

Stefanie Tellex

Week 3

First, I'd like to complain about cat rubbing being non-social. Just because rubbing can also be used to mark territory doesn't mean it's not a social behavior. My family's cat rubs against us in a totally different way than she does furniture or other objects in her territory. She does it for longer, comes back for multiple applications and meows and purrs. Similarly, a socialized cat will come running up to you if you call it, even if it has never seen you before, and rub against you, roll on the ground, and purr. I've never seen a cat react in that way to an inanimate object. (They will get quite happy about catnip, but that ties into pheromones.) Cats also rub each other when they greet each other. I think it's likely that the rubbing behavior was co-opted from territorial marking, but it's clear from my experience that a cat doing it to humans is doing it for social reasons.

Also, I thought of a more far-term ethical concern. Humans have a taste for sugar and fat because in the hunter gatherer environment it was adaptive to eat as much of them as you could get and you often couldn't get enough. Now, in the modern world, this taste has become a liability because we can indulge as much as we want, leading to obesity, tooth decay, and other health problems. What if, in building relational computers, we discover the "sugar" of relationships - an interaction that humans crave, but that leads to unhealthy behavior if it's fully indulged? For example, what if a successful relational machine becomes **better** than a human friend? What if it satisfies your need for a relationship better than any human, so you no longer need to talk to humans at all? I think it's interesting to try to imagine what the "sugar" of relationships might be...

J. Harvey and S. Hendrick. Handbook of Personal Relationships: Theory, research, and interventions, chapter Self-Report Methods in Studying Personal Relationships

They list three reasons for doing self reports. I think those reasons apply to a much lesser extent in human-machine relationships.

- "Participant verbalizations about subjective experiences are the only currently available means for investigators to tap directly such covert activities as perceptions, feelings, thoughts, expectations, and memories." While this is true of the human participant, presumably the machine can log whatever the investigator deems important about its internal state at all times in the relationship.
- "Participants can describe events that, though overt, are usually private to the relationship." The machine can record everything, and if you have the user agree before hand, the researcher can access the entire interaction. There are still issues of information retrieval and data mining.

- "Participant reports are relatively easy to obtain and involve much less inconvenience and expense than observer reports." Again, the machine can record everything, at the cost of disk space.

In human-machine relationships, issues like frequency and length of interactions will be known - the machine can easily record it. The advantage of self reports comes from aggregating the data and sorting through the raw experience in order to tell a higher level story about the relationship. Currently machines are not good at this task, although we might hope they will be some day. I suspect that human-machine relationships might be an interesting avenue to use to validate some of the self-report data. You could for example compare the human's self-reported frequency of contact to the actual length of interaction. There is some research into classifying a human voice based on emotions - you could try to quantify the emotional state of the user during the interaction and see if it matches their self report.

D. Kenny. Handbook of Personal Relationships: Theory, research, and interventions, chapter The Analysis of Data from Two-Person Relationships,

I'm not sure if intra-class correlation is relevant to us because in our case, the dyad members are so distinguishable. It would be cool and impressive if you could show that some measures were independent of human vs. machine - that in itself would be an exciting result. Also, probably the machine is identical in all cases, so it's misleading to treat the machine as a unique entity in each pair if it's the same machine across all pairs. Many measures probably can't even be made against the machine (yet).

So for human-machine relationships, it's probably completely within-dyad. You might want to compare human-machine to human-human, but then I think you would use ANOVA to ask if there was a significant difference in the two classes. If you were doing that, you run into a problem in the human-human side, so I supposed in that case you would want to compute the interclass correlation in order to compare the two groups. I think that for most machines (until they pass the Turing test), would have to be compared to some specific type of human relationship - like parent-child, human-pet, or human-therapist. In that case there's a natural division between the dyads. I can't think of many cases where you would want a machine that you would compare to a symmetrical human relationship like friendship. (I can't think of many symmetrical human relationships besides friendship (romantic or not).)

L. R. Wheelless and J. Grotz. The measurement of trust and its relationship to self-disclosure

I'm not sure how you validate a study like this? They defined a measure and decided that it measured trust based on the survey results I guess? It seems like you should have some independent measure of trust that you can show correlates with this measure? How do they know the effect isn't an artifact of the way they wrote the survey? I guess I feel like they defined trust in a vague way, and I don't think they showed that what they were

measuring is trust; it's just based on the questions in the survey which we don't even see...?

I think it makes more sense to measure individual trust in terms of actions - to try to quantify how much power you are willing to give the target over you? Self-disclosure is one way of measuring this, and asking counterfactual questions: "If you committed murder, would you ask this person for help escaping?" "If you needed a place to sleep, would this person let you?" Maybe you could quantify the cost to help, and the benefit to helping by asking people to rate them independently, and then use a survey like that to compute the amount a person believes the target will risk for them. This view is more in line with the Bickmore work we read last week, where a relationship is defined in terms of the ways members of the dyad expect it to help them.

Adam O. Horvath and Leslie S. Greenberg. Development and Validation of the Working Alliance Inventory.

I think this paper is much better than the Wheelless paper. It presents the instrument they used, and it validates it against other measures.

I'm not too impressed by the expert question filtering process; it seems like it might easily be biased by whatever fad is currently popular in the field.

I don't understand a lot of the math in both of these papers - "reliability estimate", "estimated alpha", etc.