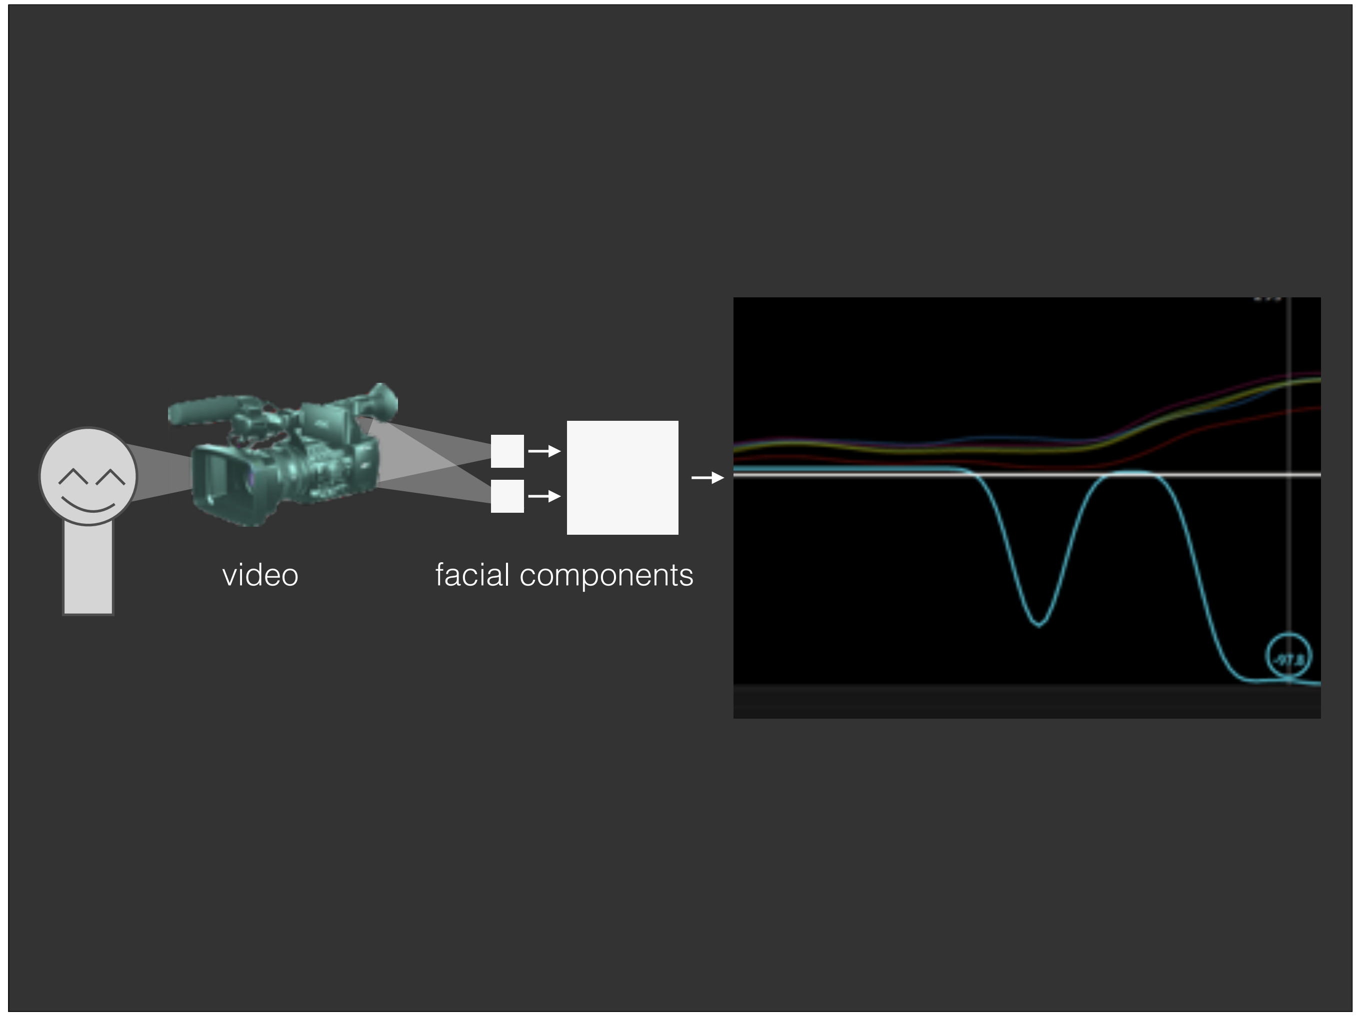

So, what's happening with this example is that video footage of the user's face is then transformed into different facial components, which are plugged into their model, which outputs predictions about values such as valence, concentration, surprise, etc.

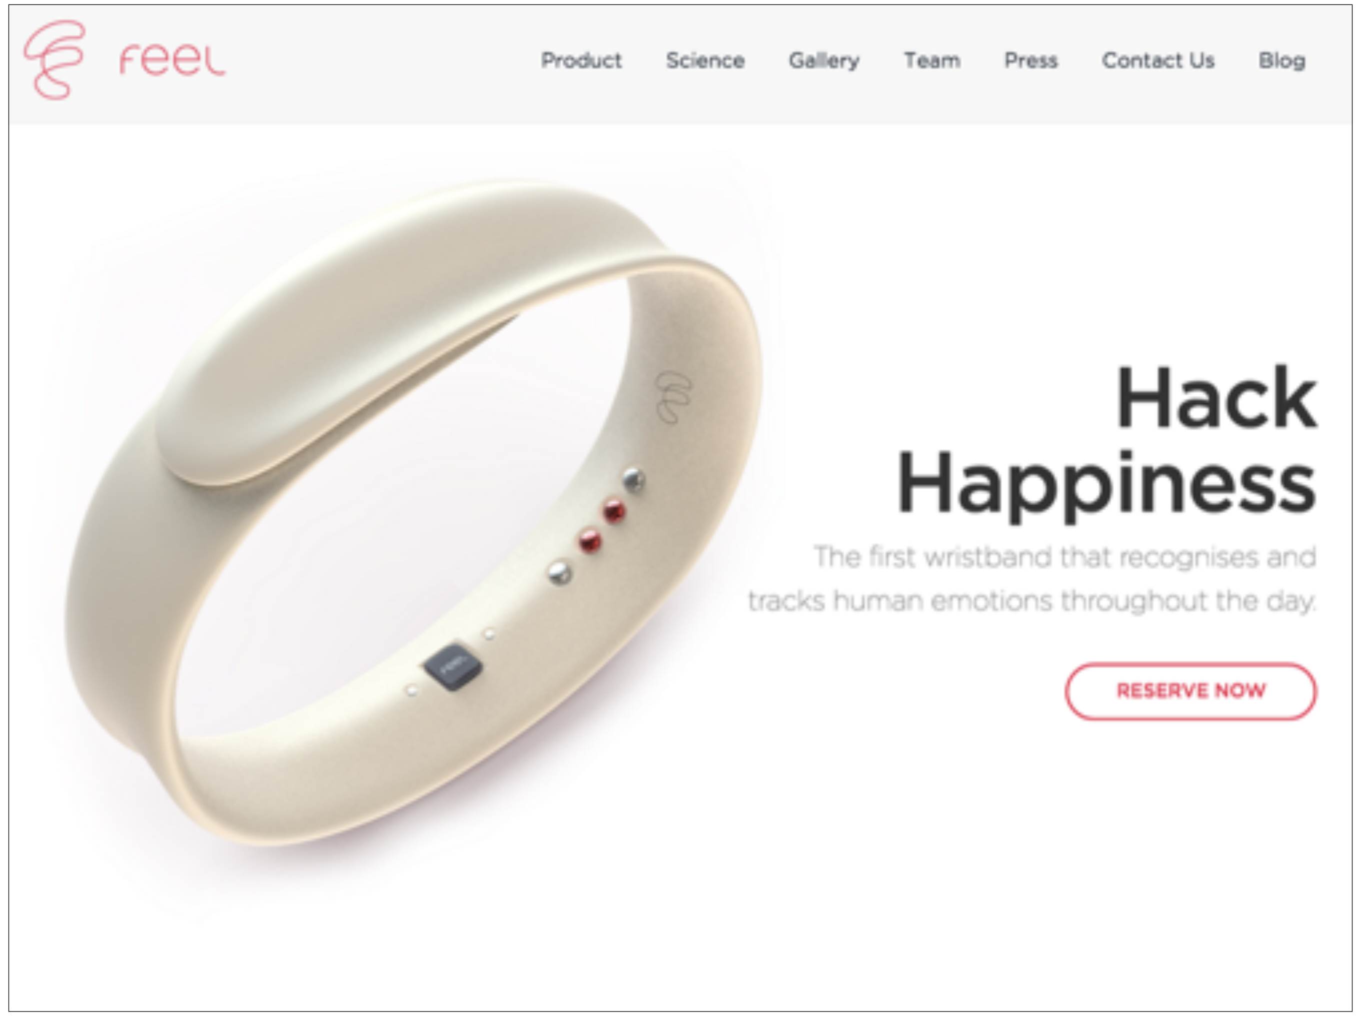

Another example that was demoed at CES of this year is the Feel wristband, that tracks the wearers' emotions throughout the day in order to help them "Hack Happiness".

They track galvanic skin response, pulse, and skin temperature. Galvanic skin response, also known as skin conductance, is a measure of how electrically conductive one's skin is. or basically how sweaty you are. but, like, micro-fluctations in that. a sudden increase in skin conductance is associated with excitement of various kinds. for example feeling nervous and having sweaty palms.

heart rate is also associated with emotions, perhaps we have more intuition about that. for example feeling afraid and your heart is pounding

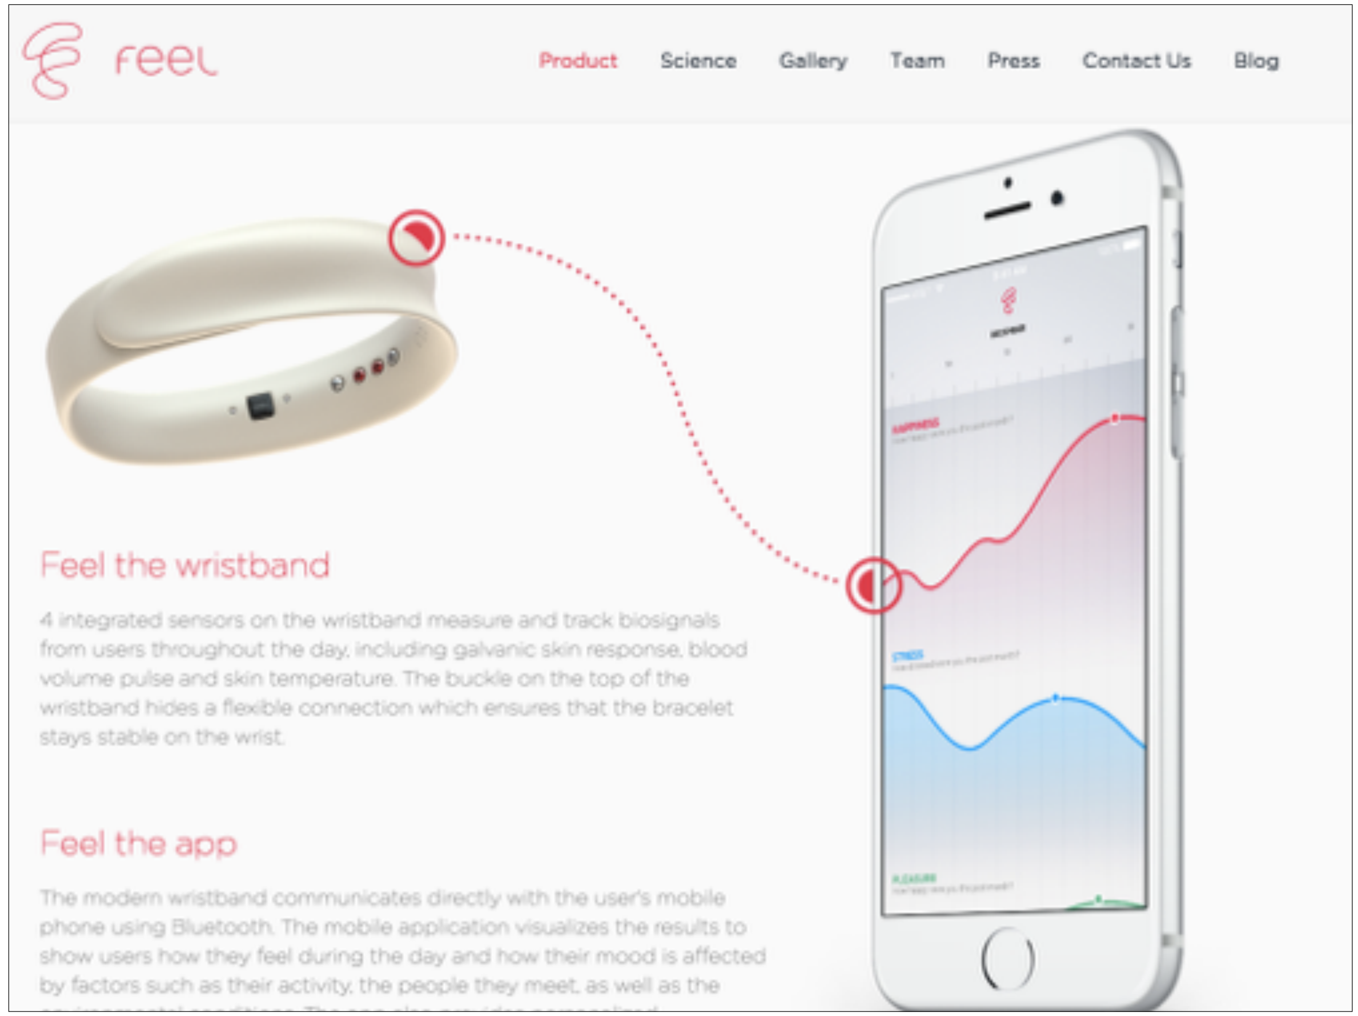

It comes with a mobile app to show you graphs of your “happiness”, “stress”, and “pleasure” over time.

So, it’s taking a quantified self approach to helping people be happy.

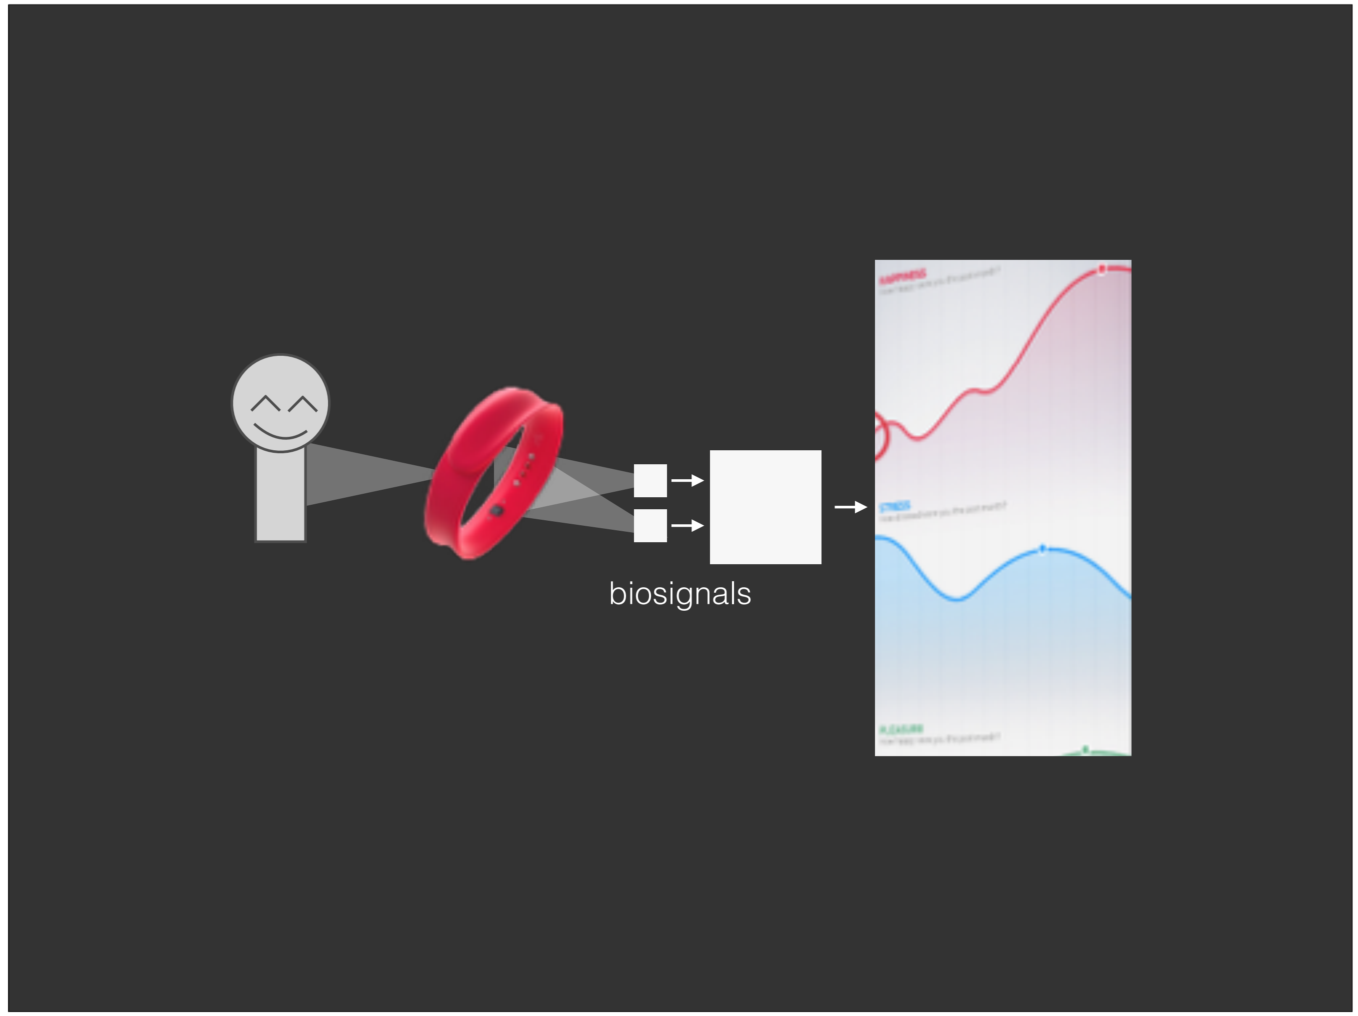

Again, to break it down, what's happening here is that a wristband is measuring various biosignals that get plugged into their model, and the model makes predictions about happiness, stress, and pleasure.

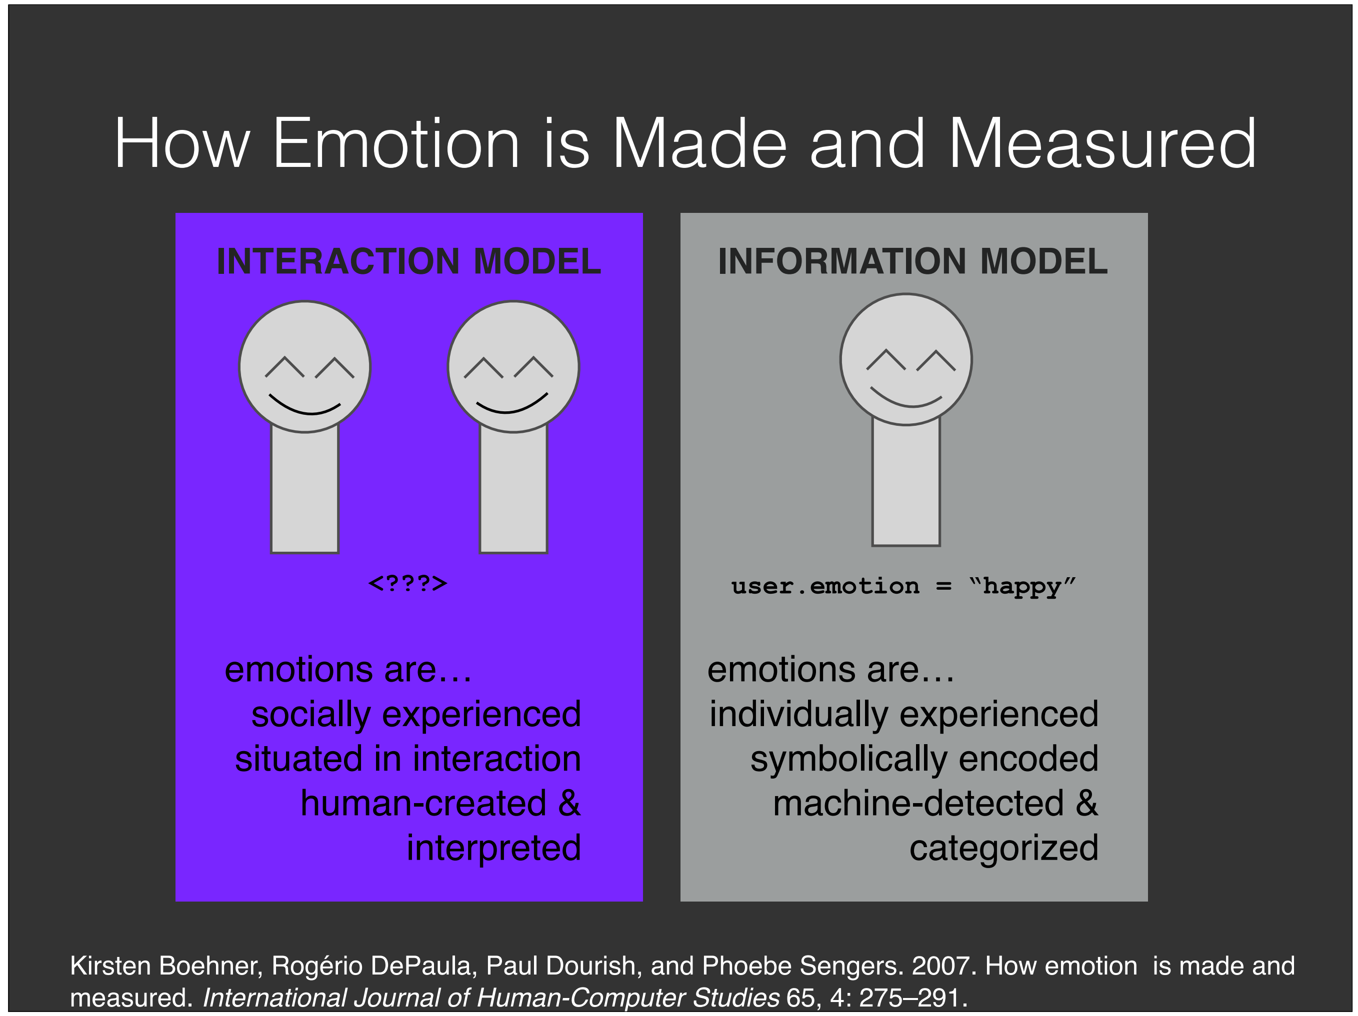

OK so after looking at some example applications, let’s look at some critique.

In the first reading, How Emotion is Made and Measured, they look at Picard’s approach to affective computing, which has largely become the standard approach, and call it the Information Model, and critique it. And they propose their own approach, which I’ll refer to as the Interaction Model.

In the Information Model, emotions are assumed to be individually experienced, measured at the level of the individual, they are symbolically encoded, and the goal is to have computers detect emotions and put emotions into discrete categories such as “happy”, “stressed”, or positive or negative valence. It’s assumed that if user.emotion equals happy, then we can transmit that information from the user to the computer and then transmit that somewhere else to another user, and the meaning of happiness is taken as being largely independent of context.

In the Interaction Model, the focus is less on detecting and categorizing emotions and more on helping other humans interpret emotion in the context of interaction. So, it’s not assumed that, say, happiness is a discrete emotion that makes sense outside of context. Rather, the emphasis is on emotions as socially situated - how we think about our own emotions, even when we’re alone, is inextricable from our culture and our interactions with others.

So, in terms of designing interactive systems with the Information Model, one might write some code like user.emotion = happy. Whereas with the Interaction Model, I think people are still trying to figure out how to design systems with that model, so I just put in some question marks.

Freaky, the alien larvae baby

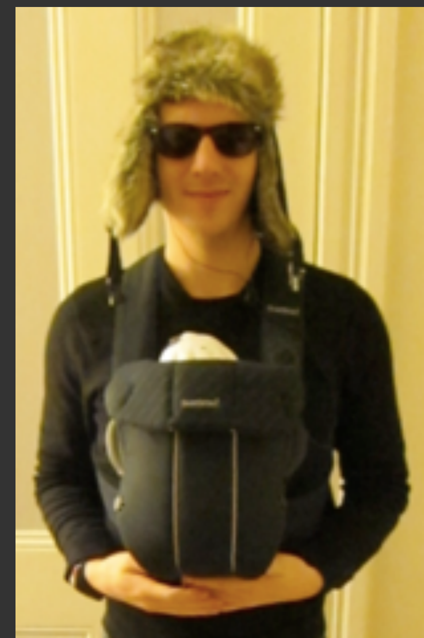

machine learning to detect fear
computer performing fear by "freaking out"
hybrid human-machine emotion

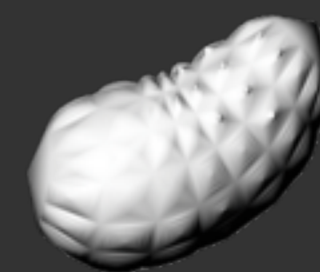

Lucian Leahu and Phoebe Sengers. 2014. Freaky: performing hybrid human-machine emotion. *DIS'14*

One example that gets away from purely detecting emotions in the user is Freaky. Freaky is an alien larvae like creature worn in a baby carrier. It uses machine learning to detect fear in its person. When it detects fear, it "freaks out" by making noises and vibrating. Its person has to pet and rock Freaky to get it to calm down.

I mean, the form factor is obviously weird, but I think that serves to show that this is an "alien" machine interpretation rather than claiming that this is the "true value" for the human's current emotions. So the system accommodates both machine interpretation and human interpretation, rather than claiming they are the same.

So, I think it could be argued that Freaky is a hybrid approach. The information-centric part is that the alien larvae does make its own judgment about "fear" based on its informational measurements of the wearer's heart rate. But, rather than assigning that label of "fear" to the wearer, Freaky enacts that fear itself. Much of the meaning and interpretation comes out in the wearer's interaction with Freaky, and with other people who are around at the time.

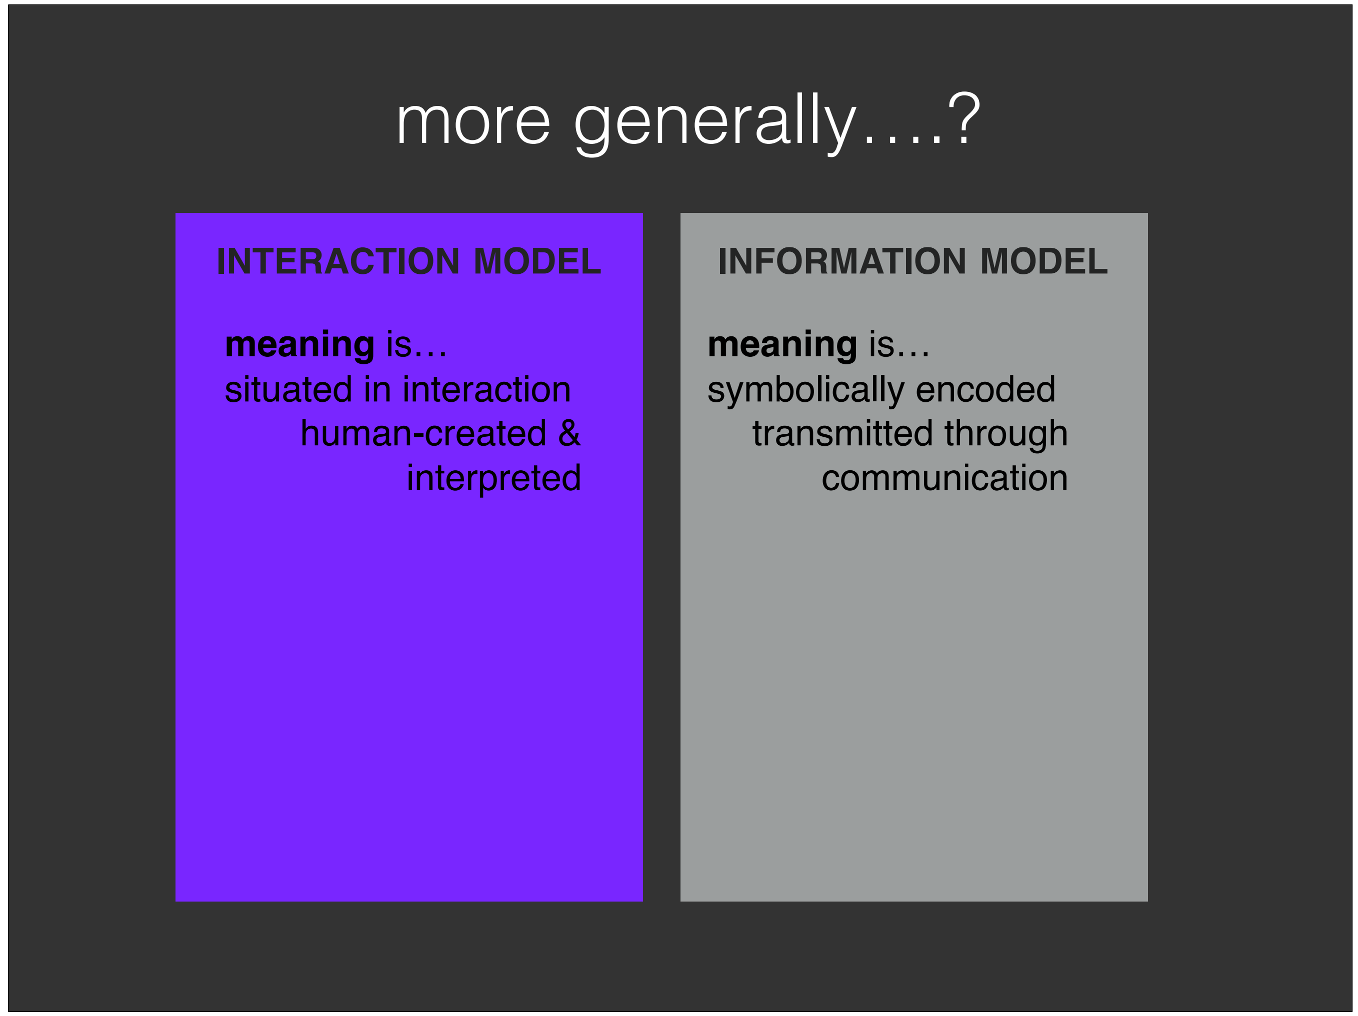

OK so here I want to take a big step back and propose something more general. This is really high level and somewhat tentative so some of the details still need to be worked out. So, work with me on this - feel free to propose refinements or critiques.

More generally looking at an interaction-centric approach versus an information-centric approach, I think a lot of it comes down to where **meaning** is supposed to be. In the interaction model, meaning comes out through interactions. Humans create and interpret meaning while interacting with each other. Or, even if we're alone, the meanings we think about are inextricable from our social upbringing.

In the information-centric approach, true **meaning** is in the symbolically encoded message. Communication is about transmitting that pre-existing symbolic meaning.

Shannon’s theory of information

- transmit signal over noisy channel
- want to reconstruct the original signal
- meaning = signal = series of symbols

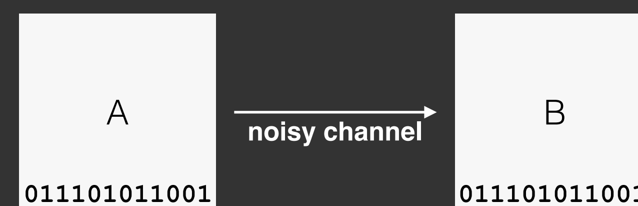

We can trace these ideas of an “information” centric approach to meaning at least all the way back to Shannon’s theory of information. It’s a mathematical theory focused on transmitting a signal from point A to point B over a noisy channel, and wanting to reconstruct the original signal at the other end.

political affiliation from Twitter

INTERACTION MODEL

political affiliation is…
performed in interaction
changes in context

study how people talk
about politics with
different people (friends,
family, coworkers, etc)

INFORMATION MODEL

political affiliation is…
either D or R
inherent to the individual

from tweets (noisy signal)
try to infer D or R

We can trace these ideas of an “information” centric approach to meaning at least all the way back to Shannon’s theory of information. It’s a mathematical theory focused on transmitting a signal from point A to point B over a noisy channel, and wanting to reconstruct the original signal at the other end.

semantic validity

"the degree to which analytical categories accurately describe **meanings** and uses in the chosen context"

Krippendorff (2004), "Validity," Content Analysis.

So, why does any of this matter?

Well, it has huge effects on semantic validity. If we think that meaning is situated in interaction but we go around measuring it like the true meaning exists in symbols, then that's pretty poor semantic validity. Or vice versa.

measurement

- need to measure something to study it
- we choose what is meaningful to measure
- system of measurement creates object of study

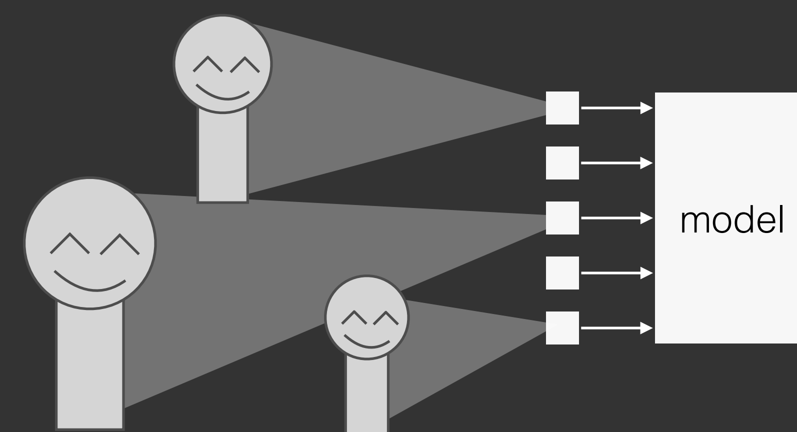

A lot of this also really comes down to decisions around measurement. Data science relies on having data as input, and as we know there is no such thing as raw data - data is not extracted, data is created, and in what we choose to measure and what not to measure, our system of measurement creates our object of study for data science. For example, if our model is looking at individual people as instances, then those people necessarily get reduced down to a vector of things that were measured about those people.

I would say that most data science tends to draw from an information-centric approach, because it's based on

Where else do you see information vs. interaction approaches?

How can we design models for data science that employ an interaction or hybrid approach?

end