Sajid Sadi

Week 2

Steve Duck. Interaction and Daily Life in Long-term Relationships. Chapter from Human Relationships.

In relation to this reading, I am mostly expressing the issues I find important in relation to designing systems that can relate with human beings. Of course, there are many other considerations as well, but I am prioritizing these thoughts. Additionally, I realize that some of this might sound like pessimist, but is meant more as problem definition. To begin, I think one of the most important messages of this reading, which is pointed out time and again, is that relationships are in fact a changing holistic phenomenon. There is some level of inherent instability within any relationship, and this instability is not merely present in the emotional aspect of a relationship, but within the social, communicational, and mental model states as well. This of course presents in some ways a challenge for those of us who intend to make machines relational, since machines are, often by design and definition, deterministic. Additionally, even when the machine is designed to act non-deterministically (at least as perceived by the user, perhaps by using the user's own non-deterministic behavior as input), there exists a potential that this non-determinism may exceed the thresholds for the level of certainly that the user expects of the system. Here, of course, I use the term user to specify a person in a relationship with a machine, though a more nebulous but perhaps more meaningful term would be "partner" or "agent".

Further along, the discussion of "snap decisions" brings forth another obstacle to systems that try to participate in relationships. While a computer generally presumes the start of a relationship as the first encounter, in fact the person is generally biased long before that time by preconceptions and hearsay about what to expect of machines, whether those facts be true of not. Additionally, the idea of social context in relationship, discussed later in the chapter, also plays into the idea of having a relationship with a machine, since generally speaking such a relationship has a potential to be just as or perhaps even more demanding than a human-to-human relationship, but the latter is generally socially perceived to be of higher rank and importance. Moreover, because the user defines the start of the relationship differently and generally never starts off in a "clean slate" state, the system must, in some ways, must meet the user's expectations in order to establish a relationship. This "pre-tainting" of an individual's perception is described as a serious factor in whether a relationship forms or not in human-human encounters, and the same obstacle faces designers of relational systems as well, though perhaps the stakes are much higher in the latter case.

This also leads to another issue regarding the disparity between what the machine and the user may consider the beginning of a relationship: the consideration of context. Context is, of course, greatly important in any interaction. Since we tend to never fully specify the context of an exchange within the exchange itself, in some sense the machine starts off at a serious disadvantage in the beginning of the interaction, because it generally lacks a lot

of the context of the introduction itself. A system that knows little starts off learning the context from the user (thus providing a negative initial impression), while an expert system starts of knowing context about some particular subject, and generally that subject is not the context of the initial interchange in a freeform environment. However, the latter does provide an advantage, because disseminating the fact that the system is an expert in a particular topic automatically biases the user towards that context, thus mitigating the context issue to some extent. However, I realize that such a bias also reduces the possibility of a complete relationship with the system, since the user then also gains a predilection to considering the system incapable of more than that initial exchange.

Some of the other issues that perhaps pertain less to my research work than those of others also caught my eye. In particular, the concept of strategy in relationships poses an interesting problem, both in that a computer has perhaps less self disclosure to itself than its human counterpart, and in the fact that such strategy seems to require a high level of understanding of the holistic context and content of relational exchanges than we seem to be able to muster up at this time. On a high note, it is noted that the most important relationships are perhaps the ones we do not categorize as such, and this offers perhaps the greatest inroads to relational systems. Lastly, I would hope that we would not have to enter into the dissolution aspects of relationships in the same context as the book, and at any rate, perhaps it is easier for people if the machines took the blame, at least for the time being.

Glyn M. Collis and June McNicholas. A Theoretical Basis for Health Benefits of Pet Ownership: Attachment Versus Psychological Support.

I think the most important thing I took away from this reading is the fact that it certainly seems to be a rich area of investigation, particularly because the theories remain, according to the reading, in large part unconfirmed. Perhaps one should not try too hard to get a relational system to become a pet of sorts, but at the same time, despite the lack of hard conclusions, it appears that there are subjective (and perhaps objective) benefits to having a relationship with a pet. In the initial sections, it is pointed out that the guiding assumption is that humans deal with pets via the same pathways used for human-human interactions. This seems particularly promising as an inroad for relational systems, since it allows for a lower-expectation pathway into the realm of social acceptance than a direct leap to human-human type relationships seems to allow.

It seems particularly important to me to note that the commonly held notion that the relationship to a pet is similar to a relationship with a child does not hold water at levels beyond the superficial. In particular, the attachment of a child to an adult is reversed with an adult and a pet. Even if the internal workings model of attachment is used, it seems that only forms an archetype for future attachment and affection patterns, which in turn suggests to me that perhaps the level of intelligence embodied in a child is not necessarily as important as feedback that activates that archetype in forming an attachment with the user. On the other hand, the social support theory seems to provide some interesting ideas to explore. In particular, the idea that an opportunity for nurturance is perhaps a good access point for the rest of the support network is a good starting point. Additionally, the

definition of companionship, as well as the lack of need for social competence, put lower demands on a system than would the full social interactions required for the same efficacy.

Tim Bickmore. Theory.

I found this reading interesting because of the sort of mathematically rigorous treatment it provides for the subject. In many ways it reflects the same conclusions that Duck arrived to via a different path and context. The social exchange interpretation of relationships also makes an interesting read, though it is perhaps a bit personally depressing to think of relationships in such way. In particular, it is pointed out that the model does need several additions to accommodate ideas such as multiple distributed goals within an interaction, accommodation, score keeping, and trust. It is interesting to see all these woven together in a theoretical model, though perhaps it is by itself a little harsh (and demanding of the understanding of computers today) for direct translation into reality. I think an interesting point made here is that trust essentially requires communication, because it is insufficient to simply accommodate another. The agent being accommodated must understand this accommodation in order to be able to trust the agent who helped out. This is, to me, a roundabout way of saying that trust only exists in the presence of transparency, though sadly the reading does not go into a way of providing that transparency that would be very friendly to human agents (ie, not simply doing so probabilistically).