MAR15-2014-020338

Abstract for an Invited Paper
for the MAR15 Meeting of
the American Physical Society

Cognitive Computing: From Breakthroughs in the Lab to Applications in the Field

GURUDUTH BANAVAR, IBM TJ Watson Research Center

In the last decade, the availability of massive amounts of new data, and the development of new machine learning technologies, have augmented reasoning systems to give rise to a new class of computing systems. These "Cognitive Systems" learn from data, reason from models, and interact naturally with us, to perform complex tasks better than either humans or machines can do by themselves. In essence, cognitive systems help us penetrate the complexity of big data, reason using rich models, and enable each of us to perform like the best. We believe this will transform every industry for the better. Artificial Intelligence (AI) has a rich history, dating back to 1950's. The desire to build a Turing machine that would imitate a human dominated the psyche of the AI researchers for many decades [1]. The complex reality of human beings and the limitations of early symbolic approaches narrowed the success of early AI technology to a few specialized and small-scale applications [2, 3]. Over the last decade, new kinds of unstructured data from social networks, streaming data, and online publications, as well as massive data emitted from sensors from the physical world have outpaced traditional forms of structured data. This will continue to grow exponentially. The insights embedded in this massive amount of data can provide unprecedented opportunities for business and social value. Data has indeed become one of our most precious resources, and with its accelerated pace of evolution, it will determine the future trajectory of business and society [4]. New tools are being developed to extract insights out of the big data, which is abundant, unstructured, noisy, and unreliable. These tools have not relied on the same techniques that helped us exploit clean, structured data of the past, using small-scale models of the world and explicitly specified reasoning mechanisms. The new tools are using more automated statistical pattern-matching techniques, called machine learning, that have come of age in the last decade [ref]. In addition to reasoning from explicitly specified models of the world, these new machine learning techniques have given rise to a new class of systems that effectively learn from patterns in big data, and simultaneously augment their world models. Such systems can also interact naturally with us, on human terms, through natural language (i.e., unstructured text data), speech (i.e., unstructured audio data), vision (i.e., unstructured video data), and other modalities. We call this emerging class of systems that reason, learn, and interact naturally with us "Cognitive Systems". IBM's Watson is a family of cognitive systems targeted to a variety of domains. The first Watson system was capable of answering factoid questions as effectively as the best professionals in that field (as demonstrated by the Jeopardy! exhibition match, see illustration). Follow-on systems answer other types of questions, e.g., those that require passage answers, as well as other domains, e.g., healthcare, insurance, and education. Yet other cognitive systems in the Watson family go beyond question answering to support discovery of insights hidden in big data, such as in huge repositories of scientific literature, reasoning with evidence to support or refute topics of discussion, and to go beyond textual data to images and videos.