

Creating Impact Through Community

ANNUAL REPORT 2021



PARTNERSHIP ON AI

From Our CEO

This year, PAI Partners came together to make important contributions to the design and deployment of AI that advances positive outcomes for people and society.

From the integrity of our information ecosystem and the future of work to the ethical risks of research and transparency through documentation, PAI is working with Partners in Africa, Asia, Europe, and North America to redesign AI for all.

To support research and action on misinformation in the public interest, our Media Integrity Program worked with Partners to develop an accessible database of different types of interventions used to address false and misleading information.

To assess the risks of AI on workers, our AI and Shared Prosperity Steering Committee published the first international agenda focused on measuring and disclosing the impact of the AI industry on the jobs market.

To catalyze a community commitment to responsible research publication practices, our Safety-Critical Team worked with researchers from industry, academia, and civil society to develop recommendations to anticipate and mitigate the potential negative impacts of AI research.

To promote transparency and accountability, our ABOUT ML Steering Committee released the leading online reference library of documentation best practices across the machine learning lifecycle.

During our strategic planning process, I learned how passionate our Partners are about building a truly inclusive AI community building deployable solutions. In all of these efforts and many others, PAI's impact in 2021 was only made possible by the diversity of perspectives, disciplines, and sectors engaged from across our international community of Partners.

I would like to thank PAI's Board of Directors and Staff Team for their passion and commitment every day.

Importantly, thank you to our Partners for your commitment to change. It is only with your passion that our vision will become a reality.

Sincerely,



REBECCA FINLAY
CEO, PARTNERSHIP ON AI

FEBRUARY 10, 2021



Our 95 Partners

North America

Canada
United States

68 PARTNERS

10 Industry
41 Nonprofit
15 Academic
2 Media

Europe

France
Germany
Ireland
United Kingdom

15 PARTNERS

2 Industry
6 Nonprofit
6 Academic
1 Media

Africa

Kenya

1 PARTNER

1 Nonprofit

Asia

China
India
Hong Kong
Japan
South Korea

9 PARTNERS

3 Industry
4 Nonprofit
2 Academic

Oceania

Australia
New Zealand

2 PARTNERS

1 Nonprofit
1 Academic

PAI is a non-profit community of academic, civil society, industry, and media organizations from across five continents addressing the most important and difficult questions concerning the future of AI.


By bringing these many stakeholders together, PAI establishes a common ground between entities that otherwise may not have cause to work together and, in so doing, serves as a unifying force for good in the AI ecosystem.

Spanning sectors, disciplines, and borders, PAI is bringing together diverse voices from across the AI community to share insights, create resources, and advance positive outcomes for people and society.

Dedicated to Impact

By bringing together diverse voices from across the AI community and around the world, PAI seeks not just to spark new ideas but to turn those ideas into lasting change.

At all points in our process—from convening stakeholders and creating resources to promoting what we’ve learned and practicing accountability—PAI is working to advance responsible governance and best practices so that AI can have the greatest possible benefit for people and society.



Creating Community

Equity in action
Collecting insights in Africa and South America
Coming together to envision the future




Informing the Public

Promoting wider understanding
Shedding light on barriers to diversity
Fostering AI literacy



Encouraging Policy Innovation

Informing AI public projects
Developing solutions to support workers
Correcting oversights in algorithmic fairness



Fostering Changes in Practice

Illuminating impacts on workers
Offering guidance to reduce harm
Making ML transparency transparent

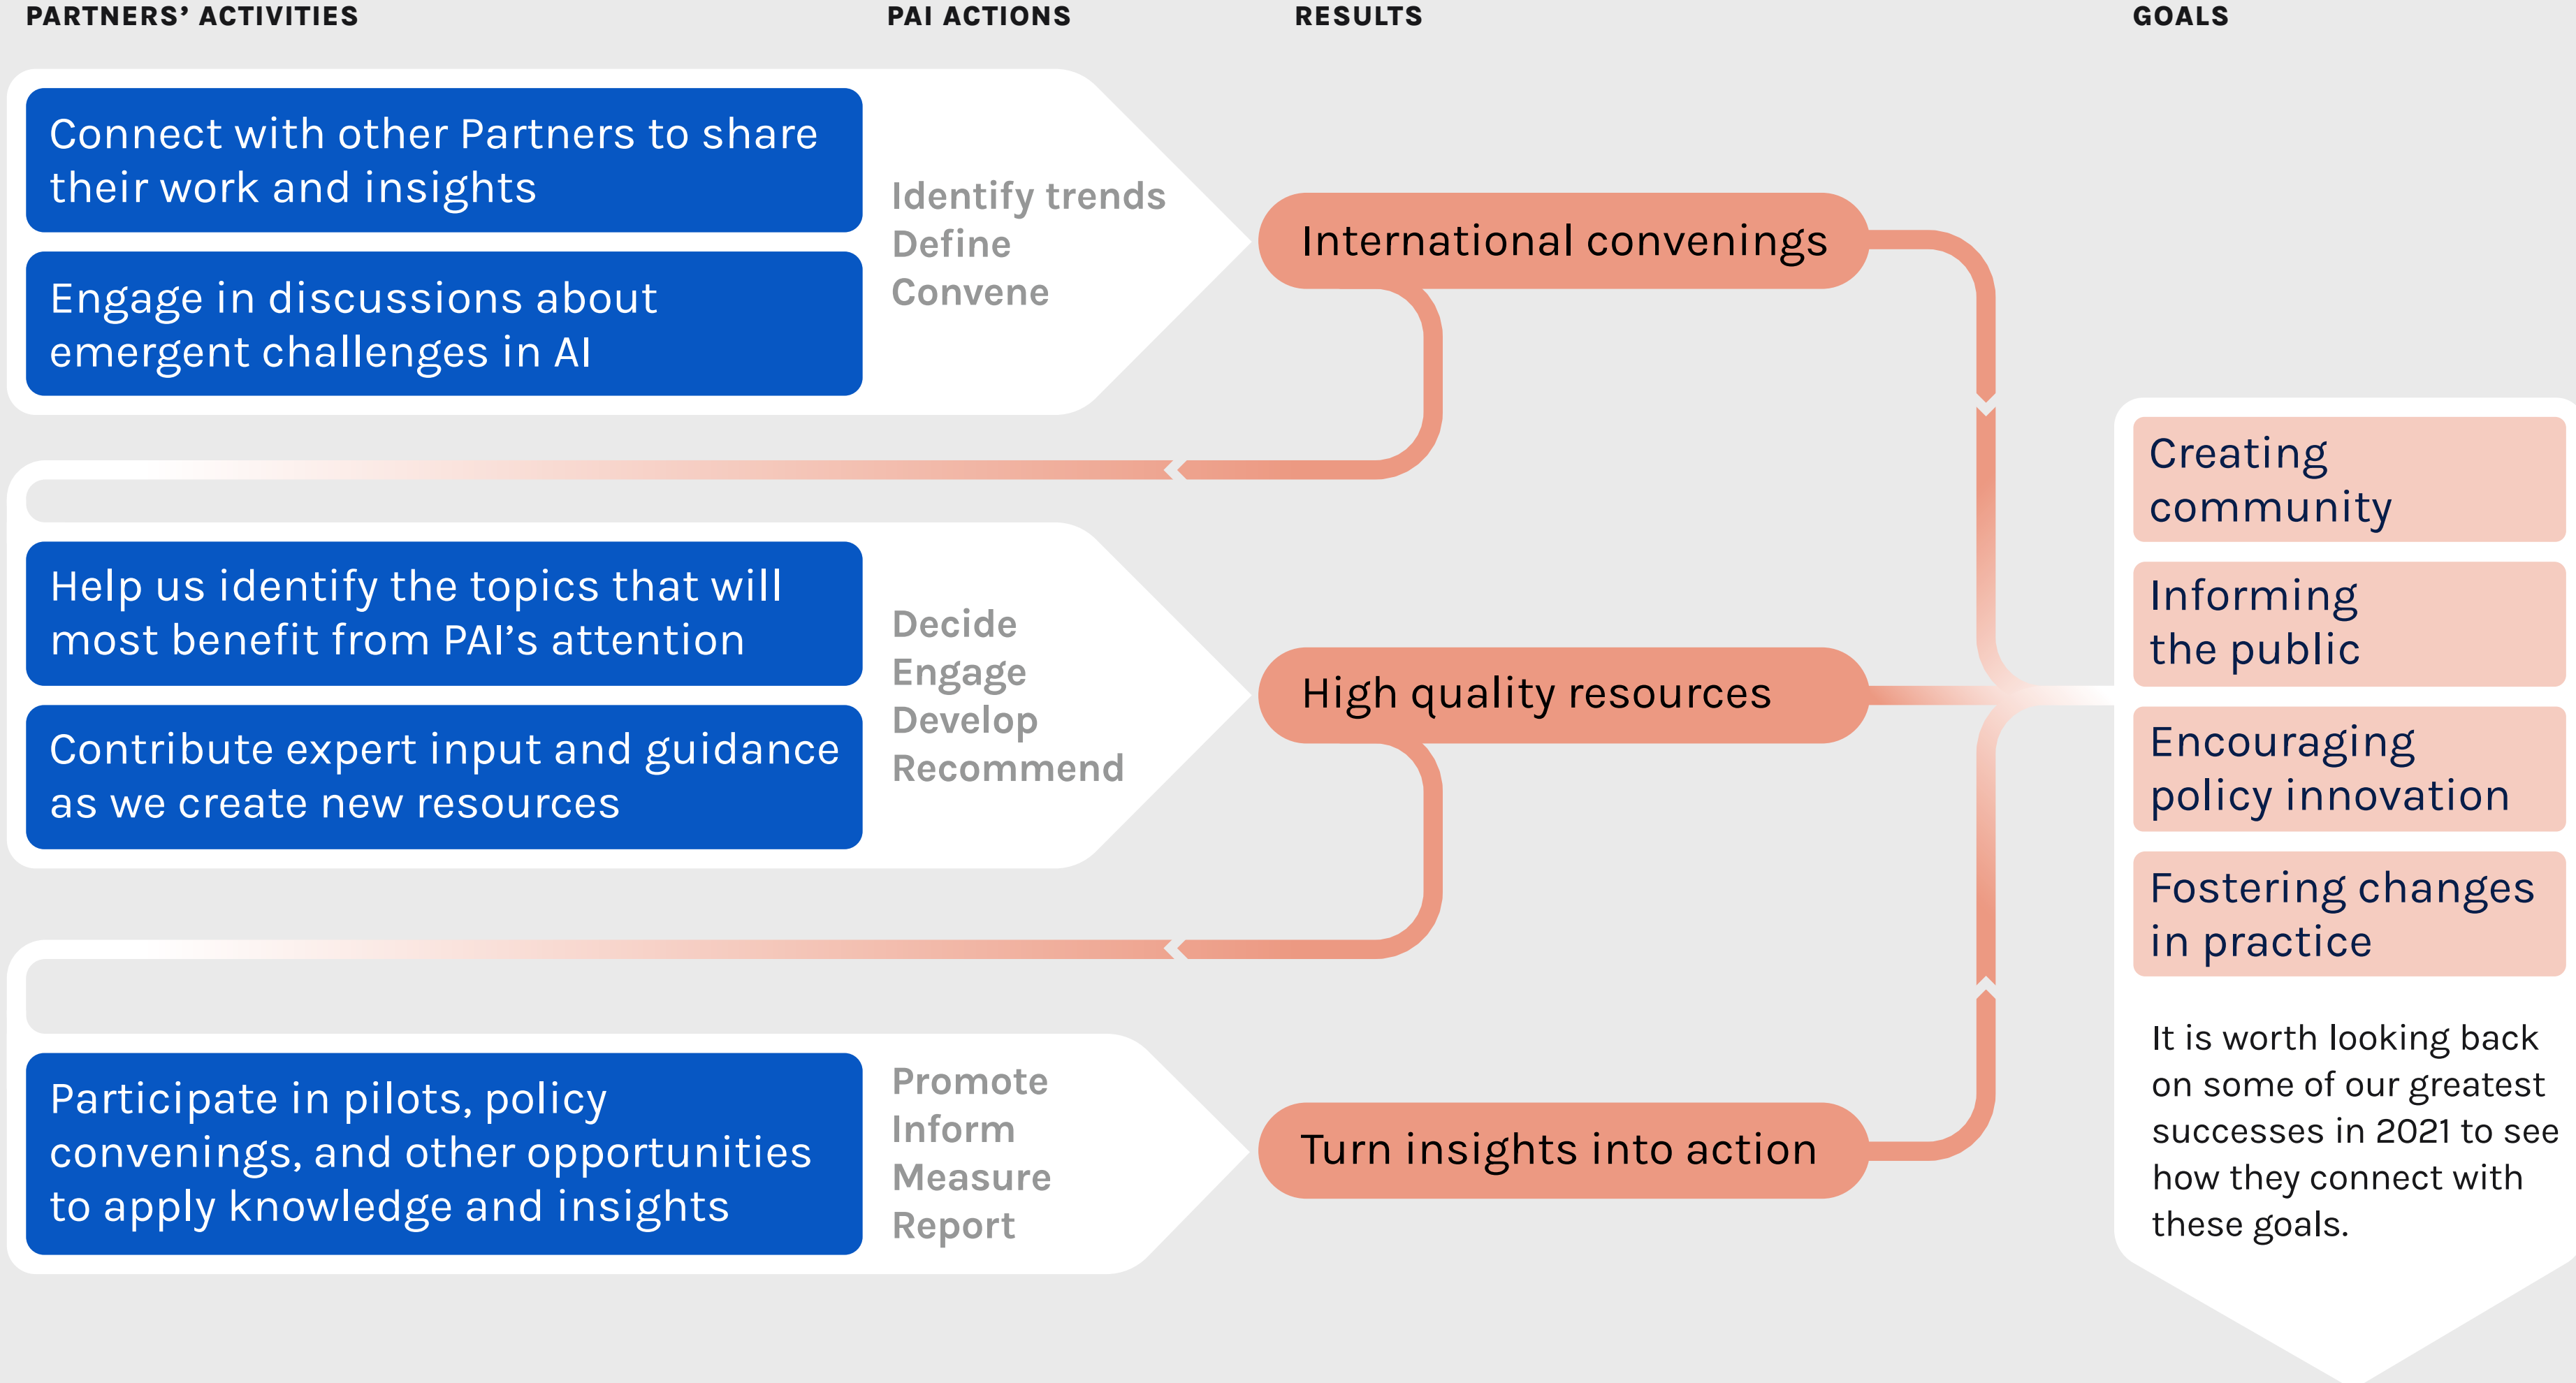
Read on to learn more about some of our most impactful work with Partners from the last year.

Get Involved

It is only through the dedication of our Partners that we can create the change we believe in.

To give our Partners more ways to engage with PAI and more clearly define our relationship with them, we renewed [our process](#) in 2021. This process has its foundation in our revised [Theory of Change](#).

As we operationalize this framework in 2022, we invite our Partners to further [connect with PAI](#) and get involved.



Creating Community

GOAL

An international, inclusive and equitable PAI community

Collecting insights in Africa and South America

PAI believes the best approaches to global challenges are informed by global perspectives, drawing from the lived and professional expertise of stakeholders from around the world. In the spring of 2021, PAI teamed up with human rights non-profit WITNESS to collect insights on one such challenge: the detection of AI-manipulated media. How can access

to the tools for detecting synthetic media be made more equitable without compromising their utility? To guide our work on this question, PAI and WITNESS hosted [two convenings](#) of journalists, activists, and researchers in South America and Africa, asking them who should have access to these tools and under what conditions.

Coming together to envision the future

In the next five years, what would we like to see happen in the AI space? And what can we do to help make those outcomes a reality? In December of 2021, PAI hosted a lively discussion addressing these questions at our Partner event [Partner Perspectives: The Next 5 Years in AI](#). To help spark this conversation, the event featured a panel discussion between representatives from PAI Partners

Data & Society, IBM, Meedan, and PolicyLink followed by breakout sessions facilitated by three of our Program Heads. Spaces like these are essential for building the connections and collaborations across the AI industry that will enable us to both envision a brighter future and make that future possible.

Equity in action

Fairness, equity, and inclusion are among PAI's core beliefs. In November 2021, PAI created a new position to operationalize our commitment to these values, naming Tina M. Park as our first [Head of Inclusive Research and Design](#). By integrating inclusive principles into

PAI's research, engagement, and communication practices, this role is designed to not only ensure our work is infused with our values, but to also provide models for the wider AI community on how to conduct research responsibly. With all our endeavors, PAI is working toward a truly inclusive AI community, creating new connections between our Partners and seeking out voices that have previously been excluded from the conversation.



Informing the Public

GOAL

A public that is better informed about the social and societal impact of AI

Promoting wider understanding

For AI systems to be truly transparent, they need to be understood not just by those directly interacting with them, but by the wider society these systems affect. In addition to PAI-produced materials, we promoted the public understanding of AI and its impacts last year through a variety of [media appearances, speaking engagements, and public events](#). In 2021, discussions of PAI's work could be found in publications such as Wired, The New Yorker, and The Washington Post.

And on podcasts like Towards Data Science and How AI Happens, PAI staff members shared their expertise with new audiences to educate the public about how AI intersects with important topics including diversity, safety, and the economy.

Shedding light on barriers to diversity

Why does the AI field struggle to attract and retain diverse talent and what can be done about it? Important answers can be found in the experiences of these professionals themselves, the focus of a [three-part series](#) authored by PAI Fellow Jeffrey Brown in the fall of 2021.

Drawing from in-depth interviews with more than 40 present and former tech workers, this series drew attention to an important (but under-examined) barrier to greater diversity in AI: worker attrition. By sharing insights from these interviews and a forthcoming PAI study analyzing common themes among them, PAI educated readers on a too-often-overlooked aspect of AI's "diversity problem."

Fostering AI literacy


In addition to talks and workshops at conferences like ACM FAccT, ICLR, and RightsCon, PAI co-produced several public events in 2021 to share our insights and provoke deep conversations about AI.

[With the USC Shoah Foundation](#), PAI hosted a panel of experts to discuss how synthetic media will affect the preservation of Holocaust testimony and other historical truths. [At the San Francisco Public Library](#), PAI moderated a discussion of AI's impact on workers and how we can guide this technology to utopian ends. And [at the JournalismAI festival](#), PAI led a conversation on how to bridge gaps between AI ethics and journalistic standards.

57

Speaking engagements in 2021

Encouraging Policy Innovation



Informing public AI projects

Central to PAI's mission is sharing our insights with key actors, including those who develop public policy. Following the passage of the National Artificial Intelligence Initiative Act last year, PAI was pleased to respond to [two requests for information](#) from the U.S. government to inform new AI projects. Our research-based resources are designed to not only synthesize knowledge but to make that knowledge actionable. As the U.S. builds its National Artificial Intelligence Research Resource and Artificial Intelligence Risk Management Framework, it will be able draw from PAI's important work concerning transparency, inclusion, risk documentation, and responsible publication norms.



“AI and data have become new pillars impacting Future of Work and workers... [PAI's Redesigning AI for Shared Prosperity] Agenda is very relevant and crucial.”

REEMA NANA VATY

Executive Director, Self-Employed Women's Association (SEWA)

Developing solutions to support workers

When the effects of automation on the job market are discussed, it is usually workers who are expected to adjust to the impact of AI. In May of 2021, PAI released [“Redesigning AI for Shared Prosperity: An Agenda,”](#) which calls instead for the creation of shared prosperity targets: verifiable criteria AI companies must meet to support the future of workers. This approach can help guide innovation

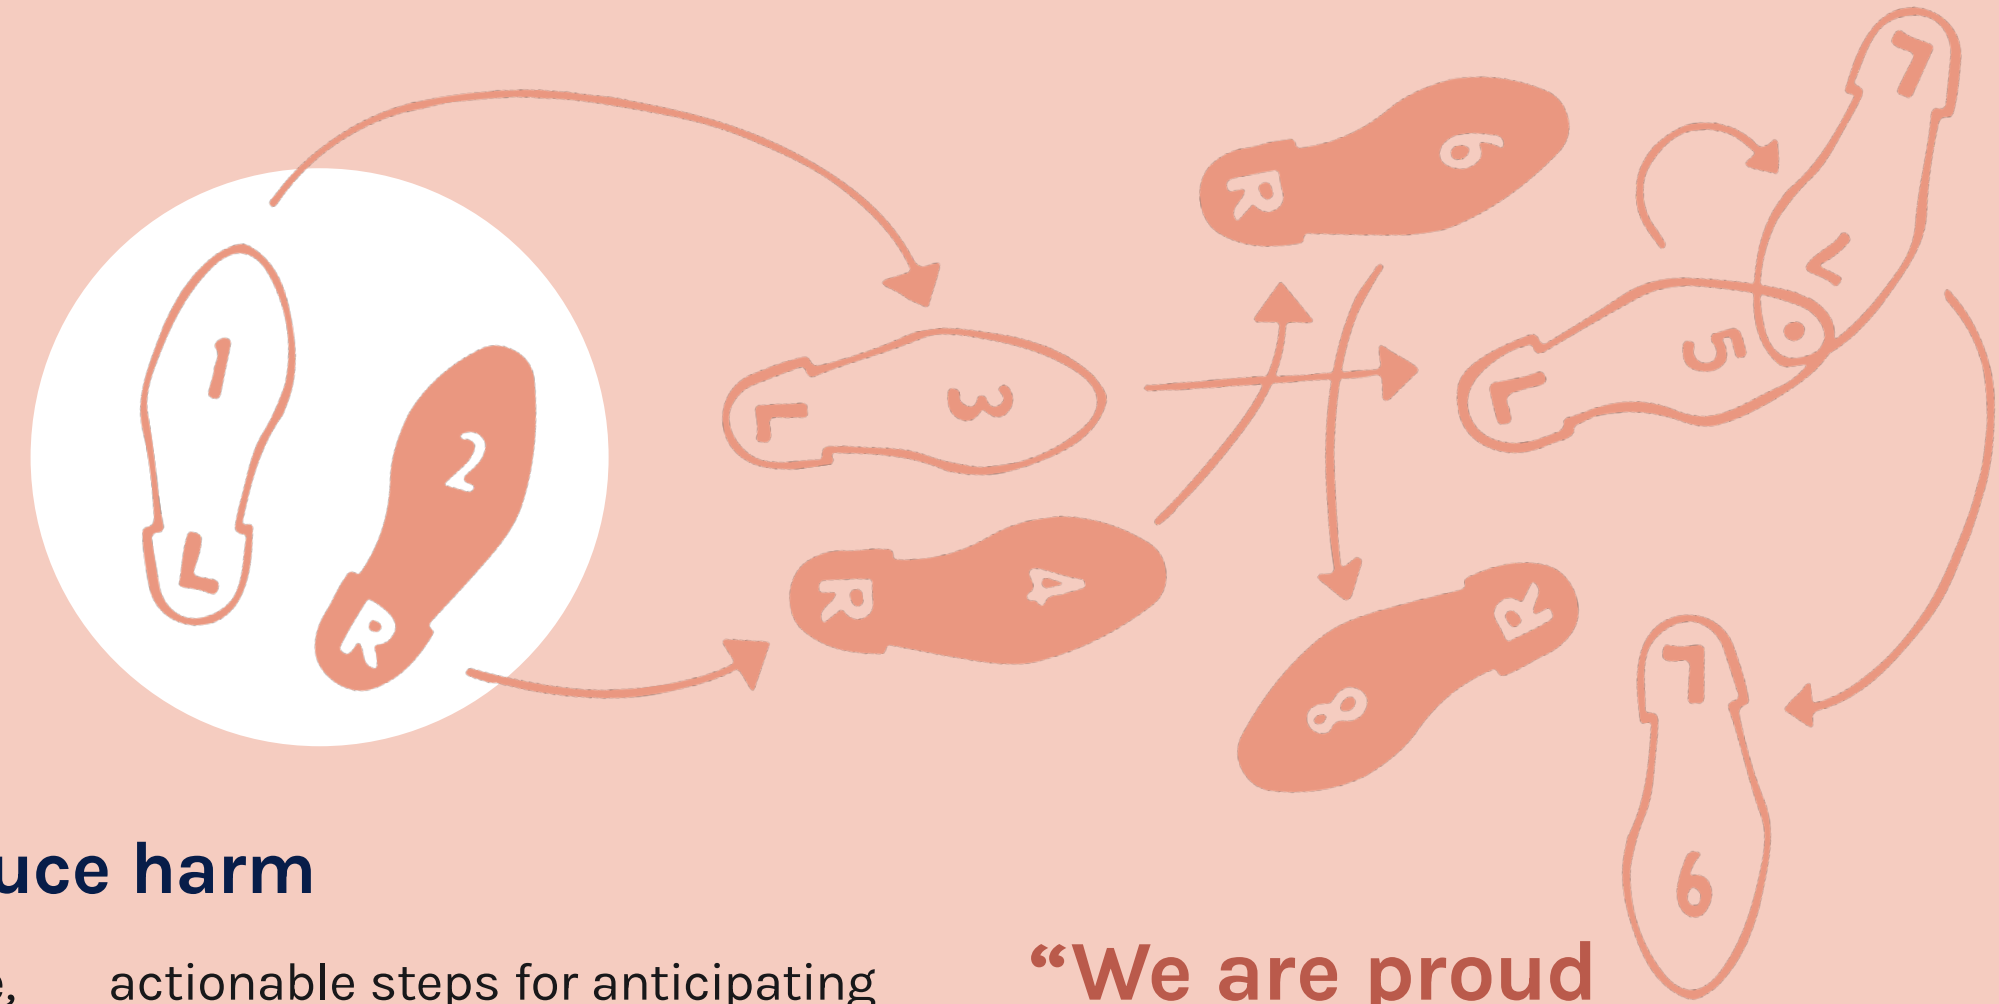
away from automating human tasks and toward complementing worker abilities, raising productivity while making jobs safer, more stable, and less physically exhausting. PAI is now working to develop these targets, a tool organizations can use to ensure new products do not harm workers and verify their commitment to a world where AI enriches us all.

Correcting oversights in algorithmic fairness

To have their intended outcomes, attempts to make AI fairer must be made with a full understanding of their consequences for marginalized groups. Common to many current approaches is the use of demographic data. Too frequently, however, the trade-offs associated with collecting and using this data are overlooked. In December of

2021, PAI published the white paper [“Fairer Algorithmic Decision-Making and Its Consequences,”](#) which seeks to correct this oversight. By cataloging the risks and benefits associated with using demographic data for algorithmic fairness, the paper pushes fairness efforts to address deeper questions about data governance and ultimate impact.

Fostering Changes in Practice



GOAL

Changes in practice by PAI Partners in all sectors and broader communities

Illuminating impacts on workers

As demand for AI models grows, so too does the demand for high-quality training datasets, making data enrichment workers a growing labor force whose well-being must be considered. In the spring of 2021, PAI released “[Responsible Sourcing of Data Enrichment Services](#),” a white paper explaining how key decisions in the AI development process can have a meaningful impact on the lives of these workers. These recommendations were informed by a five-part workshop series with more than 30 professionals in the data enrichment ecosystem, including representatives of PAI Partners BBC, BSR, IBM, Intel, Microsoft, Mozilla, and OpenAI. PAI then discussed our findings in internal town halls with multiple industry Partners.

Offering guidance to reduce harm

As AI research continues to advance, it is essential that we mitigate the risks related to its dissemination, including malicious use, accidents, and other unintended consequences. In May 2021, PAI published [a white paper](#) providing six recommendations for responsibly publishing AI research. Directed at three key groups in the research ecosystem, the recommendations offered

actionable steps for anticipating potential harms. Soon after, the journal *Nature Machine Intelligence*, which contributed to the white paper, published [an editorial](#) supporting PAI’s call to action. Pointing to its own new requirement for ethical impact statements on certain papers, the journal praised the white paper’s “timely advice.”

Making machine learning transparency transparent

For PAI’s tools, recommendations, and other resources to have the impact they need to, they must be both high-quality and highly accessible. In September of 2021, PAI published the revised [ABOUT ML Reference Document](#), which identifies transparency goals for machine learning systems and offers suggestions on how they might be

achieved. This revision features new recommended reading plans and other enhancements to make the text more consumable. Additionally, PAI created the [ABOUT ML Resource Library](#), featuring the initiative’s collected outputs designed to help organizations and individuals begin implementing transparency at scale.

“We are proud to join the Partnership on AI to further its collaborative, cross-industry approach to address the challenges of misleading and harmful content online.”

DANA RAO

Executive Vice President, General Counsel and Corporate Secretary, Adobe

Thanks to Our Funders

Philanthropic Funders

Thank you to our philanthropic funders. Your support makes PAI's work possible.



Corporate Funders



Our Funding Sources

PAI is an independent, nonprofit 501(c)(3) organization and funded by charitable contributions from philanthropy and corporate entities.

Our funding is differentiated as follows:

1 GENERAL OPERATING FUNDS

These funds support the work of the whole organization. Many philanthropic organizations choose to fund PAI in this way when our work and its intended impact align with their institutional vision.

2 SPECIFIC PROJECT FUNDS

These funds support the work of specific projects at PAI. Through strategic program development, funders collaborate with PAI on a specific project with explicit outputs and outcomes.

3 CHARITABLE CONTRIBUTIONS FROM OUR CORPORATE PARTNERS

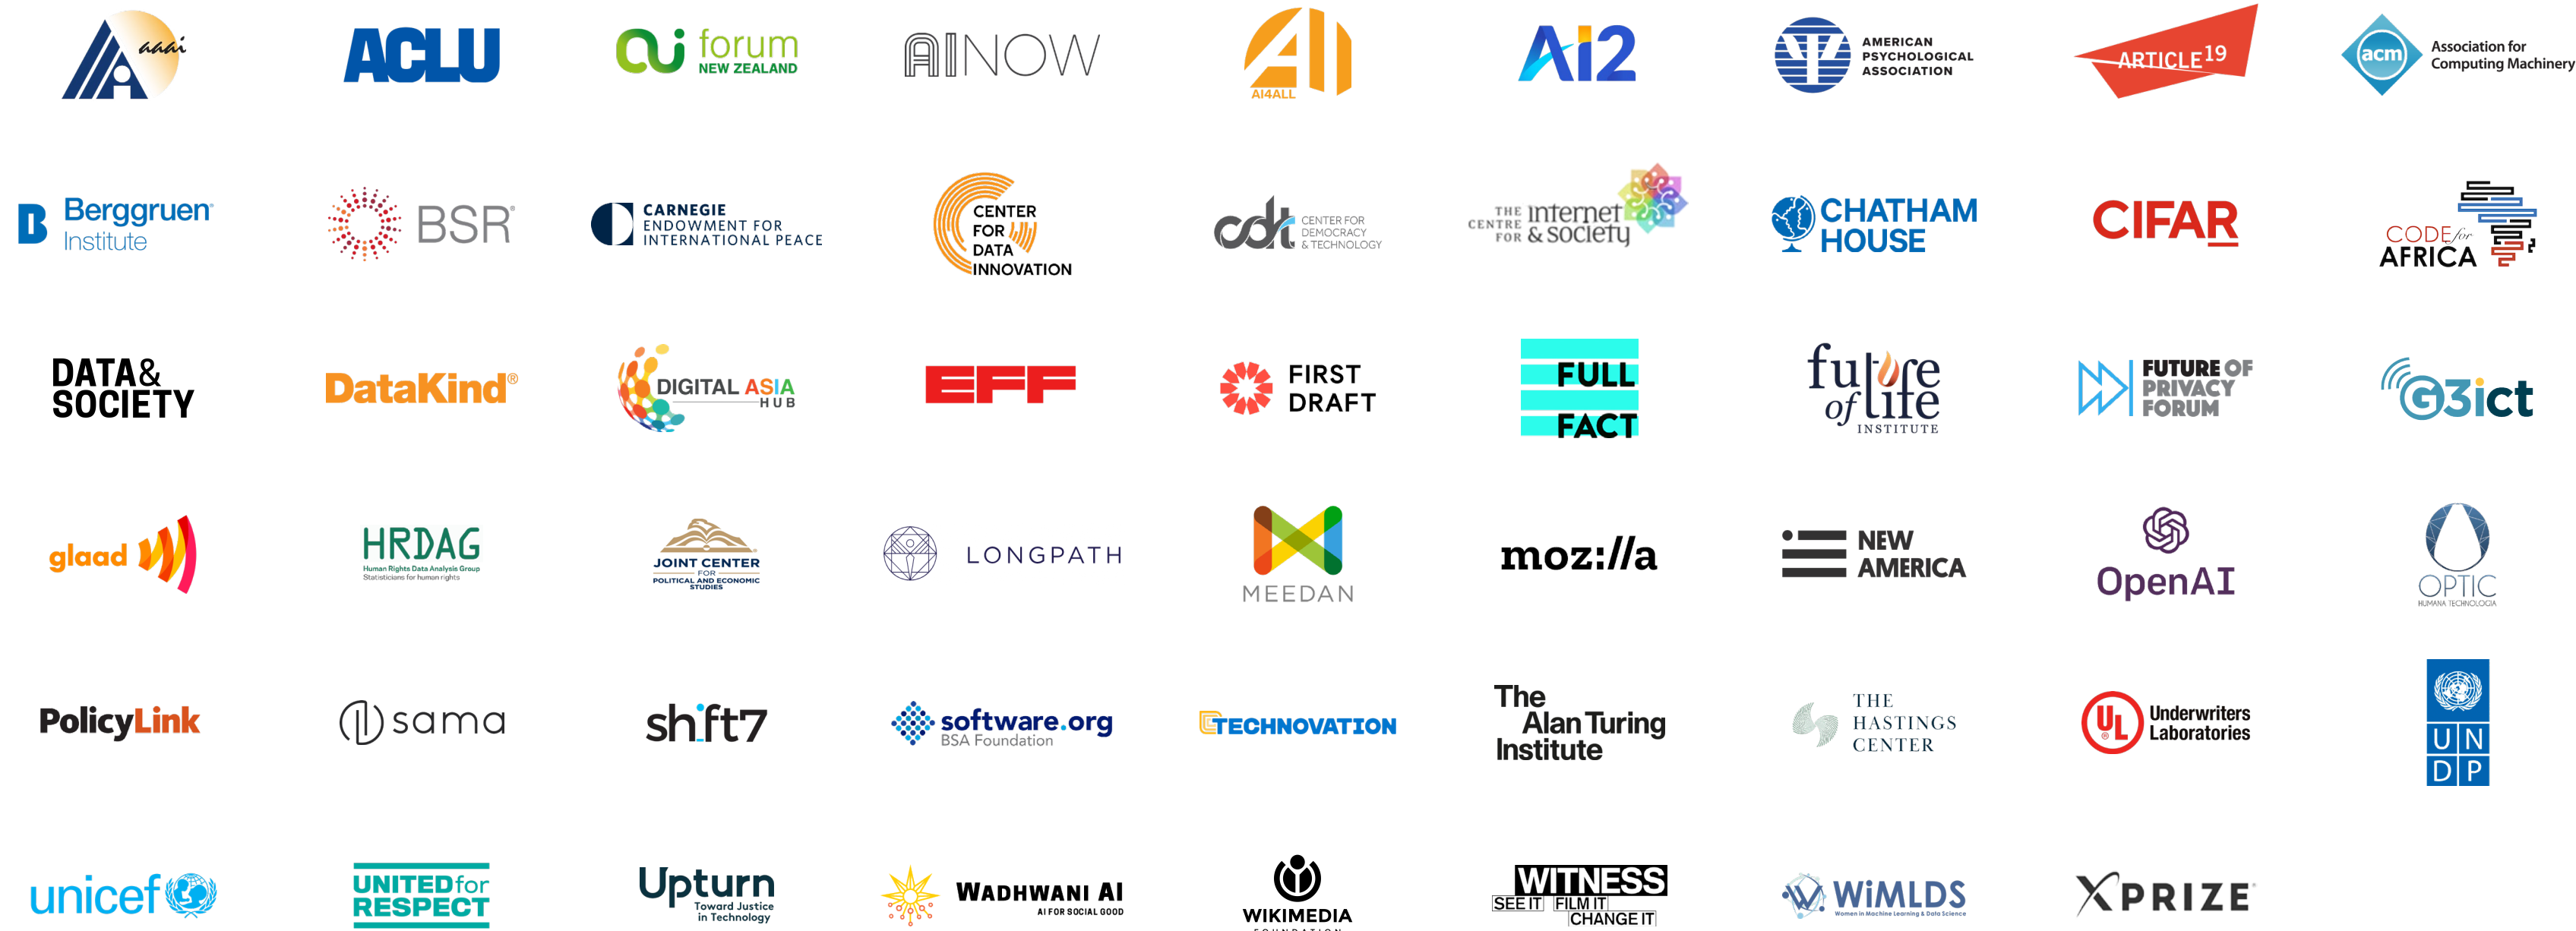
Corporate charitable contributions serve PAI as general operating support. These funds are legally classified as non-earmarked charitable contributions and not donations in exchange for goods or services, or quid pro quo contributions. This ensures our independence as an organization and in all areas of our work.

Thanks to Our Partners

PAI works with numerous organizations and individuals in a variety of sectors.

We thank all of our Partners who form this special community for their valuable insight and support.


NONPROFIT PARTNERS



MEDIA PARTNERS



INDUSTRY PARTNERS



ACADEMIC PARTNERS





PARTNERSHIP ON AI

partnershiponai.org

© 2022