Amy Eastment

Week 4

This paper makes note of how the current "bottle-neck" of intelligent computing is in the need for more common sense knowledge and the development of heuristics to manage it, but after that, we only see mention of the common sense side of things. I guess I would have liked to have seen more about what sort of heuristics are being developed for machines to make use of their "common sense."

+ While I can see the usefulness in teaching a camera to understand how sets of pictures might be related, feeding it "common sense" from the OMCS sounds like a bit much. I would think one would rarely want to apply the little snippets of common sense to annotating their photos…most people just annotate names or locations – something short to jog their memory - and not full phrases about what was going on at the time.

+ The common-sense storytelling applications sound really interesting! (In particular, the OMAdventure sounds like a really clever way to train the system). But how long will it be until the "storytelling" capabilities exceed the capacity of a child, and what sorts of augmentations could be done to help it out? What if one was doing textual searches instead of image searches to throw more complicated sentences or different topics into a story?

+ The less-ambitious common sense attempts (textual-completion, the CSDJ) sound like they have been most successful; then again, they aren't relying on a large database of common-sense, just simpler algorithms. + The Reformulator and Goose sound like really good, really practical applications of common sense computing – too few people really know how to phrase their searches well, and even fewer know to keep trying different things until they have succeeded. If these systems were set up to chain on their own (each time the reader clicked on their suggestion, it offered another related suggestion with the results of the next search), I could see it being quite useful for a novice-computer-user.

+ After reading the textual-affect sensing info further down, I'm curious to see how effective their EmpathyBuddy? system really is. And while it does sound like an entertaining application, I'm left wondering what more-practical applications it might have? It also sounds very anthropomorphized – a side effect that I don't think they were intending. It makes me wonder, when we continue to add more and more of this common-sense intelligence to applications, at what point do they start to seem like (irritatingly?) like people?

+ While common sense does sound like an interesting pursuit in AI, it also sounds like a long and potentially futile one. As the paper states – it falls flat in applications that aren't "fail-soft" and seems to work best in the "pushing" of random, helpful information. At the same time, if we want truly intelligent machines, we need to figure out a way to handle the direct "pulling" of information; forming these methods is a non-trivial task.