

Natural Language Processing and Global Development: A Future-Focused Primer

Nicholas Monaco and Samuel Woolley

Introduction

Natural language processing (NLP) is a field that examines human language as a formal structure, with the ultimate goal of helping machines understand, process and produce language with an efficiency on par with humans. NLP¹ can be viewed as both a subfield of artificial intelligence and an interdisciplinary field drawing on computer science, linguistics, statistics and probability. The past five years alone have seen a profusion of useful NLP applications available to the general public: from high-performance machine translation, such as Google Translate, to AI voice assistants, such as Amazon's Alexa, Apple's Siri, and Microsoft's Cortana, these tools enjoy widespread popularity and usage around the world in dozens of different languages. Other, less evident applications, such as information retrieval (IR) techniques that underlie common search engines, predictive texting and grammar checkers have been around for decades, and are constantly being improved. Simply put, any program that enables computers and machines to analyze human language can be considered an application of NLP.

Automatic Speech Recognition, the Internet of Things, and the Semantic Web

One of the most promising applications of NLP, automatic speech recognition (ASR), aims to help machines understand and process spoken language. ASR allows humans to use speech as an input system to communicate with computers, obviating the need for written input through keyboards. ASR holds particular promise as a means of facilitating the Internet of Things (IoT). From mundane tasks like room temperature control to more significant ones such as news consumption, ASR holds the potential to obtain a desired outcome simply by speaking it into existence. Viewed from this angle, ASR is perhaps the most promising means of fulfilling Tim Berners-Lee's vision of the "Semantic Web" – a network of devices that process data and information efficiently enough to leave humans time to carry out capital-producing and capital-enhancing tasks. This also is the lens through which NLP can be understood as a field with vast potential as a tool for enhancing development around the world.

NLP Applications as a Tool for Development: Piggy-backing Leapfrogging

Much has been written about the benefits of technological "leapfrogging"², particularly in the context of development in Africa. In several countries, the usage of mobile technologies before the installation of more traditional telecommunications infrastructure, such as telephones and television, has brought transformational resources, such as viable banking and payment systems, to rural and developing regions. We believe that NLP and the IoT hold great promise in these same areas. The usage of ASR systems and machine translation/interpretation in particular hold great promise in linguistically dense and low-literacy communities. The ability to communicate with neighbors who do not share a common tongue, or to engage in political and economic processes through audio retrieval of news, are two illustrative examples of the way in which NLP holds promise as a tool for development.

¹ The field is also known by several other names: these include *computational linguistics*, *computer speech and language processing* and *(human) language technology*. The canonical reference textbook for natural language processing is Daniel Jurafsky and James H. Martin's *Speech and Language Processing*. A full, up-to-date version of this text is available at <https://web.stanford.edu/~jurafsky/slp3/ed3book.pdf>. A complete former edition is available at http://santini.se/teaching/ml/2014/JurafskyMartinSpeechAndLanguageProcessing2ed_draft%202007.pdf.

² <https://www.economist.com/special-report/2017/11/10/what-technology-can-do-for-africa>

Languages of the world and their institutional support

While there is much room for application of NLP techniques in such areas, the necessary prerequisite base of linguistic infrastructure and documentation is often not present in such communities. Though estimates vary by source, Ethnologue, an annual catalogue profiling the languages of the world, recently calculated there are an approximate 7,111 languages in the world.³ Initially, it can be tempting to believe that all of these languages exist in the same state of development – that they all have a standardized written form and documented, teachable grammar, for instance. But this is far from true. Ethnologue estimates that only 56% of the world’s languages have writing systems⁴. Standardized, written language is of course an important precursor to several foundational aspects of development: stable laws, efficient record keeping, literacy, and journalism, to name a few. It is also a precondition for efficient NLP applications, which hold particular promise in the context of developing countries.

Importantly, because languages do not all benefit from the same level of institutional or governmental support, at present, not all languages can be processed with equal efficiency and performance. NLP first necessitates a thorough linguistic analysis of the language in question, involving, among other things, a standardized written form, and formal analysis of its phonology, morphology and syntax, at a minimum. After these stages have been completed, written texts can be used as data to teach, or “train⁵”, machines to understand and efficiently process these languages. For example, Google Translate and IBM achieved significant performance gains in machine translation in the past 30 years through the usage of large amounts of parallel texts of pairs of languages. These documents, often referred to as “bitexts”, contain the same text in two different languages, and are useful for refining machine translation applications.⁶

Languages and the Matthew Effect

Sociologist Robert K. Merton once observed that professors and scientists tend to accrue more notoriety and opportunities once they have achieved initial success in their field. Merton deemed this snowballing of opportunities and fame “*the Matthew effect of accumulated advantage*”⁷. Roughly put, scientists who have made a name for themselves once tend to continue to receive a disproportionate amount of acclaim and opportunities thereafter – this is embodied in a paraphrased excerpt from the Gospel of Matthew, “*The rich get richer, and the poor get poorer*”. English and other high-profile languages have also benefited from the Matthew effect. These tongues have historically garnered more attention, documentation, analysis and research than other languages on Earth. Languages which have not enjoyed this level of academic research are referred to as “*low-resource languages*”.

As a result of this prioritization, the lion’s share of the worlds’ languages do not benefit from basic NLP support. For instance, Python’s Natural Language Toolkit (NLTK), a popular framework for implementing NLP tasks in code, only supports 15 languages for “stemming”, a process often used in information retrieval tasks. The number drops to 11 supported languages for stopword support, the preliminary step

³See more at *How many languages are there in the world?*, available at <https://www.ethnologue.com/guides/how-many-languages>.

⁴ <https://www.ethnologue.com/enterprise-faq/how-many-languages-world-are-unwritten-0>

⁵ Train being the word primarily used within the discipline of NLP.

⁶ See for example <https://www.foreignaffairs.com/articles/2013-04-03/rise-big-data> and <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html>

⁷ For more, see *The Matthew Effect in Science* at <http://www.garfield.library.upenn.edu/merton/matthew1.pdf>.

for any kind of statistical processing.⁸ Languages with a high number of speakers, many of which are Indo-European, benefit from the most impressive and accurate NLP applications, since the body of research underlying them is vast.

For these reasons, we believe that investment in practical linguistic research with the explicit goal of developing NLP applications for low-resource languages in developing regions would be a laudable aim for USAID and other development organizations hoping to use technology to hasten and enhance development in target communities. If progress is made in this regard, NLP could be a particularly useful tool for evaluating field interventions in some of the world's most underserved, and crisis-prone, areas.

⁸ For references on NLTK's supported stemming languages, see <http://www.nltk.org/api/nltk.stem.html>. For stopwords, see <https://www.nltk.org/book/ch02.html>.

Analyzing “Social” Data to Inform Crisis & Humanitarian Response in a Networked Age

Kate Starbird, University of Washington

Widespread adoption of information and communication technologies (ICTs) and especially social media has rapidly changed how people communicate and organize—including within the context of disaster and humanitarian response. These activities, when mediated by ICTs, create “digital traces” that can be used to understand human responses to disaster events and to guide formal response efforts to better meet human needs. Network analysis methods are one suite of techniques that can be used to make sense of these traces, both for real-time decision-making during crisis events and for after-action reports and research. In this short paper, we: (1) provide an overview of network analysis techniques for understanding “social” data—defining a few terms and providing a few general examples; (2) advocate for a mixed-method approach to understanding social data in the context of crisis and humanitarian response, and describe how network analyses fit into this approach; (3) provide a few examples of specific applications of network analysis to relevant cases in this context; and (4) highlights some ethical challenges related to these kinds of analyses.

Part I: Overview of Network Analysis Techniques for “Social” Data

At a conceptual level, network analysis seeks to understand the structure of relationships between entities. When we think of network analysis, we often focus on “social” network analysis—i.e. mapping relationships between people. Social networks can be built from a variety of different sources, and can map different kinds of connections. For example, researchers can use survey data to create a graph of family ties or social support systems. Platform designers can create network representations based on the “social graph” that map explicit within-platform ties (e.g. “friends” or “followers”) between accounts. And just about anyone with basic programming skills can collect data from Twitter and use it to map retweet patterns—creating a social network graph that captures the shape of information flow across accounts. However, network analysis is not restricted to social ties. The technique can be used to map out different kinds of relationships between entities—for instance, to reveal connections between terms that co-occur in the same document or map content sharing across websites. The structure of a graph typically depends upon the questions that a researcher or practitioner wants to ask or answer, constrained by the types of data that are available.

Structurally, a network consists of two types of components. *Nodes* represent entities (e.g. social actors, accounts, words, or websites), and *edges* represent connections between those entities (e.g. familial ties, following relationships, retweets, terms that appear to together, or a shared article across two sites). Edges can be weighted—e.g. to make distinctions between an edge that represents one retweet and an edge that represents 1000 retweets. They can also be non-directional, directional (e.g. a “following” tie on Twitter), or bi-directional (e.g. a “friend” relationship on Facebook). These elements constitute a basic network graph.

Network analysis techniques include a suite of quantitative metrics that can be calculated from the information in a network graph. There are algorithms that can be used to identify different cliques or clusters within the graph that are structurally distinct from each other. *Centrality* measures can be used to find nodes that are the most “influential” within a network or cluster. There are also techniques that function to characterize the connectedness of a network and to identify *bridges*—i.e. nodes that help connect distinct network components. Taken together, these quantitative metrics alone can provide insight into the structural properties of a network and help to identify influencers.

Though the “graph” itself merely implies the structure of the relationships and can be represented simply as a list of nodes and edges, there are many software applications (e.g. Gephi (<https://gephi.org/>) or NodeXL (<https://nodexl.com/>)) that allow end-users to visually represent the graphs in ways that allow for human exploration and interpretation. These applications provide algorithms (e.g. Mathieu et al., 2014) that create a spatial layout for the nodes and edges, using an underlying physics model of attraction and repulsion to place connected nodes closer together and disconnect nodes further apart. In most of these applications, nodes and edges can be sized or colored by various other properties, defined in the metadata of the graph, which allows users to visually explore how other data dimensions intersect with the network graph. For example, nodes representing websites can be sized by the number of tweets that link to them, and colored to show whether they represent media outlets, NGOs, or government websites.

Existing applications on personal computing infrastructure (your average laptop) are quite good for visualizing snapshots of small- to mid-sized graphs (~100K nodes, ~500K edges). However, additional computational power and more niche (typically in-house) software are often needed to look at larger graphs or dynamic graphs. The latter are particularly useful for seeing how a network takes shape over time, including assessing the influence of specific entities or activities on the development of a network (useful, for example, in examining how an influence campaign reshapes a network).

Part II: Advocating for a Mixed-Method Approach for Analyzing Data in the Crisis and Humanitarian Context

Though purely quantitative solutions are often attractive for reasons of replicability and generalizability, after years of research experience studying social media use during diverse disaster events, our lab has adopted and advocates for a mixed-method approach. For us, network graphs are not necessarily findings by themselves, but instead serve as jumping off points for additional analyses. We apply a “grounded” method that integrates and iterates between quantitative and qualitative analyses to build a multi-dimensional, interpretative understanding of the data (Maddock et al., 2015). We work across and move between macro- and micro-levels of analysis, to understand the relationships between large scale patterns and specific actors/actions. In this work, network graphs are often used to map out the underlying structure of a conversation, and then serve as guiding artifacts as we dive into specific parts of the graph to understand the entities and activities that create that structure.

This approach is deeply informed by the methodological innovation in the area of *crisis informatics* led by Palen and her colleagues at the University of Colorado (Palen & Anderson, 2016). It has also been adapted and applied to the study of misinformation (e.g. Maddock et al., 2015) and disinformation (e.g. Starbird et al., 2017). And its use of network graphs to guide qualitative analyses is aligned with Howard’s approach to conducting networked ethnography (Howard, 2002).

As we described above, our work relies on many different types of network graphs—including social networks (personal ties; friends/followers), networks of information flow (re-shares or retweets), networks of shared affinities (domains shared by the same users; Youtube videos commented on by the same users), linguistic networks (co-occurring terms or themes). Each illuminates different aspects of structure, and our specific research questions or hypotheses—often drawn from previous analyses on the same data sets—inform choices as to what kinds of network graphs to use as well as what kinds of properties to map onto these structures. Properties drawn from qualitative coding (e.g. account affiliation, website type, political stance) can be particularly insightful when overlaid onto a network graph. The process is inherently iterative. And this graph analysis can be combined with other types of analyses.

For example, the network graph below (Figure 1, left) is a retweet network graph generated from a collection of tweets that were posted in Spring 2018 that contained the term “NATO” in them. In this graph, nodes are Twitter accounts, sized by the number of retweets they received in the data. After creating the structural graph (in Gephi), we ran a “community analysis” using the Louvain method (De Meo et al., 2011) to detect structurally distinct clusters. Next, using simple NLP techniques to analyze tweet content and profile descriptions, we characterized each cluster of accounts (Red = U.S. conservative, alt-right and Pro-Trump; Blue = U.S. liberal and anti-Trump; Yellow = official accounts of NATO-affiliated government organizations; Green = anti-NATO cluster of “anti-imperialist” activists who consistently echo Russian disinformation; Pink = activists critical of Turkish-led violence in Afrin, Syria). Next, we applied those cluster categorizations to temporal graphs (volume of tweets over time) to get a sense of when each cluster was active (Figure 1, right). We also created tables to show the most active users and influential users in each cluster as well as the most-cited domains in each cluster. These overlapping analyses allowed us to distinguish between distinct parts of the NATO discourse, and find some interesting intersections between them—for example in the blurry seam between red and green clusters lie a very interesting set of accounts that repeatedly mobilize to amplify pro-Russian narratives.

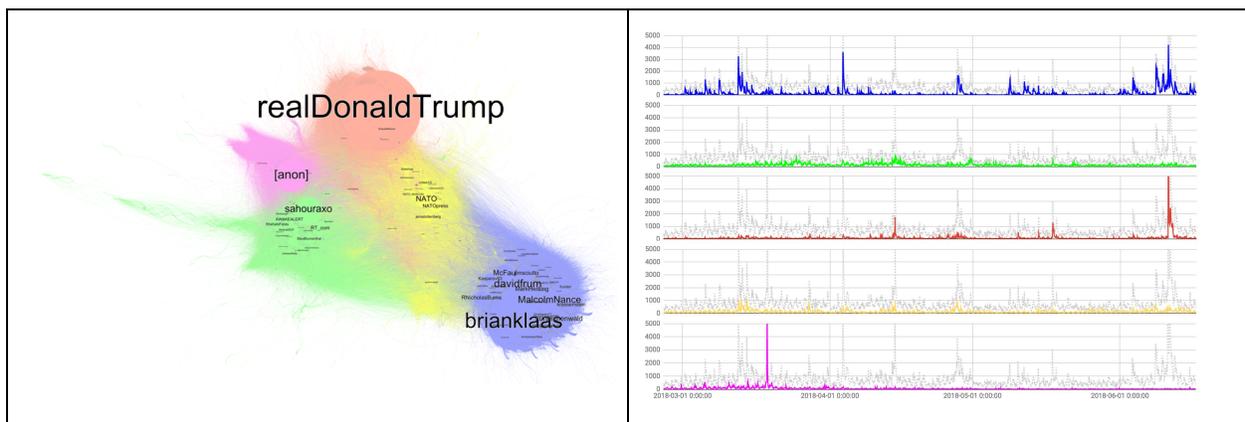


Figure 1. Retweet Network Graph (Left) and Temporal Graph (Right) for NATO Tweets.

Left = Retweet network graph generated through “NATO” tweets between February 2018 and June 2018; Nodes are accounts, sized by the number of retweets, colored by “communities” detected using the Louvain method. Right = Temporal graph separated and colored to show volume over time for accounts in each cluster.

These mixed-method analyses can be resource intensive, particularly in the human work required for interpreting the data—but this is critical for contextualizing quantitative findings, understanding anomalies, distinguishing between meaningful and unmeaningful patterns, etc. The approach we describe above has been created for the time scale of academic research, and likely maps to after-action reporting for humanitarian and crisis responders. However, we would recommend that even “real-time” network analysis tools be interpreted by trained professionals who are triangulating multiple signals from different levels (network, temporal, content-based, account or entity-based). Not every pattern or anomaly is meaningful or means what we think it means when we first encounter it in a graph.

Part III: Applying Network Analysis in the Context of Humanitarian and Crisis Response

This section highlights a few examples of network analysis being used to inform humanitarian response and/or to understand some of the challenges humanitarian organizations face in the current information

space. These are by no means exhaustive examples, but merely reflect a few different kinds of applications, ideally to highlight some of the potential of these methods in this context.

One promising application of network analysis in this context is mapping information flows about a topic or within a community (e.g. identifying hubs and bridges) to help develop communication strategies for best disseminating their messages. This can be useful for disaster response and public health campaigns. For example, researchers have proposed and prototyped a dashboard for identifying influencers in social media conversations about Ebola [Chung et al., 2015]—influencers who could be specifically targeted for public health information and counter-misinformation messaging.

Another application involves identifying and potentially countering mis- and disinformation. Rumors and misinformation (unintentional and otherwise) are persistent challenges for those responding to disaster events (Hiltz et al., 2014). Network graphs, combined with content analysis to identify rumors, can be used to identify key nodes in the dissemination of response-related rumors, and can inform strategies for countering those rumors.

Perhaps more problematically, humanitarian response efforts are increasingly targets for *disinformation* campaigns which utilize social media and “alternative” media outlets to undermine and delegitimize humanitarian response work—in some cases endangering the lives of response workers. These campaigns leverage criticism of the entanglements (both real and imagined) between humanitarian response organizations and the geopolitical goals of the sponsors of those organizations to mobilize online influencers and activists to spread their messages. These efforts include both seemingly “organic” elements as well as traces of coordination, including automated accounts, paid “troll” accounts, cultivated “journalists”, and the blending of state-controlled media. Network analysis can be used to identify and map out some of these relationships.

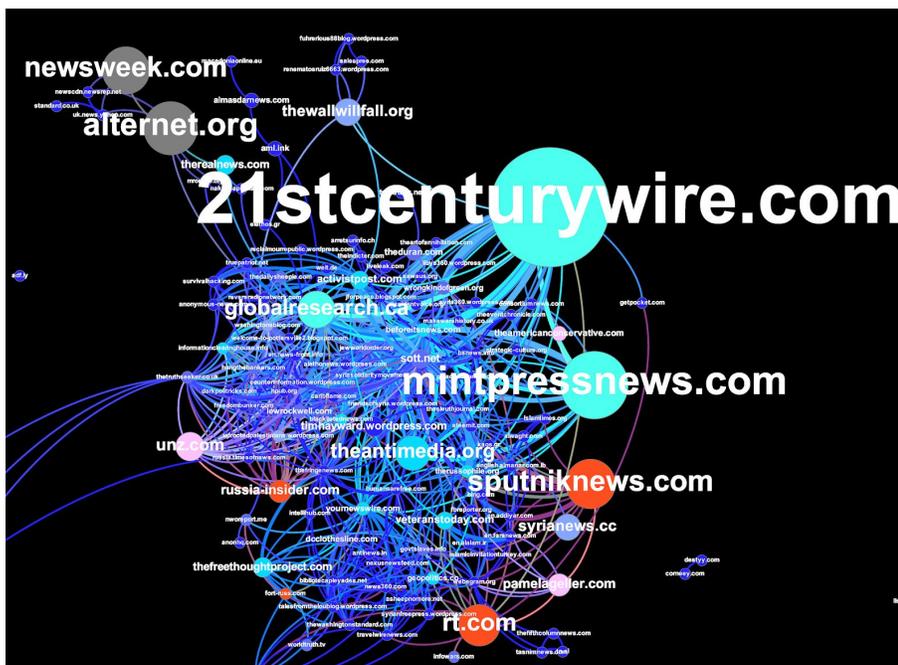


Figure 2. Content-sharing domain network graph for “White Helmets” Twitter discourse. We used links embedded in “White Helmets” tweets posted between June 2017 and September 2017 to identify articles and web domains in this conversation. Nodes here represent web domains and are sized by the

number of tweets that linked-to those domains. Nodes are connected by an edge when the two websites hosted the same article (or a highly similar article). Edge thickness is proportional to the number of articles that appear on both sites (1-7). This is a subset of the graph highlighting the “alternative” media ecosystem that consistently criticized the White Helmets.

For example, our lab has documented the ongoing efforts targeting the White Helmets, a humanitarian response organization that works in rebel-held areas of Syria (Starbird et al., 2018; Wilson et al., 2018). Our research seeks to identify and understanding how that disinformation effort is coordinated (and in some ways merely cultivated). We have created several network graphs, including a retweet network graph that helped us to see how the “White Helmets” Twitter conversation was dominated by a small number of anti-White Helmets voices, and a domain network graph that mapped content sharing across web domains cited in that discourse. Figure 2 is a view of the content-sharing, domain network graph. Each edge represents a case when the same anti-White Helmets article appeared on two different domains. Edges grow thicker when more articles appeared on both domains. This graph reveals how Russia government-funded and government-controlled websites, as well as numerous grey propaganda websites, are integrated into the media ecosystem (through content sharing) that repeatedly attacks the White Helmets. Graphs like this can help the targets of disinformation campaigns uncover and communicate the underlying connections between actors that participate in these campaigns—which can have operational and rhetorical value for defending their mission.

Part IV: Ethical Concerns

Unfortunately, the length and format constraints of this paper meant that some important topics remain under-explored. In particular, there are substantial ethical concerns that should be considered by researchers and practitioners as they integrate new data streams (including those generated by social media and other ICTs) into their work. These include geographic and demographic divides in platform access (and resulting data production) that can lead to under- or over-representation and unequal treatment for specific groups, as well as particularly concerning issues related to privacy and safety of the people who generate this data. Network data is very often about human relationships, and network analyses can reveal patterns in these relationships that were previously invisible to both those who generated the data and those studying the data. It is therefore important for those conducting and using these analyses—and those creating tools for others to conduct and use these analyses—to carefully consider how they create and interpret network graphs and how they communicate those graphs to decision-makers and the broader public.

Bibliography

Chung, Wingyan, Saikhe He, and Daniel Zeng. (2015) eMood: Modeling Emotion for Social Media Analytics on Ebola Disease Outbreak. Work in Progress Paper. Presented at the *36th International Conference on Information Systems*, Fort Worth, TX.

De Meo, Pasquale, Emilio Ferrara, Giacomo Fiumara, and Alessandro Provetti. (2011). Generalized Louvain Method for Community Detection in Large Networks. In *2011 11th International Conference on Intelligent Systems Design and Applications*, pp. 88-93. IEEE, 2011.

Hiltz, Starr Roxanne, Jane A. Kushma, and Linda Plotnick. (2014). Use of Social Media by US Public Sector Emergency Managers: Barriers and Wish Lists. In *ISCRAM*. 2014.

Howard, Philip N. (2002). Network ethnography and the hypermedia organization: New media, new organizations, new methods. *New Media & Society*, 4(4), 550-574.

Imran, Muhammad, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. (2018). Processing Social Media Messages in Mass Emergency: Survey summary. In *Companion of the The Web Conference 2018 on The Web Conference 2018*, pp. 507-511.

Maddock, Jim, Kate Starbird, Haneen J. Al-Hassani, Daniel E. Sandoval, Mania Orand, and Robert M. Mason. (2015). Characterizing online rumoring behavior using multi-dimensional signatures. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW 2015)*, pp. 228-241. ACM, 2015.

Mathieu, Jacomy, Tommaso Venturini, Sebastien Heymann, and Mathieu Bastian. (2014). ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software." *PloS one* 9, no. 6 (2014): e98679.

Palen, Leysia, and Kenneth M. Anderson. (2016) Crisis informatics—New data for extraordinary times. *Science* 353, no. 6296 (2016): 224-225.

Starbird, Kate, Ahmer Arif, Tom Wilson, Katherine Van Koevering, Katya Yefimova, and Daniel Scarnecchia. (2018). Ecosystem or Echo-System? Exploring Content Sharing across Alternative Media Domains. In *Twelfth International AAAI Conference on Web and Social Media, ICWSM 2018*.

Wilson, Tom, Kaitlyn Zhou, and Kate Starbird. (2018). Assembling Strategic Narratives: Information Operations as Collaborative Work within an Online Community. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018): Article #183.

Machine Vision

Raju Vatsavai, North Carolina State University

1. Technological perspective

Machine vision (MV)⁹ is the use of devices for optical non-contact sensing to automatically receive and interpret an image of a real scene in order to obtain information and/or control machines or processes¹⁰. From this definition, it is apparent that a MV system has two distinct functions: *sensing* and *vision*. MV starts with acquiring data (e.g., images) through remote sensors (e.g., cameras). Though many *sensing* modalities are in use by MV systems, optical imaging sensors are widely used in various applications. Sensing component of a MV system includes vision camera hardware, correcting lens distortions, coordinate systems and projections, camera calibration, and quantization and representation (image and videos). While sensing is a key component of MV systems, this technology is well matured and large number of off-the-shelf solutions are readily available. On the other hand, *vision* component of MV system deals with analysis and interpretation of data, typically scene understanding from images and videos. The following are major topics of this data analysis:

- **Image processing (IP):** Raw data coming out of cameras contains many distortions and image processing offers enhancements and corrections that result in better downstream processing and feature generation. Image processing involves color and lighting corrections, geometric corrections, denoising, contrast enhancements, etc.
- **Feature Extraction (FE):** is the process of extracting discriminative information via set of transformations to the images. Widely used features include local binary patterns (LBP), edges, edge density, textures, Fourier spectrum signatures, wavelets, image moments, histogram of oriented gradients, scale invariant feature transform (SIFT), shape-based features like Hough transform. These features shown to improve segmentation and classification performance.
- **Machine Learning (ML):** “The field of machine learning is concerned with the question of how to construct computer programs that automatically improve with experience.¹¹” ML is a subfield of artificial intelligence (AI), and is also closely related to the field of pattern recognition (PR) and data mining (DM), which are also concerned with learning from the data. The techniques from these three branches of the broad field of data analysis, such as supervised learning, unsupervised learning, semi-supervised learning, reinforcement learning and ensemble learning have been widely used in the MV systems.
- **Deep Learning (DL):** The field of deep learning, in particular convolutional neural networks (CNNs), which started with the work of LeCun et al for hand written digit recognition (1989), has become popular since winning ImageNet¹² challenge by AlexNet¹³. DL is the current state of the art in image classification, object detection, and scene understanding.

For the purpose of this brief note, we define MV system to include all aspects described above and includes both sensing technologies (both *insitu* and remote), and data analysis frameworks consisting of IP, FE, AI/ML, DM, PR, and DL. Sensing and Vision components of MV systems are analogues to humans' eye (sensing) and brain (vision). Though human visual system is still superior in many aspects, MV

⁹ We do not distinguish between machine vision and computer vision for the purpose of this brief. We refer to them as systems that use sensors (e.g., cameras) to acquire data (e.g., images) and analyze it to understand the phenomena under observation (e.g., real-world).

¹⁰ Machine Vision Association of the Society of Manufacturing Engineers

¹¹ Tom Mitchell. “Machine Learning.” McGraw-Hill Education.

¹² ImageNet: <http://www.image-net.org/>

¹³ AlexNet: <https://en.wikipedia.org/wiki/AlexNet>

systems are indispensable due to their advantages in terms of the environment in which they can operate. Key early adopters of machine vision included automotive, semiconductor, consumer products, medical/pharmaceutical, and food industry. For example, MV has been used by the food industry in the production (e.g., checking for overcooked vs. undercooked), packaging (e.g., to ensure correct labeling), and safety (e.g., checking for disease or contamination). MV systems have become indispensable to the industry as they can operate in hazardous, clean room, and high-speed production environments.

2. Application Perspective

Owing to the immense success of MV systems in the industry, and the advantages it offers, it has been widely adopted in many modern-day applications, some of which are listed below.

- **Remote Sensing (RS):** Remote sensing of Earth provides continuous observations at finer spatial, spectral, and temporal resolution. These earth observations (EO) have been widely used in to monitor and study crops, forests, urbanization, water quality, natural disasters and climate impacts on natural resources and man-made infrastructure, and identifying slums. MV techniques described above have been widely used in all aspects of the remote sensing field, from sensor design to image processing to feature selection to classification and change (damage) detection.
- **Robotics:** Robots are now widely used in the industry and as well as at homes. Robot use by agriculture industry is becoming common for tasks such as grading and sorting, packaging and loading. Such automation shown to increase productivity and reduce costs. Other important uses are medicine and rescue. Medical robots are now routinely used for patient monitoring to assisting doctors in complex surgeries.
- **Disasters:** Natural disasters, wildfires of California in the west to the Hurricane Florence in the Eastern United States, Earthquake in Papua New Guinea, heat waves in Pakistan, floods in India and Japan, Earthquake and Tsunami in Indonesia, have accounted for thousands of deaths and billions of financial losses. As per NOAA, in 2018, US alone has witnessed 14 separate billion-dollar weather and climate related disasters¹⁴ and it is estimated that global economic loss is over \$155 billion due to natural and man-made disasters across the globe¹⁵. MV systems play a critical role in acquiring to processing data and provide critical information to first responders to planners and policy makers.
- **Image (scene) understanding:** From understanding defects in manufactured parts to human face recognition to object (e.g., tumor or car) detection, to anomalous pattern (unattended bag at a public place to obstacle detection by driverless cars) detection, MV systems are playing a critical role in image understanding field. With advances in deep learning systems, MV systems performance is approaching or even surpassing human performance in many image recognition tasks.

3. Aid and International Development Perspective

Emerging MV systems play an important role for international development programs ranging from site selection, monitoring, intervention, compliance, and policy development. We now briefly identify few critical MV applications and research requirements.

- **Disasters:** This decade has witnessed increased intensity of natural disasters. Climate change is making these disasters worse. Recent UN report¹⁶ shows that climate change is behind the dramatic increase in economic losses. MV systems are critical in predicting, monitoring, and responding to

¹⁴ NOAA: <https://www.climate.gov/news-features/blogs/beyond-data/2018s-billion-dollar-disasters-context>

¹⁵ Swiss Re Group: https://www.swissre.com/media/news-releases/nr_20181218_sigma_estimates_for_2018.html

¹⁶ UN Report: <https://www.unisdr.org/archive/61121>

these natural disasters. Improved sensors and their deployment, such as seashore-based tide gauges, Deep-ocean Assessment and Reporting of Tsunamis (DART), wireless sensor networks for wildfire monitoring, dedicated remote sensing satellites for collecting global earth observations, and on demand UAVs, is critical for providing effective response and recovery to natural disasters. The devastation caused by the 2004 Indian Ocean tsunami is primarily attributed to the lack of early warning system and disaster management plan. Even in advanced countries like US, lack of early warning systems (e.g., flash floods) are leading to loss of human life and property damages. US Department of Homeland Security (DHS) ran a pilot program in Ellicott City, to deliver better flood warnings to residents via automated sensors¹⁷. Scaling such local efforts to global scale is urgently needed. Recently adopted “Sendai Framework¹⁸ for Disaster Risk Reduction 2015-2030” is an important step towards preventing new and reducing existing disaster risks, and MV systems plays critical role in achieving these goals.

- **Food, Energy, and Water (FEW):** “Water, energy and food are essential for human well-being, poverty reduction and sustainable development. With 10 billion projected population by 2050, it is expected the we need 60% more food, 10% increase in water usage by agriculture sector and more than 50% increase in global energy consumption¹⁹.” Though this is a multifaceted problem, we can think of novel solutions to meet the expected challenges by FEW sector. One important area where MV can play a critical role is food wastage and increasing crop yields. In US alone, it is estimated 30-40% of food is wasted during supply. According to UNEP, about 30% of all food produced worldwide gets wasted in food production and consumption systems. Food wastage is the single largest component of municipal landfill and major source of methane gas in US. MV systems can be deployed at each critical component of supply chain to identify and predict food wastage. Crop diseases can be monitored and predicted (see Figure 1), environment can be monitored and optimized during storage, and reduce food wastage in the kitchen. Likewise, MV systems can be used to monitor soils and waterways to reduce fertilizer runoffs. MV systems can be deployed for compliance testing. For example, US government spent about 8 billion USD during 2002-2017 on counternarcotics efforts in Afghanistan, however Afghanistan continue to produce opium. By using remote sensing images and machine learning, one can map crops grown (including opium) anywhere on the planet, not only to estimate yield but also for compliance.

Figure 1. Examples of food wastage during production, storage, and processing		
Corn leaf blight	Damages due to bugs	A rotten apple spoils the whole barrel (McGill University)

¹⁷ US DHS: <https://www.dhs.gov/science-and-technology/news/2018/05/21/news-release-dhs-st-partners-local-communities-improve-flood>

¹⁸ UNISDR: <https://www.unisdr.org/we/inform/publications/43291>

¹⁹ FAO: “The Water-Energy-Food Nexus A new approach in support of food security and sustainable agriculture.”



- **Real-time Analytics:** By integrating embeddable computing chips with sensor platforms (e.g., UAVs, Cars, Tractors) real-time knowledge can be generated, that may be critical to the end users like farmers or first responders. Recent advances in embeddable computing (e.g., Jetson TX-1/2, nano) is enabling real-time analytics (including demanding applications such as MV) as compared to traditional off-line analytics which takes days to weeks before useful knowledge can be extracted. Real-time analytics are critical for monitoring crops (e.g., identifying water stress, identifying if weeds are growing, or disease spread), transportation systems, tracking people, and during disasters. As an example, MV systems (sensors) can be deployed on tractors, and analyze data in real-time to identify potential sites for treatment and dynamically adjust fertilizer applicators.

This brief looked at machine vision systems, and associated technologies, and how they can support important applications that are of interest to the international development and aid programs. MV systems can be developed to identify regions for development by analyzing existing data, identify data gaps and deploy new sensors, track the progress, and check for compliance, develop intervention strategies and evaluate outcomes. With research and development investment in MV technologies, novel applications can be developed in support of reducing hunger to creating a sustainable society.

Creating a conducive environment for AI and ML in the developing world

Mehdi Oulmakki, Computer Science faculty at the African Leadership University

When exploring the prospects of deploying Machine learning and Artificial intelligence applications around the developing world, we must consider a complex ecosystem of academia, private sector, governments both local and foreign, as well as local and international non governmental agencies. Crucially, we must not get caught up on the act of “*deploying*” as it is but part of a process that involves the development, deployment, and ultimately maintenance of these projects over time.

From a technological standpoint, the skill set required is globally scarce, making it tremendously difficult to identify local talent. This adds to the challenge given the nature of the field itself: Context and culture matter in ML perhaps more than any other area of computer science, where data gathering decisions or omissions can translate months later into unusable results.

Beyond local talent, interactions with local industry, government, and target population are crucial from a data collection standpoint. ML’s key prerequisite is access to relevant training and modeling data, which can prove a technical, economical, and political challenge.

Finally, AI and ML come with security and ethics considerations that can be challenging to navigate, and are more often than not ignored. The perceived value of adopting the technology is not quite clear: The field still carries a connotation of foreignness, and is in need of local case studies and flagship projects to sway public opinion.

These three challenges: Local talent, access to quality data, and the optics of ML are common to all players in the field, and this note will explore how they manifest across the developing world. Before diving into these challenges, I must first make a few disclaimers:

Firstly, research on this specific issue remains burgeoning, and finding data is incredibly difficult—hopefully, this area of research can grow as our collective interest does. Secondly, this note will be heavily afro-centric given that I have spent the last few years with the African Leadership University, teaching a Pan-African cohort of computer science students.

Building up local technical expertise:

It should come as no surprise that education is a key area for capacity building. We need advocacy, initiatives, and investments around STEM education at both the university and K-12 level. The “We” used is a global one, as the tech skills gap is fundamentally a global problem: In 2013 the Obama administration predicted that by 2020, there would be a need for 1.4 million computer science related jobs and only 400,000 computer science graduates with the skills to apply for those jobs.²⁰ More recently, the European Union predicts 765,000 unfulfilled ICT jobs within the European economy by 2020²¹. While AI and ML are not the only drivers of this growth, they are significant contributors.

Why does this matter to our challenge of creating local expertise? Any initiative to upskill and train local talent must focus equally on talent retention. Local developmental projects will be competing with foreign tech institutions, big and small, scrambling for talent. We are at a historical height of labor mobility, and despite the rise of more isolationist political rhetoric, every western country is allowing room in their immigration policies to accommodate for needed talent.

²⁰ [2013 Whitehouse report](#)

²¹ From the proceedings of the [2017 E skills and jobs in the digital age conference](#)

Several institutions are tackling this brain drain issue in education. The African Leadership Academy for example takes a dual pronged approach: Financial aid comes in the form of a forgivable loan, contingent on spending a decade working on the African continent. This is complemented by intensive career development and placement efforts to make sure students find local opportunities. Since the first graduating class finished university in 2014, 65% of graduates who studied abroad returned immediately to the continent²².

We should also aim towards lessening the need for foreign education, and take a critical look at the university systems across the developing world. In the context of this note, local universities will play a crucial role in building up local capacity through two contributions: Skills acquisition and training, and Research into context-specific problems.

Research institutes and new programmes are launching at a steady pace, and the research community is growing deeper internal connections thanks to initiatives such as the deep learning indaba²³. This is crucial, as there are several unsolved problems to consider that mostly impact the developing world. In natural language processing for example, we still do not know how to handle rapid code-switching, i.e when a user switches between multiple languages within the same sentence, an extremely normal approach to language in most postcolonial environments. Targeted research grants and south-south collaboration in research can help fast track solutions to such issues, riding off the traction the research community is getting in the last few years.

As far as the role of universities in creating skilled talent is concerned, many challenges remain on the capacity and relevance front. The overall level of programmes is on the rise, building up on the research community above and partnerships with government and industry to launch programmes such as the masters in Machine Learning at AIMS²⁴, yet many programmes have not been updated for decades. Universities are also not braced to meet the sheer demographic growth that the globe is heading towards - particularly on the African continent. Out of 100 high school graduates on the continent, only 9 enroll to a tertiary institution²⁵, a data point typically attributed to a combination of capacity and affordability.

To face this, innovation is required in the very set up of educational institutions and their financing, and many organizations are moving from brick and mortar to the online space, or exploring alternative means of financing such as income sharing agreements²⁶. I believe we ought to look at the unaccredited²⁷ space as well, allowing more players to provide education to those who seek it.

²² From the African Leadership Academy [annual reports](#)

²³ 700 + researchers from across the continent are expected to gather in Nairobi—alongside satellite events—for a week of learning and teaching: <http://www.deeplearningindaba.com/>

²⁴ <https://aimsammi.org/>

²⁵ https://data.worldbank.org/indicator/SE.TER.ENRR?locations=ZG&name_desc=true

²⁶ The African Leadership University is running an income sharing agreement programme to fund its [undergraduate and unaccredited programmes](#), which so far is proving fairly successful.

²⁷ Coding bootcamps are on the rise globally. While the model has not proven successful yet for AI and ML, and can be a regulatory risk, it also provides a great place for retraining, and has shown good results for prerequisite skills such as programming and data analysis.

Ultimately to build up a talented ecosystem requires a long term vision, which balances teacher training, K-12 initiatives in STEM and english literacy²⁸, regularly updated university curricula, and space for unaccredited initiatives to provide alternative options to students. A vision that acknowledges the realities of brain drain, and builds up systematic approaches to offset its impacts.

Building up a culture of open and accessible data:

One of the most common challenges I have faced with students interested in applying ML techniques to challenges in their areas is the dearth of data. Data is the lifeblood of the field, and remains one of the scarcest, most jealously protected resource of the developing world. Over the last year, a dozen student projects I've advised have turned into data collection project as we found no relevant existing dataset.

In particular, I would like to share some of the findings of Milan Cvitovik who performed a survey of East African ML practitioners from industry and academia²⁹, and documented many gaps in available data:

- Audio recording of local languages, alongside transcriptions. This comes with the additional caveat of accounting for mixed language use.
- Written text in local languages and non-latin script.
- Pictures of non caucasian individuals for facial recognition.
- High resolution satellite imagery of local geography.
- General ground level imagery.
- Low resolution / smartphone quality pictures of documents.

Creating a strong AI and ML ecosystem starts with creating a data driven one. For ongoing projects or businesses, a data gathering and labeling strategy is crucial to best leverage ML solutions. Duolingo—a language acquisition application—is monetized through a subscription to its language learning platform, but has secured massive and varied datasets around second and third language acquisition for internal use and widespread sharing³⁰

More short term projects may be more keen on acquiring the data rather than gathering it, but this can be extremely costly. Companies who have secured data demand prohibitive prices, whereas governments may require tremendous convincing before releasing data publicly. Often, the data is not accessible in formats expected by practitioners³¹.

So what can be done? If you would benefit from an open culture of accessible data, start by sharing what you have. Contribute to establishing expectation around data availability and openness. As you partner with local businesses, governments, and ngos, consider data collection and sharing as part of the negotiation of your involvement. Support data collection initiatives, especially when they take a sustainable, entrepreneurial form.

Finally, influence on the legislative space around data protection laws is crucial - both to meet our ethical burden of privacy, but also to allow for more efficient use of resources. Kenya for example

²⁸ English is the de-facto language of computer science, while several initiatives exist to adapt material to local languages, benefiting from the global network of practitioners and resources requires english proficiency. [More context.](#)

²⁹ [Some Requests for Machine Learning Research from the East African Tech Scene](#)

M Cvitkovic - arXiv preprint arXiv:1810.11383, 2018

³⁰ <https://ai.duolingo.com/>

³¹ When finding digitized data, it is typically stored in traditional relational databases, which require non-trivial effort to turn into i.i.d real-valued vector pairs.

prevents some data from leaving the country, effectively preventing practitioners from processing information on foreign hosted cloud servers.

The optics of AI and ML, and creating a compelling narrative:

As a closing thought, I think it's important as members of this community to demystify the origin and uses of ML technologies, and have honest conversations with as broad an audience as possible on the technologies' strengths and flaws, achievements and shortcomings. The mainstream media's narrative around the technology circles around privacy issues, ethical mishaps, and overall fear mongering around the disappearance of jobs and the AI takeover.

The employment issue is key to garner political will and social support. AI and ML are perceived as tools for automation, and automation is seen as a need of western economies, whereas significant portions of the developing world are striving to secure employment opportunities for a growing young population. To put it bluntly, in many parts of the world it remains much more economically sound to hire, lodge, and feed a person rather than purchase a washing machine.

There is a dire need for case studies and success stories to balance the automation-heavy, job-threatening narrative, and provide a development centric, positive outlook on AI and ML to motivate students, businesses, and ultimately governments to change their stances towards the technology and see it as a tool they can own.

Ethical Considerations for AI Applications in International Development

Mark Latonero, PhD

Research Lead, Data & Society Research Institute | Fellow, Carr Center, Harvard Kennedy School

USAID's 2018 report, *Making AI Work for International Development*, includes an entire section titled "what can go wrong."³² The section describes the risks and harms arising from artificial intelligence (AI) and machine learning (ML) applications in development ranging from false results in predictive models to data collection as surveillance. This brief note brings an ethics frame to the discussion of what can go wrong when applying AI/ML systems in the international development context.

Ethics is a term with a multiplicity of meanings. Ethics can encompass ancient philosophical texts, normative theory, and contemporary domain applications. AI is also not a fixed term - it includes a cluster of existing technologies and methods such as machine learning, computer vision, and natural language processing; in addition to narratives of future imaginaries like artificial general intelligence. A review of the various literatures and controversies surrounding such terminology is out of scope. Instead, this discussion will revolve around the current trends around ethical research, development, and application of sociotechnical systems related to AI.

At present, AI ethics is among the hottest topics in tech,³³ precipitated by the steady stream of headlines about tech industry missteps and scandals: massive data breaches, online manipulation of elections, algorithms discriminating against people of color, and social media platforms contributing to genocide. Some have welcomed the resulting attention and the proliferation of ethics officers and ethics advisory boards at tech companies. Others have expressed concerns about ethics-washing or a co-option of ethics language with little or no meaningful process for implementation or accountability.

Academia is also abuzz with activity as a number of major U.S. and U.K. universities recently establishing multi-million dollar centers around AI ethics and values.³⁴ In the last few years, a dynamic new research area has emerged around the ACM's Fairness, Accountability, and Transparency conference, which is providing the theoretical and analytical groundwork for addressing pressing issues like bias and equity in AI/ML.³⁵ However, both industry and academia give far more attention to Silicon Valley and much less to ethical concerns of AI in developing countries.

Despite this lack of domain specific research, there are already AI/ML interventions in areas such as migration, food delivery, and counter trafficking in persons. However, the human costs of getting it wrong in these areas have higher stakes. With direct interventions in the lives of vulnerable populations, or those affected by humanitarian crises, the need for ethical standards and processes is greater. Development practitioners arguably have a higher duty of care to protect against negative outcomes. And the delicate ecosystems where development projects are implemented puts debates such as AI/ML bias in automated decision making in a unique—and more urgent—light.

³² USAID (2018). Reflecting the Past, Shaping the Future: Making AI Work for International Development. <https://www.usaid.gov/sites/default/files/documents/15396/AI-ML-in-Development.pdf>

³³ Jacob Metcalf, Emmanuel Moss, danah boyd (forthcoming). Owing Ethics: Corporate logics, Silicon Valley, and the institutionalization of ethics

³⁴ See the Stanford Institute for Human-Centered Artificial Intelligence, opened in 2019 with much fanfare and controversy.

³⁵ ACM Conference on Fairness, Accountability, and Transparency (ACM FAT*). 2018 <https://fatconference.org/>

The remainder of this note will consist of brief forays into three general topic areas where ethics, AI, and international development intersect: 1) context 2) power and 3) accountability.

Context

The challenges for AI/ML in development are not entirely new. The ICT4D community has been wrestling with overlapping ethical challenges for decades.³⁶ The folly of Western organizations expecting prosocial change by dropping a preordained technological fix into developing countries has been well documented.³⁷ Still, the idea that technology, AI or otherwise, can “solve” problems as complex as those in the development sphere remains problematic.

Tech solutionism has come to mean a faith that technology will solve the world’s most intractable problems,³⁸ or an assumption that the solution to a problem is technical rather than social, political, or normative. A related misconception for technologists is what Selbst, *et al.* (2018) call the portability trap: a “failure to understand how repurposing algorithmic solutions designed for one social context may be misleading, inaccurate, or otherwise do harm when applied to a different context.”³⁹ Causal models that focus first on the technology itself in search of a development problem to solve, ignore the human complexity that can only be discovered through a deep examination of specific social contexts. Mateescu and Elish (2019) argue that language of AI “deployment” fails to capture the realities needed for effective AI “integration” into organizational culture, staff workflows, and user practices.⁴⁰ As the USAID (2018) report notes, for AI to be integrated into development practice, “it’s not just plug and play.”

A useful way to orientate around the ethical dimensions technology in development contexts is through a sociotechnical approach, derived from the academic field of Science and Technology Studies. This approach assumes that systems are comprised of social and technical components that are intrinsically intertwined. Situating AI ethics within a sociotechnical approach gives practitioners a framework to simultaneously consider both social and technical aspects when assessing the intended and unintended consequences of technological interventions. Instead of starting with the technology, the focus is on the problems and the people *in relation to* technology. The USAID (2018) report suggests that development practitioners “bring context to the fore” as the deep expertise of local experts, who are closest to the problem, are needed to critically assess the viability of AI/ML tools.

Power

In international development, donor governments, aid organizations, and implementing partners often occupy an asymmetric power relationship to the populations they serve. The use of AI in any context raises any number of ethical issues around data protection; privacy; missing or non-representative data sets; inadequate training data; biased, unfair and discriminatory algorithms; lack of explainability, etc.

³⁶ Many of the lessons learned in the ICT4D space over the years are articulated in the Digital Principles for Development (launched in 2014) <https://digitalprinciples.org/>

³⁷ OLPC’S \$100 Laptop was Going to Change the World – Then it All Went Wrong.

<https://www.theverge.com/2018/4/16/17233946/olpcs-100-laptop-education-where-is-it-now>

³⁸ Evgeny Morozov (2013). To Save Everything, Click Here.

³⁹ Andrew Selbst, danah boyd, Sorelle Friedler, Suresh Venkatasubramian, Janet Vertesi. (2018). Fairness and Abstraction in Sociotechnical Systems. ACM FAT* Conference

⁴⁰ See Alexandra Mateescu and Madeleine Clare Elish (2018). AI in Context. Data & Society.

Should ethical considerations around the risks and benefits of AI be heightened due to existing power imbalances?

For example, privacy, which often serves as a shorthand for ethics, has been a central concern in the Data4Development space. While UN Global Pulse's privacy principles,⁴¹ released in 2016, represents how far the space has come over the years, data privacy and protection have a renewed importance for AI applications because of the vast amounts of personal and behavioral data needed for machine learning applications. But the risks of collecting, sharing, and publicizing data have been shown to be unknowable. In such cases, Arvin Narayanan et al. (2015) argue for a "precautionary approach" where the unknown risks around data privacy increases the burden and responsibility on those in control of the data.

In the development context, the data collectors and controllers will likely be large international institutions or government funded organizations who already sit in an asymmetric power relationship with the data subjects. Do the technical risks around data privacy/protection amplify power asymmetries on the ground? Should the precautionary principle prevail? A counter argument is the counterfactual: What is the cost of not intervening with technology when there is a possibility it can help someone in need? Is there a moral argument for using a tool that increases the risk to privacy for an individual or group even if it decreases the risk of physical harm?⁴²

Similar debates can be had over any number of unresolved ethical issues in AI, such as bias in data sets and algorithmic discrimination. A number of studies have demonstrated the disparate impact of algorithmic bias on the poor and disenfranchised.⁴³ A key ethical question is whether the AI/ML interventions in development contexts are designed to empower the communities in need or will these applications exacerbate and entrench existing power imbalances.

Accountability

Tech ethics is fundamentally about establishing a set of values that will guide behavior and creating processes to implement those values in everyday practice. Despite the recent proliferation of AI principles, there is no agreed upon ethical standard for AI applications in international development (or any other domain). This presents a problem of recourse for when what can wrong, does.

For AI to benefit the common good its development and usage should follow the necessary processes to safeguard against the risks and harms to fundamental human values.⁴⁴ This raises the question as to what fundamental values, principles, and minimum standards are appropriate in development contexts.

Human rights has emerged as potential governance framework for AI that covers an important class of harms based on values established through 70 years of international deliberative and legal processes.⁴⁵ For example, IEEE's 2019 report *Ethically Aligned Design* establishes the first principle that

⁴¹ <https://www.unglobalpulse.org/privacy-and-data-protection>

⁴² Mark Latonero (2018). Big Data Analytics and Human Rights: Privacy Considerations in Context. Chapter in Molly Land and Jay Aronson (Eds.) Technology and Human Rights.

⁴³ Virginia Eubanks (2018). Automating Inequality; Safiya Noble (2018). Algorithms of Oppression; Solon Boracas & Andrew Selbst. (2016). Big Data's Disparate Impacts. California Law Review

⁴⁴ Mark Latonero (2018). Governing Artificial Intelligence: Upholding Human Rights and Dignity. Data & Society Research Institute.

⁴⁵ F. Raso, V. Krishnamurthy, et al. (2018) "Artificial Intelligence and Human Rights: Opportunities & Risks," Berkman Klein Center for Internet & Society Research Publication; Mark Latonero (2018).

autonomous/intelligent systems “shall be created and operated to respect, promote, and protect internationally recognized human rights.”⁴⁶ An added value of a human rights based approach is that it is rooted in international law, and is thus linked to remedy and redress. Human rights working alongside ethical frameworks can help establish AI “governance with teeth.”⁴⁷

Other frameworks that are relevant in the development context include international humanitarian law and regulations for human subjects research. Yet the nuances in AI sociotechnical systems may not be easily subsumed under any existing framework. Who is accountable for AI systems that are making decisions autonomously, without a human in the loop? Can methods of contestation and accountability be achieved if the reasoning behind a neural network’s output cannot be explained? These quandaries may necessitate new international frameworks or new tools like algorithmic impact assessments.

During a 2018 USAID training on AI/ML in international development in Bangkok, Thailand, a key learning point emerged: “ethics in AI needs to focus beyond ‘bias’ and investigate existing problems in terms of oppression and social injustice.” AI policies that take into account context and power relationships; and seek to establish fundamental values, ethical processes, and mechanism for accountability can provide a starting point for upholding values such as justice in the international development context.

⁴⁶ Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems, First Edition (2019). IEEE

⁴⁷ Vidushi Marda (2019). Governance with teeth: How human rights can strengthen FAT and ethics initiatives on artificial intelligence. Article 19.

https://www.article19.org/wp-content/uploads/2019/04/Governance-with-teeth_A19_April_2019.pdf