ISSN : 2395-4132

THE EXPRESSION

An International Multidisciplinary e-Journal

Bimonthly Refereed & Indexed Open Access e-Journal

The Expression Journal

Impact Factor 6.4

Vol. 9 Issue 3 June 2023

Editor-in-Chief : Dr. Bijender Singh

Email : editor@expressionjournal.com

www.expressionjournal.com

The Expression: An International Multidisciplinary e-Journal
(A Peer Reviewed and Indexed Journal with Impact Factor 6.4)
www.expressionjournal.com ISSN: 2395-4132

OPPORTUNITIES AND CHALLENGES: ARTIFICIAL INTELLIGENCE IN INDIAN BUSINESS WORLD

DR. TANUJ KHATRI
PhD, MBA, MA (Human Rights), MA (Journalism & Mass Comm.)
Assistant Professor (Guest Faculty)
Department of B.B.A.
S.S. Memorial College, Ranchi University, Ranchi

::

Abstract

The future of one-sixth of the world's population rests on the shoulders of India, and the inevitable AI revolution will play a vital role in shaping its development and growth. As AI becomes ingrained in society and daily life, it possesses the potential to accelerate progress while tackling traditional hurdles such as inadequate infrastructure and bureaucracy. However, investing in AI also entails long-term societal risks that must be carefully evaluated at this early stage. In this essay, we will explore the prospects and challenges of AI in India and identify opportunities for general advancements and industry-specific breakthroughs, with a focus on healthcare. Additionally, we will shed light on problems arising from existing social norms, such as caste and gender equations. As India ventures into the AI era, we will outline crucial actions and protections necessary for fostering robust and inclusive development. AI can act as a catalyst for India's progress by streamlining processes, optimizing resource allocation, and enhancing decision-making across various sectors. General opportunities lie in overcoming linguistic divides through natural language processing and utilizing vast troves of public data for insightful policymaking. Furthermore, specialized applications in the healthcare sector can revolutionize diagnostics, treatment, and patient care, ultimately improving the overall well-being of the populace. Despite these prospects, India must confront challenges rooted in its socio-cultural fabric. Therefore, it is essential to develop AI systems that are conscious of such biases and actively work to minimize them. To encourage inclusive development, AI projects should actively involve diverse stakeholders, including women, marginalized communities, and rural populations. Collaborative efforts can help identify pressing societal challenges that AI can address effectively.

Keywords

Artificial Intelligence, India, Opportunities, Inclusive Development, Problems, Perspectives, Challenges, Policies.

::

The Expression: An International Multidisciplinary e-Journal
(A Peer Reviewed and Indexed Journal with Impact Factor 6.4)
www.expressionjournal.com ISSN: 2395-4132

OPPORTUNITIES AND CHALLENGES: ARTIFICIAL INTELLIGENCE IN INDIAN BUSINESS WORLD

DR. TANUJ KHATRI
PhD, MBA, MA (Human Rights), MA (Journalism & Mass Comm.)
Assistant Professor (Guest Faculty)
Department of B.B.A.
S.S. Memorial College, Ranchi University, Ranchi

INTRODUCTION

Artificial Intelligence (AI) is concerned with understanding the nature of human intelligence and designing intelligent artefacts which can perform the tasks which, when performed by humans, are said to require intelligence. Any major advancement in technology brings with it a range of opportunities and challenges. While AI is likely to bring substantial economic growth, it is being predicted that a number of jobs would be lost due to the automation. Therefore, it is necessary to put required policy and infrastructure in place.

Although there has been a lot of study in the subject of artificial intelligence since the word was first used in 1956, it has only lately resulted in the widespread deployment of intelligent applications for various domains and jobs. The work done in the late 1950s and early 1960s was in the direction of developing general procedures that could be used across a variety of fields. The first winter of the field began in the late 1960s and lasted until the late 1970s since the results were not particularly encouraging.

Since it was determined that domain knowledge is crucial, the research concentrated on representing and utilising knowledge. Knowledge-based systems were the name given to the developed systems. In the field of medicine, MYCIN, a KBS using rules for representing expert knowledge, has been used successfully. Many such systems have been created in various fields. These were referred to as expert systems since they were based on the knowledge of human specialists.

Numerous businesses seized the chance. By the late 1980s, it was discovered that the rules were frequently fragile and ineffective in real-world applications, necessitating the incorporation of numerous rules to address a wide range of scenarios. There was widespread dissatisfaction in the sector when a number of companies struggled to implement the technology. The second AI winter began in the late 1980s and lasted until the middle of the first decade of the present century, when high-performance computers and storage became widely accessible and reasonably priced. The development of certain algorithms, like deep learning, served as its fuel. Using the Internet, widespread distribution may be easily possible on mobile devices. It should be emphasised that, for obvious reasons, the timing of the two winters indicated above is not exact.

There are various reasons for the current high increase in AI applications. First off, thanks to developments in the electronics used in sophisticated digital circuits, computers' computing power and storage capacity have increased while falling in price over the past few decades. This has increased the affordability of computing. One of the key prerequisites for applying AI is the availability of powerful computers with lots of memory and storage, which are now reasonably priced. It has been further supported by cloud technologies. It has enabled the deployment of AI applications without a significant initial investment.

The development of AI methods such as deep neural networks, etc., is the second crucial factor. Although ideas like machine learning and neural networks have been around since the 1960s, the sophistication of algorithms has significantly increased throughout the past few decades of research.

The development of machine learning has had a tremendous impact on the use of AI in everyday life. A discussion about the Master Algorithm, which can learn any knowledge from the provided set of data, has also resulted from this (Domingos, 2015). Natural language processing, which is used in machine translation, user interfaces, etc., is another field of artificial intelligence where significant progress has been made. With the help of these technological advancements, a variety of applications were created to let individuals communicate naturally. These were known as chatbots, and they served the purpose of assisting users with chores like answering questions. These are referred to as virtual helpers.

The Internet of Things (IOT), low-cost sensors, high-speed communication networks, and mobile devices are some examples of the technical advancements that have led to the generation, storage, transmission, and processing of high volumes, high speeds, and high varieties of data. A large amount of data can offer a variety of insights that have not before been possible. For instance, Google uses enormous corpus of texts in numerous languages for two-language machine translation. The material is divided into phrases, which are then used to learn how to translate each sentence between two languages.

Businesses have revised their strategies in response to the industry's recognition of AI's worth. First AI policy has been adopted by Google. It is investigating every aspect of the application for its firm.

It declared in 2014 that it would start operating as an AI company by 2024. It should be mentioned that during the past few years, brand value rankings for IT companies have risen steadily. The most valuable brand in 2016 was Apple, but Google overtook it in 2017.

Main Thrust

AI-Driven Development Opportunities

AI has the potential to expedite development in India and help the country overcome long-standing obstacles including inadequate infrastructure and bureaucracy. Applications for AI can be found in almost every industry, including finance, healthcare, law enforcement, transportation, agriculture, and environmental protection. The Indian government recently established a task force with the specific purpose of identifying opportunities for AI across sectors and directing policy In this part, we discuss a few (at least, typical) Indian problems that are amenable to AI. We narrow our focus to three exemplary problems, which we discuss in some depth, as opposed to enumerating a lengthy list. The first two cross multiple industries, is specific to one industry.

NLP/ASR Scaling Up for Indian Languages

India is among the nations with the highest number of languages spoken, with over 700 different languages. Indo-Aryan, Dravidian, Austroasiatic, Sino-Tibetan, Tai-Kadai, and Great Andamanese are among the six families represented by the languages. Over a million speakers of at least 20 languages consider them their native language. Due to the prevalence of monolingualism and bilingualism, language naturally becomes a barrier to communication and information access. 87% of Indians cannot access this material because it is written in English.

Both Automatic Speech Recognition (ASR) and Natural Language Processing (NLP) have a lengthy history as AI research subjects. Systems for spoken dialogue, sentiment analysis, social media analysis, and machine translation have all made significant progress in these areas. The English-speaking audience of this work is probably accustomed to receiving pertinent search results, large training corpora are a major factor in the success of contemporary NLP systems like Google Translate. Sadly, when compared to the extent of data sets available for the most popular Western languages, most Indian languages have quite small data sets. For example, a public corpus of parallel text in 11 European languages contains tens of millions of words in each language, whereas Indian language researchers must make do with only around a hundredth of that amount of data. Thus, despite NLP (Natural Language Processing) being a popular study topic in the Indian AI community, the lack of digitized data limits the productivity of this field."Utilising resource-rich languages as "pivots" to create applications for resource-poor languages is one tactic that has gained popularity in the NLP community. This method can produce results quickly and may be of independent importance to linguistics.

Building systems to gather and transmit data in Indian languages is the only option we can see in the long run. We suggest that the AI community take up this effort as a serious endeavour.

For instance, ASR and innovative crowdsourcing techniques could offer avenues for digitising language data. Strangely, these regions complain about a lack of information and resources themselves. Under-resourced languages have been designated as a special issue in the field of ASR, similar to NLP. Languages that are suitable candidates for linguists to study are identified in a recent assessment by Pavlick et al. (2014) because they can be translated on Amazon Mechanical Turk.

In contrast to the use of practical methods, our plan views the linguistic differences in India as an immovable obstacle that must be overcome. The focus does not necessarily need to be on difficult activities at the beginning, like translation. More content in each language in the digital sphere, services like search and speech interfaces, and same-language subtitling in videos to boost functional literacy might all pave the way for healthy local language digital ecosystems.

Building and Exploiting Public Data

Every division of the government produces records that are freely accessible. The "Right to Information" act, which the Indian government passed in 2005, allows people to ask governmental institutions for specific kinds of information. This feature has already been successfully utilised by people and civil society as a positive step towards introducing transparency. But in order to make accountability and effectiveness fundamental to government-affiliated institutions, standardised data delivery pipelines must be constructed. In this context, Berners-Lee's classification of free data into five stars , which is presented in Table 1, is instructive. Any data, regardless of format (scan, picture, table), encoding, or availability on the Internet under an open licence, falls under the lowest category. Naturally, organised material that is linked to additional pertinent sources and presented in a table format, for example, is more immediately accessible than unstructured data that is only offered in natural language. The foundation upon which pertinent apps and services can be created is reliable, structured data. It is obvious that data system architecture going ahead must aim for a 5-star rating. It's interesting to note that looking backward also reveals an outstanding potential for AI.

There is a sizable quantity of data, notably from the last few decades, which does, even though the majority of legacy material does not even match the 1-star condition (being available in digital form, and accessible through the Internet). This data may include important information, but until it can be processed into an organised form, it will tragically stay undiscovered. The type and extent of The data makes it impractical for human annotators to perform the structuring exercise, but modern AI tools (such machine vision and NLP) can certainly handle this.

To demonstrate that even "macro" trends in different data sets are frequently unknown, and gleaning them can provide important inputs for course correction and policy formulation, we go into the specifics of a case study.

It is well known that court proceedings in India can drag on for years [34] at various appeals levels. We concentrate on income tax disputes since they have specific appellate authorities, such as the Assessing Officer, Commissioner of Income Tax Appeals (CIT(A)), and Income Tax Appellate Tribunal (ITAT) [18], but nevertheless experience lengthy court processes. The High Court (HC) and Supreme Court (SC) hear appeals from the ITAT. Table 2 displays, as of March 2015, the number of appeals and the dispute amounts reserved at various levels of litigation.

Contrary to what one might anticipate, the last column of the data demonstrates that instances with bigger dispute amounts were not appealed to higher levels. Keep in mind that although the average disagreement amount at ITAT is more than twice as much as at CIT(A), it decreases by a third when moving from ITAT to HC and even more so when moving from HC to SC. It is obvious that recognising this trend is essential to developing strategies to shorten wait times at the various appellate levels. One explanation could be that the government appeals more frequently at levels above ITAT in order to establish precedent, and another could be that many cases have been filed over the past ten years and are still pending at lower levels, which affects averages. It would seem to be a fairly simple task to confirm which of these hypotheses is true, yet curiously, the answers to the straightforward questions mentioned below are still unknown!

• At each level of appeal, what percentage of cases are brought by the taxpayer or the government, respectively?
• What percentage of appeals are successful for the government and taxpayers at each level?
• How much does the typical dispute in taxpayer appeals and government appeals cost?
• How long does it typically take for a case to be assessed and then resolved?

The ITAT and SC have independent websites that house material in various formats; up until a decade ago, their sites were not even searchable on the internet. An essential tool for legal research in India is Indian Kanoon (https://indiankanoon.org), which created specialist scrapers and provided free search services. But even when pertinent documents are located (such ITAT rulings), their lack of organisation still presents a challenge. Rupee numbers are frequently mentioned in typical ITAT rulings; the only way to determine the dispute amount is from a statement in natural language. The dispute amount in M/S Jain Furnishing v. ACIT, for instance, is the total of the two amounts mentioned in the sentence that begins, "The assessee in this appeal challenged the addition of Rs. 15,609/- on account of municipal taxes and addition of Rs. 4,80,000/- disallowing part of the rent."

Even if it wouldn't be simple, it seems doable to train an NLP approach to extract pertinent fields from tax decisions, such as the disagreement amount, especially if domain knowledge can also be used. Eventually, this straightforward technical intervention may assist in identifying bottlenecks in the tax appeal hierarchy and help conserve valuable time and resources. There are other potential in other areas of India's legacy data. For instance, the Trivedi Centre for Political Data (accessed November 11, 2017, https://tcpd.ashoka.edu.in/newabout- us/) examines a number of opportunities in the political sector.

Table 1: 5-star categorisation of open data, reproduced from the web page maintained by Berners-Lee (2010)

*	Available on the web (whatever format) but with an open licence, to be Open Data
* *	Available as machine-readable structured data (e.g. excel instead of image scan of a table)
***	as (2) plus non-proprietary format (e.g. CSV instead of excel
****	All the above plus, Use open standards from W3C (RDF and SPARQL) to identify things, so that people can point at your stuff
*****	All the above, plus: Link your data to other people's data to provide context

Table 2: Appeals at different levels of litigation [20] (L = lakh = 105; C = crore = 107).

Appellate Number of authority	Amount in Average perappeals	Amount in dispute (Rs)	Average per case (Rs)
CIT(A)	2.32 L	3.84 LC	1.6 C
ITAT	37,506	1.45 LC	3.9 C
HC	34,281	37,684 C	1.09 C
SC	5,661 4,	654 C	82

Healthcare

AI technology has the potential to significantly ease the difficulty of gaining access to high-quality healthcare in developing nations. Lack of qualified medical professionals ready to work outside major cities is one of the industry's biggest issues. The WHO's standard for a nation's healthcare workforce is 22.8 to 59.4 skilled health workers per 10,000 people; India falls between those two levels.

Modern AI makes it possible to find ways to improve the skills of scarce staff and, to some extent, make up for the lack of standard lab facilities. For instance, Gann et al. (2017) have shown that features that human pathologists are often not trained to identify can be used to predict the recurrence of prostate malignancies. Similar to this, Beck et al. (2011) use computational techniques to extract and use newer, more useful information for the prognosis of breast cancer. The creation of a software-controlled microscope that can accurately diagnose malaria in the field is yet another achievement of computational pathology. A significant factor in neonatal mortality is neonatal sepsis. According to studies, time-series data from common non-invasive metrics, such heart rate and breathing during a preterm baby's first few hours of life, can predict morbidity with an accuracy comparable to invasive (and frequently expensive and unavailable) lab tests.

Epidemiological study can be aided by data-driven algorithms to better identify disease load and treatment options. The well-designed POSEIDON study collected data from clinics in 880 Indian cities and towns over the course of one day. This data, which was collected from over 200,000 individuals, even with a preliminary study reveals patterns in the frequency of visits to healthcare facilities across gender and age groups, classifications of ailments, etc. Additionally, there are observable variances from comparable data sets acquired in other nations like Singapore and Sri Lanka. As a result, large-scale data analysis can give healthcare policy non-trivial inputs.

The use of automated capturing techniques, such as IoT-enabled medical devices and app-based forms with location- and image-based inputs, can help digitise health data. The goal would be to build pipelines that offer genuine and correct data with the least amount of human involvement. Theft of equipment and theft of prescription drugs. Such delivery losses might be reduced by computer vision-based solutions for tracking inventory and workers.

Healthcare has been used as an example "vertical" in this section, but it should be highlighted that both issues and practical solutions frequently cross boundaries. For instance, bad infrastructure and corruption may result in poor service delivery, limited access to education and information, a debt trap from excessive out-of-pocket costs, and poor health in a community. Understanding the connections between different problems would be a good overall strategy before jumping to solutions.

An AI-Centric Approach Has 3 Risks

There is optimism and hope since AI can advance development in a variety of (sometimes unexpected) ways. However, failing to foresee and mitigate the possible risks of growth driven by AI would be naive. The key issues that stem from India's socioeconomic setting are addressed in this section, including labour displacement. The global AI tsunami, which is starting to remove workers from their occupations, does not exclude India . According to a recent report by McKinsey & Company (2014), improvements in machine learning and natural language interfaces (voice recognition) could have an impact on 6–8 million individuals who are currently employed in typical administrative, customer support, and sales professions." A big number of individuals who may be dependent on these wage earners could be negatively impacted by a loss of jobs on this magnitude, which would be a serious consequence for a middle-income country trying to lift many people out of poverty. India's esteemed IT sector is already facing the effects of automation, warning that a crisis brought on by the loss of jobs may affect the populace in the coming years. Other AI adverse effects can take longer to become apparent.

Societal discrimination is strengthened. India's caste system is a societal structure with a long history. Unfortunately, it still tolerates discrimination in sneaky and covert ways that have an impact on salaries, employment, incarceration rates, and access to bank credit . The possibility that data-driven algorithms may inadvertently take up biases from the data they are fed has grown in importance with the development of AI. For instance, recidivism rate assessment systems in the US are thought to exhibit racial biases. Names and addresses contain caste and religious markers, which can easily influence data-driven algorithms that are used to evaluate applications for jobs, loans, or bail. In a study conducted a few years ago, Banerjee et al. (2009) discovered evidence of caste-based discrimination in employment applications for call centres.

Even if we assume that the decisions in this instance were made by human judges, it is still problematic if they are later used to train an algorithm that will screen applications. increasing the gender pay gap. Both the total number of Internet users and the total number of mobile Internet users in India are projected to increase in 2017, reaching 420 million and 300 million, respectively.

In India's rural areas, where 60% of Internet access occurs via mobile phones, mobile phones are the main access point to the Internet. While the widespread use of mobile devices appears to be a benefit for AI generally, it may unintentionally exacerbate the gender gap.

Due to patriarchal and misogynistic social standards, women in South Asia are 38% less likely than men to own a mobile phone . As a result, the actual access rate may be substantially lower. As a result, the application of AI may be divided along gender lines (along with other divisions resulting from economic and geographical constraints).

Another issue is the significant level of gender inequality in India's software sector. Therefore, there is a genuine danger that the AI that the entire population will use will be biassed heavily in favour of men. Unwanted long-term effects could result from this imbalance. Targeting the underprivileged and excluding them. Due to the significant costs of creating AI-based apps, it's possible that private firms will provide the initial impetus. With no special need to address socially important concerns like equal access, it is reasonable for corporations to seek profits from sectors with sizable profit pools. As a result, the requirements of the less lucrative may not be taken into account. It is useful to use the example from Figure 1: It is improbable that Google will give its Tamil–Hindi translation engine the same priority as its English–Mandarin engine. There is a chance that the poor would become even more marginalised when commercial interests are combined with AI-based platforms. A compelling explanation of this unsettling possibility can be found in a recent piece by Calo and Rosenblat (2017).

Steps and Safeguards for AI for Development

In this part, we offer a few guiding ideas for the development of an effective AI ecosystem in India. Without the simple thermometer, neither automotive engines nor air conditioners could have been created. Building the tools to measure India's "vital statistics" is essential at this point so that they can later be enhanced. AI requires accessibility to pertinent digital data in order to function effectively. We advise that the development of 5-star data pipelines be given priority, as was already stated in Section 2.2. The "Digital India" and "Open Government Data" programmes of the government (accessed October 6, 2018, at https://www.digitizeindia.gov.in/) are good moves in this direction. Additionally to publicly available data from governmental agencies, it would be beneficial to develop locally pertinent public open data sets about things like language, health, crops, markets, and so forth. In some circumstances, AI tools like machine vision and crowdsourcing may be used to kickstart the endeavour to produce such data sets.

If the activity of creating AI-based solutions is restricted to a few people and locations, it would neither be efficient nor sustainable. It is crucial to actively train a larger segment of the community to develop and maintain AI systems for their own needs, particularly women, linguistic minorities, and rural communities. Our illustration of Google's erroneous translations in Figure 1 still serves as an example. Clearly, local speakers who are aware of the gender biases present in their own languages can correct such errors more efficiently and promptly [10]. However, in order to create and operate their own translation engines, they will need both the data and the technological know-how. The suggestion made by Jain (2002) to actively supplement the Gandhian, bottom-up model of information-generation and distribution with the Nehruvian, top-down model is particularly pertinent to the development of the AI knowledge network. The development of AI frameworks, standards, and APIs can be facilitated by expanding the open source movement, which has had some success in India.

India benefits from having a well-established educational system and a skilled labour force. However, a large, diversified, and developing country's demand exceeds the supply of information and skill. Flagship presentations like Deepmind's AlphaGo programme [56], especially if they can be contextualised in the Indian setting, can inspire young people to work in the field of artificial intelligence. It would also be beneficial to publish intriguing data sets and hold competitions. Domestic research centres of excellence may play a leading role not just in transdisciplinary fields but also in fundamental AI technologies. If AI become the new electricity, society would require both electricians and electrical engineers.

It would be a good idea to take steps to train a sizable workforce to produce apps employing vision, voice, and other technologies, as this could lessen the impact of job losses by absorbing some of the shock.

Industry, particularly startups, will be crucial in finding and realising the advantages of AI across various industries. India has a thriving ecosystem for tech entrepreneurship, with access to talent, funding, and big markets. As of May 2017 , there were over 300 startups in India with an emphasis on AI, and since 2014, more than USD 100 million has been invested in them. However, this amount pales in comparison to the US and China, where investments total more than $4 and $3 billion, respectively. Startups will have to overcome the challenges of a lack of data sets and skills; stronger collaboration with universities may be helpful in the latter case.

Startups that must limit risk can concentrate on high-volume, low-margin industries. For instance, even a 5% decrease in electricity use, raw material waste, or rejection rate can make a significant difference in the manufacturing sector. India will need to develop regulatory frameworks, such as safety and quality standards, legal frameworks addressing data security, privacy, and liability, and ethics review committees, to keep up with the development of AI.

Conclusions

In this study, the contours of AI in India are firstly defined. Our main finding is that AI entails both opportunities and risks, with opportunities typically being obvious and alluring and risks maybe taking longer to exhibit their consequences. We believe that with careful planning and management, AI may not only have a net positive impact on India's development but also assist in overcoming current growth barriers.

Given the significant stakes, it is essential to conduct a thorough scholarly analysis of the development of AI in India. This publication serves largely as a springboard for further research in this area. The offered template and the various references are provided in the hopes that they can assist the curious researcher in beginning more extensive research.

It is true that we haven't covered all the bases (we've only briefly examined healthcare). We point the reader to other recent papers that devote entire sections to agriculture, transportation, education, urban planning, security, employment, entertainment, manufacturing, robotic automation, and environment for a more thorough list of verticals.

Works Cited

- Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2017. Machine Bias. Pro Publica (2017). Accessed October 10, 2017,
- https://www.propublica.org/article/machine-bias-risk-assessments-in-criminalsentencing.
- Abhijit Banerjee, Marianne Bertrand, Saugato Datta, and Sendhil Mullainathan. 2009. Labor market discrimination in Delhi: Evidence from a field experiment. J. Comparative Economics 37, 1 (2009), 14–27.
- Biswajeet Banerjee and J. B. Knight. 1985. Caste Discrimination in the Indian Urban Labour Market. Journal of Development Economics 17, 3 (1985), 277–307.
- Michele Banko and Eric Brill. 2001. Scaling to Very Very Large Corpora for Natural Language Disambiguation. In Proceedings of the 39th Annual Meeting on Association for Computational Linguistics. Association for Computational Linguistics, 26–33.
- Andrew H Beck, Ankur R Sangoi, Samuel Leung, Robert J Marinelli, Torsten O Nielsen, Marc J van de Vijver, Robert B West, Matt van de Rijn, and Daphne
- Koller. 2011. Systematic analysis of breast cancer morphology uncovers stromal features associated with survival. Science Translational Medicine 3, 108 (2011), 108ra113.
- Tim Berners-Lee. 2010. Linked Data. (2010). Accessed October 28, 2017, www.w3.org/DesignIssues/LinkedData.html.
- Laurent Besacier, Etienne Barnard, Alexey Karpov, and Tanja Schultz. 2014. Automatic Speech Recognition for Under-resourced Languages: A Survey. Speech Communication 56 (2014), 85–100.
- Nicola J. Bidwell. 2016. Moving the centre to design social media in rural Africa. AI & Society 31, 1 (February 2016), 51–77.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In Advances in Neural Information Processing Systems 29. Curran Associates, 4349–4357.
- Erik Brynjolfsson and Andrew McAfee. 2016. The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies. W. W. Norton & Company.
- Ryan Calo. 2016. Robots as Legal Metaphors. Harvard Journal of Law and Technology 30, 1 (2016), 209–237.
- Ryan Calo and Alex Rosenblat. 2017. The Taking Economy: Uber, Information, and Power. Columbia Law Review 117 (2017), 1623–1690.
- Subrata Chattopadhyay. 2015. Corruption in healthcare and medicine. Indian Journal of Medical Ethics 10, 3 (2015), 153–189.
- Mike Cooley. 1995. The myth of the moral neutrality of technology. AI & Society 9, 1 (March 1995), 10–17.
- Kate Crawford. 2016. Artificial Intelligence's White Guy Problem. The New York Times (2016). Accessed October 4, 2017,
- https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligenceswhite-guy-problem.html.
- Credit Suisse. 2014. Global Wealth Databook 2014. (2014). Accessed October 27, 2017, http://publications.credit-suisse.com/index.cfm/publikationenshop/ research-institute/global-wealth-databook-2014/.

- Pratik Datta, B. S. Surya Prakash, and Renuka Sane. 2017. Understanding Judicial Delay at the Income Tax Appellate Tribunal in India. National Institute for Public Finance and Policy (2017). Accessed October 28, 2017.