

Remembering the Past for Meaningful AI-D

Julie S. Weber¹ & Kentaro Toyama²

¹Computer Science & Engineering, University of Michigan

²School of Information, University of California, Berkeley

Abstract

This position paper describes how the nascent area of AI for development can learn from the challenges and successes of its parents: artificial intelligence and information and communication technologies for development (ICT4D). AI suffered from overly ambitious beginnings and years of stumbling before finding its footing, and achieving impactful ICT4D has been an equally challenging endeavor. We describe the history and challenges of both AI and ICT4D research, and present three broad suggestions for AI-for-development researchers: (1) that they spend as much time as possible with the kind of site or the organization they are hoping to impact; (2) that they be ambitious but humble in their goals and expectations; and (3) that they put AI in the service of existing, well-intended, competent development organizations.

Those who cannot remember the past are condemned to repeat it.

– *George Santayana (1863-1952)*

Introduction

What do the fields of artificial intelligence (AI) and information and communication technologies for development (ICT4D) have in common? They both overpromised and underdelivered, at least in their initial years.

The early visionaries of artificial intelligence research set out to solve the problem of creating human-level intelligence. In 1970, Marvin Minsky, for example, suggested that AI would meet the intelligence of the “average human” within three to eight years; he revised this to a slightly more conservative “one generation” in 1977. A generation has passed, and although AI has made tremendous strides – in search, machine learning, speech, natural language, etc. – there is not yet, for instance, a vision system that can today solve the problem of the apocryphal “vision summer” – identifying objects on a table with video imagery. Indeed, the entire field of computer vision, spun off from AI, does not appear likely to solve even just this one aspect of human-level intelligence anytime soon.

Similarly, early ICT4D work set goals such as “double the income of rural India,” perhaps using extremely affordable PCs. But, again, despite significant progress in, for example, the worldwide penetration of the mobile

phone, rural incomes have not doubled, \$100 PCs have not materialized, and increasingly, observers question the wisdom of such projects as technology solutions looking for a problem.

In this position paper, we summarize the history of both the fields of artificial intelligence and information and communication technologies for development, discuss where they face their biggest challenges, and suggest how AI applied to global development can learn from the history of its parent fields. Our suggestions are that AI be applied to the development context with ambition, but also with humility and realism; from a practical point of view, we believe that AI’s greatest contributions can come when viewed as a way to serve and amplify existing, successful development efforts, and not as technology that solves any of the myriad challenges associated with poverty.

A Brief History of AI

The field of artificial intelligence emerged in the late 1950s. Alan Turing claimed that computers would be considered intelligent by the end of the 20th century, and he proposed the Turing test of computational intelligence, which remains an open challenge to this day. Despite some early progress, such as Newell and Simon’s work on automated problem solving and McCarthy’s work in logic programming and question-answering, it became increasingly clear that capturing human intelligence in computer form was not quite as simple as these early luminaries had imagined: many problems were simply beyond the frameworks developed by early AI researchers, and equipping computer systems with large common-sense knowledge bases was met with both technical and philosophical hurdles.

Despite his immense optimism in the 1970s, by 1982, Minsky had made a dramatic about-face: “[T]he AI problem is one of the hardest science has ever undertaken” (Kolata 1982). At this point, the field had experienced notable, public failures (e.g., the General Problem Solver, or GPS (Newell, Shaw & Simon 1959)), funding for AI research endeavors had shrunk, and research progress was at a standstill. Then, over the next couple of decades, artificial intelligence as a field gave birth to many subfields, such as vision, learning, speech, and robotics. While each sub-community has achieved individual

successes, together they still appear far from achieving human intelligence (Cohn 2006).

A Brief History of ICT4D

Richard Heeks, an observer of ICT4D well before the abbreviation was used, claims that the first use of information technology in India can be traced back to 1956 (Heeks 2008), which is also the year in which John McCarthy claims to have coined the term “artificial intelligence” (Skillings 2006). For a long time, computers in the developing world were primarily restricted to governmental organizations and academia. It was in the 1990s, with the explosion of the Internet and commoditization of digital technology, that the global development community began to see ICTs as a way to achieve development ends. The two World Summits on the Information Society (WSIS) were expressions of this new interest.

At around the same time, multinational technology companies noticed emerging markets as a new frontier, and started to explore technology’s relevance in the developing world. For example, Intel initiated several product-development efforts in China and India, and Microsoft set up its “Unlimited Potential Group” to address emerging markets.

Riding on this wave of interest in technology and development, numerous projects gained rapid momentum and spouted rhetoric that was not founded on good evidence. One example of this early hype was the rural telecenter. Its proponents insisted that providing rural areas access to information via computer kiosks would close the “digital divide” and in turn alleviate poverty. But, research over the last decade has stubbornly refused to find reliable impact on the communities that telecenters serve (Sey & Fellows 2008). Another prominent example is the One Laptop per Child project, which seeks to put an inexpensive laptop in the hands of every child as a way to solve deficiencies of existing educational systems. However, here, too, a technology alone appears unable to overcome the exact problems that face those same educational systems: poor infrastructure, absentee teachers, poor teacher training, and low spending on education over all (Kraemer, Dedrick & Sharma 2009; OLPC News 2010).

AI Challenges

So, what is it that makes AI – the ambitious, human-level-intelligence brand of AI – so difficult? Many of the challenges are deep – but from the perspective of the researcher, even the potential route to progress is strewn with roadblocks. Difficulties range from the need for great amounts of data to train learning systems; to usability challenges that allow imperfect AI systems to complement users; to resource challenges inherent in building large, deployable systems.

Data Challenges

Large Training Set Requirements. As our understanding of machine learning becomes increasingly sophisticated, some researchers have come to believe that data is king: Good data corpora are frequently more valuable than even significant improvements in algorithms (e.g., (Rajaraman 2008)). Unfortunately, vast amounts of data are often required to train a useful, intelligent system. The actual amount of data required for a given purpose is also highly variable and domain-dependent.

Labeled Data Requirements. The type of data required to train an AI-based system is also important. Not only must there be large quantities of data, but in many cases (e.g., for any application of supervised learning in particular), the data must be labeled in order to be meaningful. Labels come at a price, however. In some cases, labels must be manually generated by people who understand the context, which can be especially cumbersome. Fortunately, in other settings, data sets come with labeling: diagnostic medical data sets, for example, often contain a list of symptoms and the doctor’s diagnosis, all of which is collected in the process of providing patient care, rather than in a post-hoc labeling process.

Usability Challenges

When expectations are not met, users of AI-based systems are likely to be disillusioned and discontinue use of a system (Whitworth 2005). Furthermore, while AI algorithms often operate in the background, unnoticed by users, for a system as a whole to be usable, it must also have a suitable front end. A recent AI Magazine special issue on *Usable AI* describes usability challenges for AI systems in much more detail (Jameson, Spaulding & Yorke-Smith 2009).

Transparency & Control. Beyond data requirements and ease in interfacing with an intelligent system is a requirement for transparency into a system’s decision process and control over its decisions. People want their tools to be deterministic, providing them with a clear understanding of why a system does what it does (Schneiderman 1998). And the ability to scrutinize the reasoning behind a system’s decisions is only useful when the mechanisms behind that reasoning and those decisions can be modified to a user’s specifications (Höök 2000).

User Satisfaction. Perceived usability is strongly affected by user satisfaction, which may itself be affected by various features of an intelligent system (including the existence of control and transparency). For example, users tend to appreciate an intuitive system that serves a need, avoids making blatant or irreversible errors, fails gracefully when it does make errors, and deals effectively with various other situations frequent in interactions with a user in the loop.

Resource Challenges

Finally, the resources required in the development of working AI systems tend to be extensive. For instance,

many recent, high-visibility AI-based initiatives have required large teams working for years (e.g., CALO (Mark & Perrault 2005), the 2005 DARPA Grand Challenge (Thrun et al. 2006)) – and these are simply for prototype projects that end up in closets or museums. The teams themselves are diverse, often involving collaboration among interface designers, computer scientists of various flavors, and domain experts (e.g., automotive engineers). It is not merely an intelligent algorithm that must be invented but rather a complete end-to-end system, including the requisite hardware and/or interfaces, either with respect to a user or to hardware or other software technologies.

AI Successes

Of course, AI has had many successes – underdelivering on great expectations does not preclude significant progress. Intelligent collaborative filtering techniques, for instance, have brought about more satisfying user experiences in the web search and online shopping domains (Google 2009, Netflix 2009). These inventions have overcome all of the challenges listed above, primarily due to availability of resources and strong command of their respective markets. Decades of steady research progress have also brought handwriting recognition to the commercial market. Tablet PCs now generate text from scribbled, handwritten notes. And, increasingly, customer service by phone is being delivered by speech recognition systems that use clever dialog design to avoid situations that speech still cannot discriminate. These successes have been fueled by large research and development teams, clear usability requirements, and widespread access to training data via field testing.

ICT4D Challenges

Unfortunately, the successes of AI in the developed world do not necessarily translate immediately to usefulness in global development. On top of the standard challenges of ICT4D, in general – deficient infrastructure, under-trained human capital, socio-political barriers, lack of capacity to maintain technology, etc., the challenges of AI applications are only heightened in developing-country scenarios. On the “ground,” in poor communities themselves, there is little use for assistive online search tools or self-driving cars when neither PC nor automobile are affordable. And, automatic medical diagnosis is of little value if quality medicines are not available or local clinics are not able to perform surgery. Yet, on the side of national or international policy, the best tools for decision making are just as liable to meet misuse as good use, if put in the hands of corrupt ministers and bureaucrats.

The points we make below intentionally parallel the discussion about AI. Although we believe the challenges for technology are generally greater in global development contexts, we also discuss a few advantages of developing-country contexts for AI.

Data Challenges

If data is hard to gather in the developed world, it is even more difficult in the developing world. Surveys are noisy and unreliable due to poorer human capacity in both delivering and responding to questionnaires. Languages are many and dialects splintered: the one-time cost to collect a sufficient parallel corpus for machine translation between English and Spanish could be amortized over the hundreds of millions of people who might benefit, whereas a similar corpus from English to Great Andamanese – a language said to be spoken by 24 people, none of whom are particularly wealthy – would never be worth it.

Labeling, too, is a challenge. In many Internet tasks, labeling happens through a crowd-sourced mechanism that takes advantage of a large population of well-educated users who don't have to work two jobs just to feed their children – Page Rank, for example, could be said to be taking advantage of data labeling that is done incidentally when websites are linked to one another. But, among poor communities, there is rarely an online crowd to source – they aren't spending leisure time on the web. Search engines are notoriously bad at finding relevant links to obscure material in foreign languages, except in those cases where keywords match.

On the other hand, many poor communities have human labor in great abundance, so it is becoming possible for data to be obtained at a much lower cost, often using the mobile phone, which has been embraced by even very poor communities at a stunning rate (Eagle 2009).

Usability Challenges

Usability challenges also suffer extra hardship in the developing world. Illiteracy is high, and due to unfamiliarity with PCs, potential users are often intimidated or wary of interacting with equipment they fear they might break (Medhi 2007). Cognitive models of how the machine works are similarly underdeveloped, if only because of lack of exposure. Thus, any additional features are even harder for users to learn and adopt. Finally, what might be considered intuitive to a seasoned developed-world PC user might come as a conceptual novelty. What meanings do “desktop” and “folders” have to a farmer whose life revolves around soil, compost, and the weather?

Resource Challenges

Finally, we come to the most obvious challenge in poor communities, namely, deficiency in resources. If working AI systems require expensive manpower and equipment in the developed world, there's no reason to suspect that they can be more easily afforded in the developing world. A data-analytics system that mines data from government corpora to help a rich country determine healthcare policy might be completely out of reach for the government of a poor nation. Similarly, if an AI application requires customization for each site or user, or technical or instructional support for unexpected events or behaviors, this adds the need for a qualified technician or programmer

– again, less available in worlds without a significant IT industry.

AI for Development: Recommendations

Moving forward, there is much that can be learned from the history and challenges faced by the fields of AI and ICT4D. In particular, we present what we think are the key takeaways for researchers who are new to working on projects for the developing world. We also suggest some specific possibilities for AI that could prove beneficial in development. Our recommendations are based not only on the history of the two parent fields, but also on our experience as researchers in the areas of AI and ICT4D.

First, nothing beats actually spending time in the environment you hope to impact. All of the following recommendations are negligible in comparison to this one point. If a picture is worth a thousand words, then a visit is worth a million, and an extended stay is worth trillions.

Second: Be ambitious but realistic about what can actually be accomplished. This also includes being humble in anticipating AI's potential contribution to development, and patient in seeing value and resultant impact. Both AI and ICT4D can claim important contributions only as a direct result of the ambition and determination of the people who worked in these fields, so enthusiasm is required, despite the apparent appearance of “mountains beyond mountains.”

Finally, technology is best viewed as *support* or *amplification* of people and organizations who are well-intentioned and competent. In the case of AI, it's likely that as an additional prerequisite, the person or organization should also have demonstrated capacity to operate and maintain technology. According to David Waltz, a prominent AI researcher, “[F]or the most part, AI does not produce stand-alone systems, but instead adds knowledge and reasoning to existing applications, databases, and environments, to make them friendlier, smarter, and more sensitive to user behavior and changes in their environments” (Waltz 1997).

Thus, both for technical and contextual reasons, AI should build intelligence into existing systems and institutions rather than starting from scratch, or hoping to replace existing systems, however broken – an institution that cannot fix itself is unlikely to be able to support and use a complex technology properly.

Examples of projects that could work might be a system for providing effective data analytics to sound government institutions, or providing customized tools for a corps of competent rural healthcare workers. But, for example, smart tutoring software for a school with absentee teachers? Well, it's expecting a lot from a school that can't get its teachers to show up, to keep track of its technology, maintain it, and use it well, especially in an environment where infrastructure is unreliable, technicians are scarce, and the ongoing costs of technology are expensive.

Challenges to Researchers

We next highlight two challenges we anticipate, specifically for *researchers*, working on AI for development. These are speculations based on hard-won experience in ICT4D, but certainly not scientific assertions. In fact, it would be nice to be proven wrong.

New AI Research?

While the developing world could benefit from AI in various ways, it is not all clear that *new* AI research is required for AI to be successful. ICT4D itself has struggled to take hold in computer science departments, largely due to the difficulty of finding technical research challenges that would be considered good computer science (Toyama & Ali 2009). As a logical subset of ICT4D, AI for development is sure to face similar challenges.

AI on the Ground?

We contend that there aren't many intersections between on-the-ground projects in development, where technology could directly impact a poor community, and the use of artificial intelligence techniques. As above, any project requiring data for training is unlikely to find large-scale data at anything other than an aggregated level. And, projects that hope to ease the burden of human labor or intelligence through technology will find it difficult to undercut the low cost of labor. Finally, even in situations in which more intelligence could help, it's rarely the case that lack of intelligence is the bottleneck. Thus, a medical diagnostic system to help rural nurses diagnose patients seems sensible in practice, but in reality, the bottlenecks are, more often than not, access to human physicians (with whom patients can establish trust), access to genuine pharmaceuticals (and not quack pills), or persuading patients to visit clinics in the first place.

We, thus, again counsel towards working with good organizations who have established a beachhead on the battle against these recurring challenges.

Research Horizons

The presentation thus far may be discouraging, but there is room for AI to make an impact on development. Below, we list some of the opportunities.

Digitizing Paper Forms. Paper forms are produced in abundance in government offices and non-profit organizations across the developing world (Singh 2009). Despite attempts to close the “digital divide” for these organizations, it's likely that a transition to digital data will not occur any time soon. In fact, some researchers suggest that in working with technology novices, paper remains a preferred medium because of its very physicality – paper receipts are more trusted than a transient LCD display (Parikh 2009).

This presents us with a unique opportunity for AI: How about systems that bridge a paper-to-digital gap by a combination of good AI and good UI design?

For example, one possibility is to use electronic tablets, which allow a user to write with pen and paper while simultaneously capturing the input digitally. One such recent project for microfinance incorporated machine learning to perform online digit recognition, as well as pattern recognition techniques to recognize simple pen gestures (Chakraborty & Ratan 2010).

Another possibility is to re-design paper forms so that they simultaneously retain their pen-and-paper ease of use, and are machine readable after scanning. There has been some preliminary work toward this end, for numeric input only, that is described in (Singh 2009), but much more can yet be done.

This kind of research is very much in line with our recommendations. On the one hand, there is ambition both with respect to the size of the problem (some government offices have warehouses of paper forms that they must look up for each citizen transaction) and the technical challenge – general handwriting recognition is not yet perfect, yet it seems within reach to buttress it with careful design and an adequate UI. On the other hand, there is realism in respecting the value of paper. Second, this type of work seeks to aid existing organizations. In fact, the authors of all three papers cited above did their research in close collaboration with competent non-profit organizations.

Data Processing and Analysis for Policy. Artificial intelligence techniques, and specifically machine learning algorithms, can be effective tools for supporting *data processing* and *analysis* especially at the level of regional or national policy. The ability to perform automatic pattern recognition may prove especially useful when experts are either costly or unavailable in a specific country or region, even if the right intent is there.

This seems like one of the more promising areas for AI, but our earlier caveats still hold: A good AI system is no use in the hands of a corrupt or incompetent bureaucracy, and it may be that standard statistical tools such as multivariate regression more than suffice for the task at hand. The latter “problem,” of course, is only a problem for the AI researcher intent on contributing AI *research* for development.

Conclusion

To reiterate our recommendations, we suggest the following for new researchers in AI for development: (1) First and foremost, visit the sites or the organizations that you hope to impact, and, if possible, stay for an extended period to gain good intuition; (2) Be ambitious but realistic, humble and patient with AI-D projects; and (3) Partner with an organization that is already doing successful development using technology, and look for ways for AI to amplify their impact.

References

- Chakraborty, S. and Ratan, A. 2010. Managing Microfinance with Paper, Pen and a Digital Slate. *In submission*.
- Cohn, D. 2006. AI Reaches the Golden Years. *Wired News*: July 17, 2006.
- Eagle, N. 2009. txtEagle: Mobile Crowdsourcing. In *Proceedings of Human Computer Interaction International (HCII)*, San Diego, CA.
- Google ©2009. <<http://www.google.com>>.
- Heeks, R. 2008. ICT4D 2.0 – The Next Phase of Applying ICT for International Development. *IEEE Computer* 41(6).
- Höök, K. 2000. Steps to Take Before Intelligent User Interfaces Become Real. *Interacting with Computers*. 12(4).
- ICT4D 2009: Proceedings of the 3rd International Conference on Information and Communication Technologies and Development. Doha, Qatar: April 2009.
- Jameson, A.; Spaulding, A.; and Yorke-Smith, N., eds. 2009. Special Issue on “Usable AI”. *AI Magazine* 30(4).
- Kolata, G. 1982. How Can Computers Get Common Sense? *Science*, 217.
- Kraemer, K.; Dedrick, J.; and Sharma, P. 2009. One Laptop Per Child: Vision vs. Reality. *Communications of the ACM*. 52(6).
- Mark, W. and Perrault, R. 2005. CALO: A Cognitive Assistant that Learns and Organizes. <www.ai.sri.com/project/CALO, 2005>.
- Medhi, I.; Sagar, A.; and Toyama, K. 2006. Text-Free User Interfaces for Illiterate and Semi-Literate Users. In *1st International Conference on Information & Communication Technologies for Development*.
- Netflix ©2009. <<http://www.netflix.com>>.
- Newell, Allen; Shaw, J. C.; and Simon, Herbert 1959. Report on a General Problem Solving Program. *Proceedings of the International Conference on Information Processing*, 256-264.
- OLPC News <<http://www.olpcnews.com>>.
- Parikh, T. 2009. Engineering Rural Development. *Comm. of the ACM* 52(1): 54-63.
- Rajaraman, A. 2008. Datawocky: On Teasing Patterns from Data, with Applications to Search, Social Media, and Advertising. *More Data Usually Beats Better Algorithms*. <<http://anand.typepad.com/datawocky/2008/03/more-data-usual.html>>.
- Sey, A. and Fellows, M. 2009. *Literature Review on the Impact of Public Access to Information and*

Communication Technologies. Seattle: University of Washington Center for Information & Society.

Shneiderman, B. 1998. *Designing the User Interface*. Addison Wesley.

Singh, G.; Findlater, L.; Toyama, K.; Helmer, S.; Gandhi, R.; and Balakrishnan, R. 2009. Numeric Paper Forms for NGOs. In *3rd International Conference on Information and Communication Technologies and Development*.

Skillings, J. 2006. Getting Machines to Think Like Us. cnet news, July 23. <http://news.cnet.com/Getting->

[machines-to-think-like-us/2008-11394_3-6090207.html](http://news.cnet.com/Getting-machines-to-think-like-us/2008-11394_3-6090207.html).

Thrun, S., et al. 2006. Stanley: The Robot that Won the DARPA Grand Challenge. *Journal of Field Robotics*.

Toyama, K. and Ali, M. 2009. Computing for Global Development: is it Computer Science Research?. *ACM SIGCOMM Computer Commun. Review* 39(5): 40-43.

Waltz, D. L. 1997. Artificial Intelligence: Realizing the Ultimate Promises of Computing. *AI Magazine* 18(3).

Whitworth, B. 2005. Polite Computing. *Behaviour & Information Technology* 24:353–363.