A Conversation with Akhil Mathur: Toward the vision of equitable AI

INNOVATION LEADERS

by Bushra Anjum

Editor's Introduction

Ubiquity's senior editor Dr. Bushra Anjum sits with Akhil Mathur, a principal research scientist at Nokia Bell Labs in Cambridge (United Kingdom). Both discuss improving the robustness of AI models deployed on mobile and wearable devices to minimize the accuracy degradation as these devices are used in new domains.

A Conversation with Akhil Mathur: Toward the vision of equitable AI

INNOVATION LEADERS
by Bushra Anjum

Akhil Mathur is a principal research scientist at Nokia Bell Labs in Cambridge (United Kingdom) where he does research on Mobile Sensing and Deep Learning for Edge Devices. His current focus is on developing techniques to improve the robustness and execution latency of deep learning algorithms on edge devices. Previously, he has done research in the fields of Ubiquitous computing, HCI and Technology for Development. Mathur is a member of the ACM Future of Computing Academy (FCA) and serves on the editorial board of ACM IMWUT journal. He has been awarded with the mBillionth Award South Asia, Wolfond Fellowship at the University of Toronto, and the President's Gold Medal for his undergraduate studies. His research has been covered by several media organizations including the *New Yorker*, *Financial Times*, *Livemint* and Canadian Broadcasting Corporation.

What is your big concern about the future of computing to which you are dedicating yourself?

My long-term vision is to see AI technologies work equitably for users around the world and in particular, I focus on those AI technologies that process sensor data collected from a user's mobile and wearable devices to generate meaningful insights.

In recent years, we have witnessed the exciting capabilities and opportunities that AI can bring to end users. For example, there is the Seeing AI [1] app from Microsoft, which assists users with visual impairments in identifying people, text, or objects near them, and the AI-based biometrics platform by Element Inc. [2], which aims to provide a digital identity to infants and children in the Global South. While such technologies are undoubtedly promising, I believe it is important to ensure that they work—universally—for users around the world and not just be limited to a subset of the population, otherwise we risk creating a world of AI haves and have-nots, i.e., people who will gain advanced capabilities in their lives with AI and others who will not.

What is the challenge in fulfilling this vision? Many of today's AI solutions, such as those based on supervised machine learning, require collections of large-scale labelled data to train an AI model. Moreover, a fundamental assumption behind these solutions is that the distribution (or characteristic) of the data will remain the same between training and deployment stages. If there is any divergence between the training and test data distributions (a phenomenon known

as dataset shift), the accuracy of supervised machine learning algorithms can degrade. Naturally, as AI technologies become more widespread, such divergences between training and test data are more likely to occur—for example, users in different parts of the world will have diversity in the devices they use (ranging from low-end phones to very expensive smartphones), their environment conditions (e.g., ambient noise levels, natural lighting), and their technology usage patterns—all of which are likely to introduce unique variabilities in the user data, thus making it different from the original training data. This dataset shift, in turn, can degrade the accuracy of the AI model when it is deployed in new scenarios different from those in which it was trained.

Now one solution is to re-train the model for each scenario in which it will be deployed, however this involves repeating the collection of labeled training data, which makes this approach extremely time-consuming, expensive, and unscalable. Wouldn't it be better if we can adapt an existing AI model to new scenarios without collecting large-scale labeled training data from them? How do we develop such solutions? How will these solutions manifest in practical software systems? Answering these questions in the context of AI models on mobile and wearable devices is a key focus of my research.

While this is clearly an interesting computing problem to solve, I also believe it has very significant implications for wider society. Collecting large-scale labeled data to train AI models is expensive, and further scaling this data collection exercise to diverse user groups around the world will be infeasible for most AI developers. What will they do then? Will they build AI solutions that only benefit selected user populations, perhaps by taking commercial considerations into account? This would be particularly concerning and may worsen the technological and income inequalities that already exist in the world.

How did you become interested in creating equitable technical solutions?

In the early phases of my career, I did research on mobile-based human-centered computing. One of my projects, named CommunityWeb [1], focused on creating a mobile web platform in India that allows local content creators (e.g., citizen journalists and radio presenters) to share their audio content with users who own low-end feature phones with limited memory and computational capabilities. A core technical component of our solution relied on web-based streaming of audios, and while it worked flawlessly on our personal “high-end” smartphones using standard APIs, when we deployed it on “low-end” feature phones owned by our target users, the audio streaming quality was patchy and resulted in a very poor user experience. Lessons like these taught me that many technical components that we often consider as “standard” in modern computing systems are actually not robust and do not work well universally. I realized that it is important for technology researchers to ensure that our solutions work—not just on devices and in environments familiar to us—but also in diverse and often technically-challenging environments in order to reach large segments of the population.

Later while working on AI-based human sensing systems during my Ph.D. and career at Bell Labs, I started investigating similar challenges with robustness of AI models on mobile and wearable devices. Based on my experiments, I uncovered scenarios where an AI model for speech recognition works well on the smartphone for which it was trained, but its accuracy drops when it is tested on a different smartphone, or a model trained for human-activity recognition using accelerometer sensors suffers severe accuracy degradation if users change the placement or orientation of the sensor on their body. Although these challenges seemed familiar at a high-level, I was surprised that even state-of-the-art AI models, trained with hundreds of hours of data, could not deal with such subtle variations between training and test data. Clearly, this performance degradation will only get worse when AI models are deployed in more challenging and diverse real-world scenarios.

What project are you currently leading to improve the robustness of AI models so that they are able to adapt to changing scenarios?

A key focus of my research is to minimize the accuracy degradation of AI models when they are deployed in new scenarios (or domains) that are not represented in the labeled training dataset. To this end, we are exploring solutions based on unsupervised domain adaptation where the idea is to collect a small amount of unlabeled data in the new domain and re-train the AI model using adversarial learning [2] to boost its performance in the new domain. Using this approach, a user can adapt and personalize an existing AI model to work accurately in their own environment, even if it is different from the environment in which the AI model was initially trained. Most importantly, as this solution relies only on unlabeled data, it is cheap to deploy and significantly reduces the burden on users or AI developers to label the data to adapt the model for their needs.

We are also looking at the cases where it is not possible to re-train the AI model, perhaps because it is proprietary and its architecture is not publicly known. In these cases, we have designed a “data translation” solution [3], which transforms the test data collected from users in a way that it appears closer to the training distribution. This translation component can be added into any AI system pipeline and boosts the overall accuracy of the system.

My team is also working on making these solutions easy to integrate in any AI system. To this end, we have developed a software framework named SensiX that provides high-level abstractions for system developers to seamlessly include our domain adaptation and data translation solutions in a mobile or embedded system.

These are however still early-stage solutions tested in controlled settings. In my role as a member of the ACM FCA, I am exploring possibilities of collaborating with other researchers and practitioners to test these solutions in more realistic settings. I am optimistic and excited

that these solutions will mature and reach end-users in the near future, and move us closer to the goal of equitable deployment of AI technologies in the society.

References

- [1] Mathur, A., and Jaiswal, S. Exploring the interplay between community media and mobile web in developing regions. ACM MobileHCI 2013.
- [2] Chang, Y., Mathur, A., and Kawsar, F. Unsupervised domain adaptation for inertial sensing models. To appear in ACM IMWUT 2020.
- [3] Mathur A., et al. Mic2Mic: Using cycle-consistent generative adversarial networks to overcome microphone variability in speech systems. ACM IPSN 2019.

Biography

Bushra Anjum is a software technical lead at Amazon in San Luis Obispo, CA. She has expertise in Agile Software Development for large scale distributed services with special emphasis on scalability and fault tolerance. Originally a Fulbright scholar from Pakistan, Dr. Anjum has international teaching and mentoring experience and has served in academia for over five years before joining the industry. In 2016, she has been selected as an inaugural member of the [ACM Future of Computing Academy](#), a new initiative created by ACM to support and foster the next generation of computing professionals. Dr. Anjum is a keen enthusiast of promoting diversity in the STEM fields and is a mentor and a regular speaker for such. She received her Ph.D. in computer science at the North Carolina State University (NCSU) in 2012 for her doctoral thesis on Bandwidth Allocation under End-to-End Percentile Delay Bounds. She can be found on Twitter [@DrBushraAnjum](#).

DOI: 10.1145/3397260