

2027: Human vs Machine Literacy

UMASS
AMHERST

PROJECT
LITERACY

Commissioned by Pearson, founding partner of Project Literacy



Foreword from Project Literacy

Human literacy levels have stalled since 2000. At any time, this would be a cause for concern, when one in ten people worldwide – including 5 million British adults – still cannot read a road sign, a voting form or a medicine label.

But at a point in human evolution when technological progress has never been more fast-paced it becomes more damning, and the need for remedial action more urgent.

This report explores comparisons of human and machine literacy, suggesting that we are close to reaching a point where the smartphones in our pockets will be better at reading and writing than 758 million people alive today.

We need to upgrade our people in the same way that we upgrade our technology, in order to break the vicious cycle of illiteracy. While each new generation of smartphone improves, illiterate parents tend to have illiterate children – perpetuating a devastating cycle.

To break that cycle, we need to teach adults to read first, so they can pass it on to the next generation. [Project Literacy](#) wants to break the cycle of illiteracy so that by 2030 no child is at risk of illiteracy.

Technology is part of the solution. Project Literacy has proven that the capacity for technology to be a force for good against illiteracy through a number of initiatives. One project, with not-for-profit Worldreader, uses mobile technology to help parents in India read more to their children. Another partnership, with Unreasonable Group, has seen us develop the world's first accelerator program supporting start-ups in emerging economies that aim to help close the global illiteracy gap by 2030.

The pace of machine learning and technological change provides society with unprecedented opportunities including the chance to tackle the global humanitarian crisis of illiteracy.

Kate James
Chief Corporate Affairs & Global Marketing Officer, Pearson

Introduction

According to this brochure, why is it difficult for people to know if they have high blood pressure?

<p>Black adults has high blood pressure, according to a two-year survey by Public Health Service in the 1960's. Other studies show as many as one out of three Black adults has high blood pressure.</p> <p>High blood pressure is the most common chronic disease treated by practitioners in the Black community.</p> <p>More Black people die as a result of high blood pressure than any other disease.</p> <p>For every Black person who dies of sickle-cell anemia, at least 100 others die from the effects of high blood pressure.</p> <p>The rate of death from the effects of high blood pressure for Black people is nearly one and one-half times the rate for White people.</p>	<p>Checked regularly</p> <p>Unfortunately, high blood pressure is a silent killer and crippler. At least half of the people who have high blood pressure don't know it because symptoms usually are not present. The only way you can be sure is to have the doctor check your blood pressure. You should have your blood pressure checked at least once a year, especially if:</p> <p>(1) you are Black, (2) if you are over 40, (3) if members of your family or close relatives have had high blood pressure or the complications of high blood pressure (stroke, heart attack, or kidney disease), or (4) if you have frequent headaches, dizziness, or other symptoms that may occasionally be related to high blood pressure.</p> <p>Follow your doctor's instructions</p>
--	--

This test question appeared in the National Assessment of Adult Literacy, a large-scale study conducted by the National Center for Education Statistics (NCES) to assess adult literacy levels in the United States. The NCES found that 1 in 7 of U.S. adults performed at the "Below Basic" level of prose proficiency, on a four-point scale below the Basic, Intermediate, and Proficient categories. This group of people had a very difficult time answering this question, getting it correct only 11% of the time – as opposed to more literate adults who could answer with 88% accuracy (Kutner et al., 2005, 2007).¹

This question is a text search problem – the answer is stated in a plain manner with a fairly direct correspondence to wording of the question, but it has to be found among a larger amount of irrelevant text. It represents a basic reading skill that is necessary to participate in modern society; for example, people need to understand medicine instructions, voting ballots, bus schedules, product labels, etc. The NCES study found that below-basic literacy rates barely changed from 1992 to 2003, staying stubbornly high around 14%, corresponding to 30 million adults. Within that group, 7 million (3% of adults) were considered nonliterate,

¹ The original question was a bit harder than shown here, since it contained a larger full page.

struggling even to recognize words and letters (Baer et al. 2009). A more recent study found similar rates of illiteracy (OECD 2013).

While the past 25 years have seen disappointing progress in U.S. literacy, there have been huge gains in linguistic performance by a totally different type of actor – computers. Dramatic advances in natural language processing (Hirschberg and Manning, 2015) have led to the rise of *language technologies* like search engines and machine translation that “read” text and produce answers or translations that are useful for people. While these systems currently have a much shallower understanding of language than people do, they can already perform tasks similar to the simple text search task above – exceeding the abilities of millions of people who are nonliterate.

In popular discussion about advances in artificial intelligence, it is easy to be drawn into comparing computers against the, say, the very best chess or Go players that humanity has to offer. But human intelligence and information processing skills – such as reading and writing – are very unevenly distributed, and it is worth paying attention to the needs of millions of people whose abilities are weaker than average. In this review we consider the comparison between computational reading skills versus human literacy. While “machine reading” is not close to mastering the full nuances of human language and intelligence, it is rapidly advancing. Already, it exceeds the abilities of the nonliterate adult population, and within the next decade, it may achieve reading skills similar to young children – potentially surpassing the below basic literacy levels of 1 in 7 U.S. adults today.

The following sections detail the relationship between what it means for humans versus machines to read, take stock of current machine capabilities in comparison to humans learners, and assess future prospects.

Language versus literacy

The NCES’ National Assessment of Adult Literacy studied several components of literacy, including quantitative literacy (i.e. arithmetic problems) and visual document literacy (i.e. signing forms or interpreting diagrams). We focus on *prose literacy*: the ability to interpret and comprehend written language such as sentences, essays, or narratives.

While humans can struggle with written language, we have an astonishing, innate ability for spoken language. Children can learn language without any formal educational instruction, understanding and producing a near-infinite variety of sentences by combining their meanings in different ways — a hallmark of human intelligence (e.g. Pinker 1994). But reading and writing are not natural skills for humans, requiring years of painstaking, hard work to learn to read with enough proficiency to be functionally literate in modern society.

Do machines exhibit prose literacy? In one sense, no, because computers do not deeply understand any type of language. It is easy for a human to fool a spoken dialogue system like Apple Siri or Amazon Echo, since they only understand a very shallow level of linguistic information from a person's utterance, and are pre-programmed to recognize specific types of requests. Deeper understanding of dialogues and text is a long-standing but incomplete goal of the field of natural language processing; for example, Etzioni et al. (2007) define *machine reading* as the autonomous understanding of text, without a software's creator carefully prescribing how the system should understand it.

On the other hand, as evidenced by humans' discrepancy between spoken versus written language ability, machines in some sense are strong at working with the textual, written form of language, despite their incomplete understanding of it. Written language is an explicit symbolic encoding of language, and computers are terrific at mechanical manipulation of sequences of symbols – characters, words, sentences, documents.

While humans start with a core of linguistic intelligence then build literacy on top, language technologies can instead start with the symbolic, written form of language.² In some sense, machines already “read” text in a shallow manner, which is incredibly useful since it can be performed at scale. For example, consider web search: it is difficult for humans to read through hundreds of millions of web pages to manually create a directory of them. But computer algorithms can automatically process all of them, fetching important keywords and phrases from these documents, indexing and further processing them in order to quickly find relevant web pages in response to a user's search query.

Web search is an extreme example to illustrate machine reading, since it is difficult to imagine human reading as an alternative. But there are many other domains where human reading has been traditionally widespread but machine reading is now outpacing humans. For example, *e-discovery* systems can replace the thousands of person-hours of document reading that lawyers routinely do, by helping search and filter massive numbers of documents. Newspaper articles warn “Software Does in Seconds What Took Lawyers 360,000 Hours.”³

² Spoken language technologies also exist, which listen to a person's speech and try to comprehend it. Typically, current technologies first map a person's speech to text before further processing.

³ *Bloomberg Markets*, February 28, 2017. Also: “Armies of Expensive Lawyers, Replaced by Cheaper Software,” *New York Times*, March 4, 2011.

Basic reading skills

Computers are skilled at the more basic, mechanical levels of reading at the level of words and phrases which (literate) humans typically learn as children. The National Institute for Literacy (Armbruster et al., 2001) outlines several major components of reading education at the kindergarten through third grade levels (ages 5 through 9). All but the last can be done fairly well by computers:

- *Phonemic awareness and phonics*: understanding how written letters and groups of letters correspond to sounds.
Computer performance: Computers are fairly good at this in the context of speech synthesis, especially for words they already know about. These steps are not necessary for direct machine reading of text.
- *Fluency*: recognizing words automatically when reading silently, and fluidly reading a sentence out loud, instead of reading one word at a time; this indicates the learner is starting to understand the structure of a sentence.
Computer performance: Computers are excellent at recognizing words, fairly good at recognizing grammatical phrases of text (with around 90% accuracy), and fairly good at reading aloud.
- *Vocabulary*: being able to recognize and understand a large set of words.
Computer performance: Computers are excellent at memorizing very large dictionaries of words, phrases, and the names of people, places, organizations, etc. – such databases within typical natural language processing systems often contain millions of such entries. In such a database, computers can even hold a rough gist of words' meanings and track what other words are similar to them, or are synonyms or otherwise related.
- *Text comprehension*: Understanding the meaning of a text.
Computer performance: Computers are not as good at this, but progressing. While it is difficult for a computer to understand things not explicitly stated in a text, some systems can translate between different ways to say the same thing, or find answers in a text if they are directly stated.

Instructional standards for children are relevant for thinking about adult literacy – millions of adults struggle with the most basic reading skills. For example, the NCES found that 3% of adult test takers were not able to complete any of the written tasks such as the one shown in the introduction. They instead administered a very basic oral examination to these nonliterate respondents. The questioner pointed to food packaging or other familiar materials, and asked questions such as:

- What letter is this? (*Letter reading*)

- Please read this [number] for me. (*Word reading*)
- Please point to the word "water." (*Word identification*)

These respondents (the lowest performing among all test takers) gave correct responses with low accuracy rates: 57% for letter reading, 46% for word reading, and 70% for word identification. This poor level of performance is representative of 3% of the U.S. adult population – some 7 million people (Baer et al. 2009, ch. 4).

Computers can also tackle harder reading tasks that are still difficult for adults who have better literacy skills than the above nonliterate group, but are still at a below basic level. For example, in the text search task in the introduction, in order to answer the question,

“Why is it difficult for people to know if they have high blood pressure?”

the reader must match it with this document text:

“... people who have high blood pressure don't know it because symptoms usually are not present.”

Here, the clause following “because” is a correct answer. Computationally, this text search problem can be partially solved with direct phrase matches (underlined) between question and document text. A little bit of understanding of sentence structure is necessary to recognize that the *because* clause can provide an answer to a *why* question; but this type of inference is well within the bounds of what current question-answering systems can do. Indeed, the IBM Watson Jeopardy-playing system performed more difficult versions of question-answering with great skill, often searching Wikipedia and related resources for answers to the posed questions (Ferrucci, 2012), and current web search engines like Google can sometimes find textual answers stated directly in a web page.

The future of machine literacy

Current research on computational question answering and reading comprehension is making substantial progress, though still below human performance. For example, algorithms can answer multiple choice questions from the New York Regents fourth-grade science examination with 75% accuracy (Clark and Etzioni, 2016); on another multiple choice dataset, current computer systems achieve 77.5% accuracy, versus at least 95% for humans (Sugawara et al., 2017). Computers have already progressed to match human performance in more specific tests like TOEFL synonyms or SAT word analogies in the past twenty years (Turney and Pantel, 2010; Landauer and Dumais, 1997). Accelerating advances in computational power, and refined techniques like deep

learning, are leading to fast-paced improvements in algorithms that process and compose language meaning.

One way to think about basic literacy versus machine reading is to consider what tasks people and computers can do in the world. A person without basic literacy may struggle to drive if they cannot easily read traffic signs like “Slow, children” or essential information like street addresses. By contrast, cars increasingly drive themselves. (While self-driving cars do not currently read road signs, it is certainly possible to build such capabilities.) Smartphones enable artificially intelligent language skills to be embedded in real-world situations – for example, a class of smartphone apps can recognize text from the camera and perform natural language processing on it, such as translating to another language (e.g. Google Translate) or reading it out loud (e.g. KNFB Reader).

Software that reads text aloud is a tremendously useful assistive technology, usually intended for blind or visually impaired users. But as we have seen, it in fact has stronger reading skills than millions of (non-visually impaired) American adults. This raises questions about how societal efforts ought to support efforts to improve artificial intelligence versus human language abilities. Surely, we should make stronger efforts to improve human literacy. But would it be helpful to use assistive reading technology for people with reading problems? Or is it dangerous to rely too heavily on such technology if it substitutes for learning literacy skills?

A more symbiotic relationship may exist in efforts to develop software to help teach language skills – for example, Duolingo for foreign language learning – where natural language processing technologies can help the software better adapt to the learner’s needs. Researchers have investigated how to construct intelligent reading tutors that help functionally illiterate adults learn to read and write (Ramachandran and Atkinson, 2008); related work has developed software to help children (Mostow and Aist, 2001) and high school students (McNamara et al., 2007) improve their reading abilities as well.

Conclusion

Rapid advances in machine reading technology serve as a surprising comparison point to the prevalence of nonliteracy and below-basic literacy levels among people. As natural language processing systems begin to outpace the reading skills of millions of people, it is likely that within the next ten years we will see everyday devices, such as our smartphones, take on capabilities that make them more literate than hundreds of millions of people across the world.

But illiteracy is not inevitable – literacy education dramatically enhances human abilities by endowing us with the ability to read, write, and fully participate in modern society. We often speak of investing in technology to enhance artificial

intelligence – but through better education, we can invest in *humans* to enhance *our* intelligence, and strive to achieve literacy for all people.

References

Bonnie B. Armbruster, Fran Lehr, Jean Osborn (2001). *The research building blocks for teaching children to read: Put reading first*. National Institute for Literacy. <https://lincs.ed.gov/publications/pdf/PRFbooklet.pdf>

Justin Baer, Mark Kutner, John Sabatini, Sheida White (2009). *Basic Reading Skills and the Literacy of the America's Least Literate Adults: Results from the 2003 National Assessment of Adult Literacy (NAAL) Supplemental Studies*. National Center for Education Statistics, NCES 2009481. <https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2009481>

Peter Clark, Oren Etzioni (2016). My Computer Is an Honor Student — But How Intelligent Is It? Standardized Tests as a Measure of AI. *AI Magazine*, 37(1) 5-12.

Oren Etzioni, Michele Banko, Michael Cafarella (2007). Machine Reading. *AAAI Spring Symposium*, AAAI Press.

David A. Ferrucci (2012). Introduction to “This is Watson.” *IBM Journal of Research and Development*, 56(3/4).

Julia Hirschberg, Christopher D. Manning (2015). Advances in natural language processing. *Science*, 349(6245):261-266. DOI: 10.1126/science.aaa8685.

Irwin S. Kirsch, Ann Jungeblut, Lynn Jenkins, Andrew Kolstad (1993). *Adult Literacy in America: A First Look at the Findings of the National Adult Literacy Survey*. National Center for Education Statistics, NCES 93275. <https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=93275>

Mark Kutner, Elizabeth Greenberg, Justin Baer (2005). *A First Look at the Literacy of America's Adults in the 21st Century*. National Center for Education Statistics, NCES 2006470. <https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2006470>

Mark Kutner, Elizabeth Greenberg, Ying Jin, Bridget Boyle, Yung-chen Hsu, Eric Dunleavy (2007). *Literacy in Everyday Life: Results from the 2003 National Assessment of Adult Literacy*. National Center for Education Statistics, NCES 2007480. <https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2007480>

Thomas K. Landauer, Susan T. Dumais (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2) 211-240.

Danielle S. McNamara, Yasuhiro Ozuru, Rachel Best, Tenaha O'Reilly (2007). The 4-pronged comprehension strategy framework. In McNamara, Danielle S. (Ed.), *Reading comprehension strategies: Theories, interventions, and technologies* (pp. 465-496). Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.

OECD (2013). *OECD Skills Outlook 2013: First Results from the Survey of Adult Skills*. OECD Publishing. <http://dx.doi.org/10.1787/9789264204256-en>

Steven Pinker (1994), *The Language Instinct*. Harper Perennial Modern Classics.

Sowmya Ramachandran, Robert Atkinson (2008). An Evaluation of Intelligent Reading Tutors. In: Woolf B.P., Aïmeur E., Nkambou R., Lajoie S. (Eds.), *Intelligent Tutoring Systems*. Lecture Notes in Computer Science, vol 5091. Springer.

Luz Rello, Horacio Saggion, Ricardo Baeza-Yates (Ed.) (2013). *Proceedings of the Second Workshop of Natural Language Processing for Improving Textual Accessibility*. Association for Computational Linguistics.
<http://www.aclweb.org/anthology/W13-15>

Saku Sugawara, Hikaru Yokono, Akiko Aizawa (2017). Prerequisite skills for reading comprehension: Multiperspective analysis of MCTest datasets and systems. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*.

Peter D. Turney, Patrick Pantel (2010). From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37 141-188.

About the Authors

Brendan O'Connor, Assistant Professor, University of Massachusetts, Amherst

Brendan O'Connor is an assistant professor in the College of Information and Computer Sciences at University of Massachusetts, Amherst, and researches in natural language processing and machine learning, especially when informed by or applied to cognitive and social science. Brendan received his PhD from Carnegie Mellon University's Machine Learning Department, and has previously been a Visiting Fellow at Harvard IQSS, and worked with the Facebook Data Science team.

About Project Literacy

Project Literacy is a global campaign founded and convened by Pearson to make significant and sustainable advances in the fight against illiteracy so that all people - regardless of geography, language, race, class, or gender – have the opportunity to fulfill their potential through the power of words.

Official partners include: Worldreader, Room to Read, GOOD Magazine, Reading Partners, BookAid, BookTrust, 826National, Jumpstart, First Book, Raising A Reader, Reading is Fundamental, Reach out and Read, Asia Foundation, We Need Diverse Books, Microsoft, the National Literacy Trust, the Bill, Hillary, and Chelsea Clinton Foundation, UNESCO, Unreasonable Group, CENPEC, Results for Development, the Center for Knowledge Societies, War Child, Lessons for Life, Doctors of the World, World Literacy Foundation, 28 Too Many, WomanKind, The Big Issue, The Hunger Project, Nutrition and Education International, and Veerni.

Learn more at ProjectLiteracy.com

About Pearson

Pearson is the world's learning company, with expertise in educational courseware and assessment, and a range of teaching and learning services powered by technology. Our mission is to help people make progress through access to better learning. We believe that learning opens up opportunities, creating fulfilling careers and better lives. For more, visit www.Pearson.com.