Available online at www.sciencedirect.com
SciVerse ScienceDirect
Int. J. Human-Computer Studies 71 (2013) 191–194
International Journal of Human-Computer Studies
www.elsevier.com/locate/ijhcs

Nature, nurture, and knowledge acquisition

Thomas R. Gruber

Tomgruber.org, Emerald Hills, California, United States

Accepted 25 September 2012
Available online 23 October 2012

Abstract

The nature vs. nurture dualism has framed the modern conversation in biology and psychology. There is an analogous distinction for knowledge acquisition and artificial intelligence. In the context of building intelligent systems, nature means acquiring knowledge by being programmed or modeled that way. Nurture means acquiring knowledge by machine learning from data and information in the world. This paper develops the nature/nurture analogy in light of the history of knowledge acquisition, the current state of the art, and the future of intelligent machines learning from human knowledge.

Keywords: Nature vs. nurture; Artificial intelligence; Knowledge acquisition; Inductive bias; Epigenetics of knowledge

1. Introduction

In his epic contribution to this issue (Gaines, 2012), Brian Gaines places the history and trajectory of the field of knowledge acquisition (KA) research in an appropriately broad context. Knowledge is a heady topic, even when studied in the embodied forms of media, expert systems, and the web. This community has always embraced the theoretical and practical together, drawing insights from history, philosophy, mathematics and other disciplines as well as the shiny objects of modern engineering. Brian offers a beautifully woven tapestry of these connections, placing the contributions of knowledge acquisition research in the context of the evolution of knowledge media and representation, information technology and artificial intelligence, and the collective body of human knowledge available online.

Within that contextual fabric, I see a familiar pattern, and think it is not pareidolia. Squint hard, and ask yourself: in what field have we seen the following progression before?

- In the early days, most people thought it was magic. The established authorities said that it was supernatural, that it comes from some kind of divine or universal source.
- A few people could work the magic, and the people were impressed. The magic workers got status and power and could start charging for tickets.
- Then others wanted the magic, and started asking questions that the magic workers could not answer. So they started inventing their own tricks, and since they were not in power, they were willing to share them.
- With this new information available, others started to notice how some tricks worked. They made theories.
- At first, the main product of theory-making was a good story, and an audience developed for listening to good stories. It helped make sense of the magic, and it was entertaining.
- The audience for stories naturally diversified into groups, each of which is drawn to hear stories that confirm their own identity and mindset. Audiences who believe that things are predetermined liked the stories about the magic coming from fundamental, universal structures. Audiences who believe that things are created by work liked the stories about the magic coming from the practiced craft of the magician.
- Meanwhile, people kept inventing tricks, and the stories and audience clusters helped spread the most impressive tricks, so that practices emerged. Each theoretical camp explained the successes with their best stories. The inventors liked the praise, and started using the theories to invent new tricks.

E-mail address: pubs@tomgruber.org

1071-5819/$ - see front matter © 2012 Published by Elsevier Ltd.
http://dx.doi.org/10.1016/j.ijhcs.2012.10.004

- With the synergy of theory and practice, inventing became engineering. The structural theories gave engineers building blocks with which to build impressive artifacts. The behavioral theories gave engineers a way to measure and analyze, to know how things behave in practice.
- And with time, very impressive things were created.

Brian gives a more elegant framework for this kind of progression, applied to the history of computer technology, information science, knowledge science, and beyond. What struck me was how this seems to happen in other fields.

2. Biology, Psychology, and AI

Consider biology and medicine. In the early days, people thought life and animate activity came from the gods or the equivalent in nature. A few people were good at creating or destroying life and influencing behavior, and they got to be priests, shaman, warlords, etc. Other people got tired of paying dues and figured out how to raise crops and animals to do medicine on their own. A lot of them died. (A lot of everybody died.) The independent thinkers started noticing patterns, some of them intermarried with practitioners, and camps of story/practice/knowledge transmission emerged (think Ayurvedic, Chinese, and European traditions in medicine). As best practices started to improve the quality of life, science and engineering emerged, and knowledge became power. Systems of knowledge acquisition and dissemination that worked better got more power. The structural theories of biological mechanisms such as germs and genes led to breakthrough therapies. At the same time, evidence-based medicine helped reward the practices that worked, even when not understood. Today, we have both systems working together. The combination has produced impressive artifacts, like the human genome database and therapies constructed by molecular engineering.

I see the same pattern in the history of psychology and cognitive science. Fast forward from early ideas about human nature and the craft of illusionists to the emergence of theories that begin to predict and prescribe. Social scientists got better at devising experiments; therapists got better at applying results. Again, camps of story/practice/knowledge transmission emerged. Structuralists explore genetic and biological determinants of behavior; behaviorists explore the nature of learning. With technology, we can now test structuralist theories of cognition in an MRI, and behavioralist theories on the web. Impressive indeed.

Now to bring this home, apply this to the field of artificial intelligence (AI) and knowledge science. 25 years ago, we were in awe of the magicians. It was amazing! Computer programs could play board games with people, diagnose infectious diseases, and talk as if they could see. We were amused by the appeal of ELIZA (Weizenbaum, 1966) and wanted our own professor's assistant from Apple's Knowledge Navigator video (1987[1]). We had high priests and grand theories, but they did not do the tricks. The tricks were where the action was, and the knowledge acquisition community asked a lot of questions. Is there a man behind the curtain in that demo? If it is an expert system, how do you model expertise? If knowledge is power, why is a large program more brittle than a small one? How do we make an engineering discipline out of knowledge representation? How do we make tools for building systems that can reason? What are the tricks for machines learning from people?

As in other fields, the dance of theory and practice emerged. Generic methods for representation and reasoning led to tools for the same, with corresponding architectures for ontology-based classification, task-based problem solving, and evidence-based inference. While mainstream AI conferences rewarded increasingly dry-erase whiteboard results, the culture of the KA community always stayed close to practice. As Brian recounts, the workshops demanded equal facility at writing a peer reviewed article for a journal and giving a demo under a poster. The culture of working systems and application-grounded research gave the KA community momentum and adaptability. When AI went into winter recess, we had sled races. When the web happened, we embraced it—even though the architects of HTTP, HTML, XML, and the rest did not know about "proper" knowledge representation. They created magic, and we wanted to see what we could do with it. Soon, the architect of the Web was proposing a *semantic* web (Berners-Lee et al., 2001), adopting the tools of our trade. And AI people started building programs to learn from the "collected intelligence" of the web.

Where are we today in this developmental progression? Has the field of artificial intelligence and knowledge science created impressive things? Well, we have Watson[2] – and Siri.[3]

Watson is a master question–answering machine, better than most humans at answering factual questions. Watson does not "know" more than people. Rather, I think of Watson as a brilliant research librarian; it is an expert at understanding questions and finding potential answers from an array of sources. Although it is tuned for performance in the game of Jeopardy, its architecture is capable of absorbing content from a large variety of sources and adapting to new types of questions (Ferrucci et al., 2010). Watson demonstrates convincingly that intelligent performance can be achieved by harvesting human knowledge in written form.

Siri is a virtual personal assistant, realizing a longstanding dream of AI as a consumer product used by millions of people every day. If Watson is a research librarian, Siri is like a personal secretary or concierge: it helps users solve tasks such as making a schedule or reserving a restaurant. Siri (as

[1] http://www.youtube.com/watch?v=-jiBLQyUi38&feature=gv&hl=en

[2] http://www.ibmwatson.com/

[3] http://siri.com