Explaining AI: Fairly? Well?

Margaret Burnett
School of EECS
Oregon State University
Corvallis, Oregon USA
burnett@eecs.oregonstate.edu

ABSTRACT

Explainable AI (XAI) has started experiencing explosive growth, echoing the explosive growth that has preceded it of AI becoming used for practical purposes that impact the general public. This spread of AI into the world outside of research labs brings with it pressures and requirements that many of us have perhaps not thought about deeply enough. In this keynote address, I will explain why I think we have a very long way to go.

One way to characterize our current state is that we're doing "fairly well", doing some explaining of some things. In a sense, this is reasonable: the XAI field is young, and still finding its way. However, moving forward demands progress in (at least) three areas.

(1) *How* we go about XAI research: Explainable AI cannot succeed if the only research foundations brought to bear on it are AI foundations. Likewise, it cannot succeed if the only foundations used are from psychology, education, etc. Thus, a challenge for our emerging field is how to conduct XAI research in a truly effective multi-disciplinary fashion, that is based on an integration of foundations behind what we can make AI algorithms do, with solid, well-founded principles of explaining the complex ideas behind the algorithms to real people. Fortunately, a few researchers have started to build such foundations.

(2) *What* we can succeed at explaining: So far, we as a field are doing a certain amount of cherry picking as to what we explain. We tend to choose what to explain by what we can figure out how to explain—but we are leaving too much out. One urgent case in point is the societal and legal need to explain fairness properties of AI systems.

The above challenges are important, but the field is already becoming aware of them. Thus, this keynote will focus mostly on the third challenge, namely:

(3) *Who* we can explain to. Who are the people we've even tried to explain AI to, so far? What are the societal implications of who we explain to well and who we do not?

Our field has not even begun to consider this question. In this keynote I'll discuss why we have to explain to populations to whom we've given little thought—diverse people in many dimensions, including gender diversity, cognitive diversity, and age diversity.

Addressing all of these challenges is necessary before we can claim to explain AI fairly and well.

CCS CONCEPTS

• **Computing methodologies** → *Intelligent agents*; • **Human-centered computing** → **Human-Computer Interaction (HCI)**

KEYWORDS

Explainable AI; explaining to diverse populations; biased explanations; XAI challenges

ACM Reference format:

Margaret Burnett. 2020. Explaining AI Fairly? Well? In 25th International Conference on Intelligent User Interfaces (IUI '20), March 17–20, 2020, Cagliari, Italy. ACM, New York, NY, USA, 2 pages.
<https://doi.org/10.1145/3377325.3380623>

BIOGRAPHY

Margaret Burnett (<http://web.engr.oregonstate.edu/~burnett/>) is an OSU Distinguished Professor at Oregon State University. She began her career in industry, where she was the first woman

IUI '20, March 17–20, 2020, Cagliari, Italy
© 2020 Copyright is held by the owner/author(s).
ACM ISBN 978-1-4503-7118-6/20/03.
<https://doi.org/10.1145/3377325.3380623>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

software developer ever hired at Procter & Gamble Ivorydale. A few degrees and start-ups later, she joined academia, with a research focus on people who are engaged in some form of software development. Together with her collaborators and students, she has contributed some of the seminal work on explaining AI to ordinary end users. She also co-founded the area of end-user software engineering, which aims to enable computer users not trained in programming to improve their own software, and co-leads the team that created GenderMag (<http://gendermag.org>), a software inspection process that uncovers user-facing gender biases in software from smart systems to programming environments. Burnett is an ACM Fellow and a member of the ACM CHI Academy.

ACKNOWLEDGMENTS

This work has been supported in part by DARPA #N66001-17-2-4030 and by NSF #1528061 and #1901031. Any opinions, findings and conclusions or recommendations expressed are the authors' and do not necessarily reflect the views of NSF, DARPA, the Army Research Office, or the US government.