

Intelligent Agents for Helping Humanity Reach Its Full Potential

Jason Hong (jasonh@cs.cmu.edu), Human-Computer Interaction Institute

Introducing Maslow

Within fifty years, we will build and deploy highly personalized intelligent agents that can help us find, set, and meet hard goals to improve our lives in meaningful ways that we choose. Think of it as a cross between a lifelong coach, a caring uncle, and an honest and supportive friend. Or, if you are into science fiction, consider it as a combination of Samantha in the movie *Her*, the Young Lady's Primer from the book *The Diamond Age*, and Minds from Iain Bank's *The Culture* series. Let's call this agent Maslow.

Healthcare is a clear case where humanity needs significant help in achieving hard goals. In the early 20th century, the primary global health issue was controlling infectious diseases like tuberculosis and diphtheria. Now, in the early 21st century, the primary global health issue is managing chronic conditions, which are ongoing conditions (like heart disease, diabetes, and depression) that require long-term and sustained changes in people's behaviors in terms of diet, exercise, and medication. The World Health Organization estimates that 60% of all deaths worldwide are now due to chronic conditions [1].

Let's say that you wanted help in being healthier. Maslow examines your lifelogs (video, photos, call logs, emails, location history, sensor data) to analyze your existing behaviors, draws on the vast data repositories of what others did and what their outcomes were, finds people who are most similar to you and analyzes the changes that they successfully made, and incorporates your personal history to personalize the interventions. If you like gamification, it will offer interventions in the form of points and achievements. If Maslow found that social proof (following what you believe everybody else is doing) is more likely to work for you, it might suggest that you chat with Alice down in accounting, since she has recently started a health program that she really enjoys.

Maslow also doles out interventions in small chunks, with personalized recommendations that are just within your grasp. It reinforces these behaviors over a long period of time until you have mastered the new habits necessary to sustain your goal. For example, it might initially suggest linking a new desired behavior (e.g. doing light exercise) with existing habits (e.g. watching television), and slowly nudge you over time to more intense kinds of exercise.

Furthermore, given that many of these changes depend on one's social and environmental context, Maslow also interacts with other people's agents to suggest small and subtle changes. Examples might include offering your spouse coupons that lead to healthier food in the house, notifying you about a batch of fresh Honeycrisp apples at the farmer's market since those are our favorites, or recommending to a friend that they post a positive message about their improvements in health to their social media account (which Maslow will be sure to find and highlight for you).

The scenarios go well beyond health as well. What if we decided we wanted to be more green? Or wanted to learn painting? Maslow might even help us find compelling new goals to set for ourselves, in forms that are fun and engaging. A young professional might want to be more like a trusted mentor, while a young teen might set a goal to be more like a popular celebrity. Maslow could incorporate deep ideas from psychology to help motivate us and sustain changes in our behavior, all while ensuring that the interventions we get are commensurate with our ability and level of motivation.

Maslow will also be used in conjunction with help from professionals. For example, if a person is facing depression, he would work with his therapist on identifying positive reinforcing activities (a common technique in Behavioral Activation Therapy [2]). His agent could then help track his goals, offer timely

interventions, and help measure his change in mood over time. He could then use Maslow to work with his therapist in reviewing the data, adjusting the interventions, and setting new goals.

Will Maslow be Possible in Fifty Years?

There are three reasons why I believe Maslow is tractable. First, there will be exabytes of sensor data, video data, and communication data about human behavior, some of it in publicly visible forums, some of it gathered by corporations, most of it in our private data repositories. This data will allow us to build incredibly detailed models of human behavior at a scale and fidelity that has never before been possible. In particular, today's smartphones are already the most intimate computing devices ever created. They already know where we go, who we know, and a rough sense of what we do, and are with us almost the entire day, from when we wake to when we sleep. The future equivalent of smartphones will also know what we eat, how we feel, and what our tasks are.

Second, we will have new algorithms for sifting through all of this data, allowing us to reliably extract low-level human behaviors, e.g. opening pill box, getting water, and putting pill into mouth. These low-level behaviors will then be synthesized into high-level tasks, e.g. taking morning medication. These high-level tasks will in turn be linked to goals, e.g. addressing my heart condition, as well as short- and long-term outcomes, e.g. this person has more physical activity over his baseline and is happier than he was one year ago.

Third, instances of Maslow can continuously evaluate the effect of small interventions by conducting large-scale A/B tests and measuring the results. Companies today are already using these kinds of tests to tweak home pages so as to optimize click-through rates, revenues, mood [3], and even cybersecurity behaviors [4]. There are also lines of research looking at how to balance between exploration / exploitation of choices. It is not too much of a stretch to imagine that these techniques can be extended, allowing agents to continually test the link between interventions and behaviors, and how those behaviors lead to better outcomes. What kinds of reinforcement work best for which kinds of people? Are intrinsic or extrinsic motivations more effective for this kind of intervention? How should this class of interventions be spaced over time for optimal results? For people with this set of physical, social, and environmental characteristics, what kinds of approaches have worked best?

Maslow also offers the potential to data mine and more quickly disseminate best practices and lifehacks to people who can make use of them most. That is, once one person finds a potentially better way to do something, the agents could selectively share the practice with a few others, rigorously test its effectiveness in practice, and choose to diffuse it further if it really works. Success here would lead to new ways of speeding up the diffusion of innovations and continually improving performance, within an organization, across an entire community, or with the public at large.

Research Challenges in Building Maslow

There are several major technical challenges in building Maslow. The most straightforward is better sensors and better context models for representing people's goals, activities, motivations, and social relationships, as well as the physical environment. Today, we only have the basic primitives, and will need far better hardware and models that are accurate, robust to errors, and low-power. Maslow will also require significant advances in data mining to cull meaningful patterns from large heterogeneous data sets, to link behaviors to outcomes, to evaluate the effectiveness of these behaviors, and to semi-automatically run a large series of A/B tests.

More difficult, however, are the social, economic, and ethical challenges. For example, today, data is highly siloed across individual web services, apps, and devices. Making this data accessible will require companies to offer new kinds of data sharing APIs, and it's not clear if there is an incentive to do so.

Maslow also poses new kinds of thorny privacy and security challenges. To address these, we will need many kinds of significant technical advances, including new ways of anonymously sharing and data mining sensitive personal data, simpler user interfaces for blocking undesired access (e.g. from nosy family members, overzealous marketing, government surveillance), and new methods for preventing people from manipulating interventions in their favor (like black hat search engine optimization organizations do today).

Lastly, Maslow raises numerous questions about freedom, choice, and agency. What kinds of interventions are acceptable? How far can and should these interventions go? Would people blindly follow offered advice uncritically (in the same way that some people blindly follow their GPS)? Would society end up homogeneous, as Maslow pushes us all to what is deemed "normal"? Might Maslow also make life too easy, taking away the challenge and joy of discovering things on one's own?

While Maslow is intended to help us achieve hard goals, the deeper underlying idea comes from our intelligent agent's namesake. The psychologist Abraham Maslow believed that one source of true happiness comes from achieving our full potential and creativity, which requires overcoming our fears, our inertia, and self-imposed limitations. As such, I believe something like Maslow is not only possible, but is something that should also be done for the betterment of humanity.

References

1. World Health Organization, Chronic diseases and health promotion. Accessed July 22, 2015. <http://www.who.int/chp/en/>
2. Portland Depression Treatment Program, Behavioral Activation Treatment of Depression. Accessed July 22, 2015. <http://www.portlanddepressiontreatment.com/behavioral-activation-treatment-of-depression/>
3. Kramer, A.D.I., J.E. Guillory, J. Hancock. Experimental Evidence of Massive-Scale Emotional Contagion through Social Networks. *Proceedings of the National Academy of Sciences*, 111(24), 8788-8790. 2014.
4. Das, S., A.D.I. Kramer, L.A. Dabbish, J.I. Hong. Increasing Security Sensitivity with Social Proof: A Large-scale Experimental Confirmation. *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 2014.