

EUROPEAN
**BIG DATA
VALUE** FORUM
VERSAILLES, NOV 21-23 2017



Partnership on AI

Bernard Ourghanlian
Chief Technology & Security Officer
Microsoft – Member of Partnership on AI

What is it?



Partnership on AI
to benefit people and society

Partnership on AI

Partnership on AI (full name Partnership on Artificial Intelligence to Benefit People and Society) is a technology industry consortium focused on

- 1. Studying and formulating best practices on AI technologies,**
- 2. to advance the public's understanding of AI,**
- 3. to serve as an open platform for discussion and engagement about AI and its influences on people and society.**



What is it?

Partnership on AI

Publicly announced September 28, 2016, its **founding members are Amazon, Facebook, Google, DeepMind, Microsoft, and IBM**, with interim co-chairs Eric Horvitz of Microsoft Research and Mustafa Suleyman of DeepMind.

Apple joined the consortium as a founding member in January 2017. In January 2017, Apple head of advanced development for Siri, Tom Gruber, joined the Partnership on AI's board

Why? Why now?

Partnership on AI

Over the past few years, there has been a surge in real-world applications of AI with the rollout of technologies that use advances in learning, perception, and natural language. Today, there are vibrant ongoing discussions about how to maximize the value of these new applications and services, and about the potential societal influences of the technologies, including issues around

- **Ethics,**
- **Economics,**
- **Privacy,**
- **Transparency,**
- **Bias and Inclusiveness,**
- **Trustworthiness,...** of the technologies.

Why? Why now?

Partnership on AI

The founding research scientists care about these issues too, and see the opportunity for ongoing cross-industry discussion and the development of best practices to help gain the most from AI technologies for the benefit of people and society.

This group foresees great societal benefits and opportunities ahead, but we also understand that as with every new technology there will be concerns and confusion associated with new applications and competencies.

We look forward to working together on these important issues, including ethics, safety, transparency, privacy, biases, and fairness.

Why? Why now?

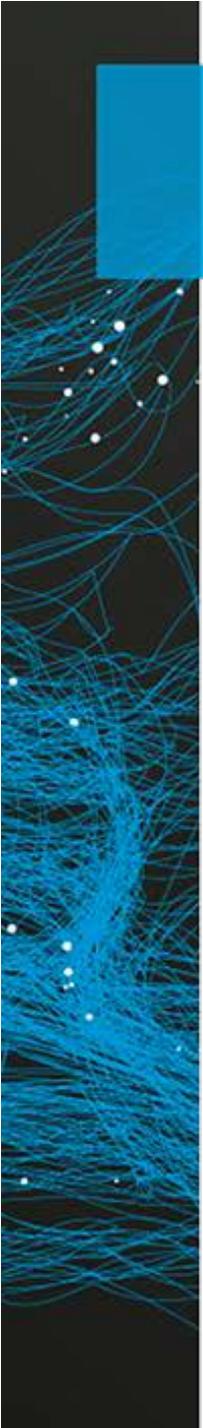
How did the idea for this organization come about?

Amazon, Apple, DeepMind/Google, Facebook, IBM, and Microsoft have been developing AI-related technologies for years.

Beyond their work on core R&D efforts, research scientists at these companies have been thinking about and discussing the potential societal impact of AI systems and how potential concerns might be addressed.

Conversations have occurred at workshops, conferences, and smaller meetings. Several research scientists at the companies kicked off informal discussions about the possibility of bringing their companies together to form a non-profit organization charged with exploring and developing best practices.

We are pleased those discussions culminated in the formation of the Partnership on AI.



Thematic pillars

1. SAFETY-CRITICAL AI

Advances in AI have the potential to improve outcomes, enhance quality, and reduce costs in such safety-critical areas as healthcare and transportation. Effective and careful applications of pattern recognition, automated decision making, and robotic systems show promise for enhancing the quality of life and preventing thousands of needless deaths.

However, where AI tools are used to supplement or replace human decision-making, we must be sure that they are safe, trustworthy, and aligned with the ethics and preferences of people who are influenced by their actions.

We will pursue studies and best practices around the fielding of AI in safety-critical application areas.

Thematic pillars

2. FAIR, TRANSPARENT, AND ACCOUNTABLE AI

AI and Data can be harnessed to develop useful diagnostic systems and recommendation engines, and to support people in making breakthroughs in such areas as biomedicine, public health, safety, criminal justice, education, and sustainability.

While such results promise to provide great value, we need to be sensitive to the possibility that there are hidden assumptions and biases in data, and therefore in the systems built from that data. This can lead to actions and recommendations that replicate those biases, and suffer from serious blindspots.

Researchers, officials, and the public should be sensitive to these possibilities and we should seek to develop methods that detect and correct those errors and biases, not replicate them. We also need to work to develop systems that can explain the rationale for inferences.

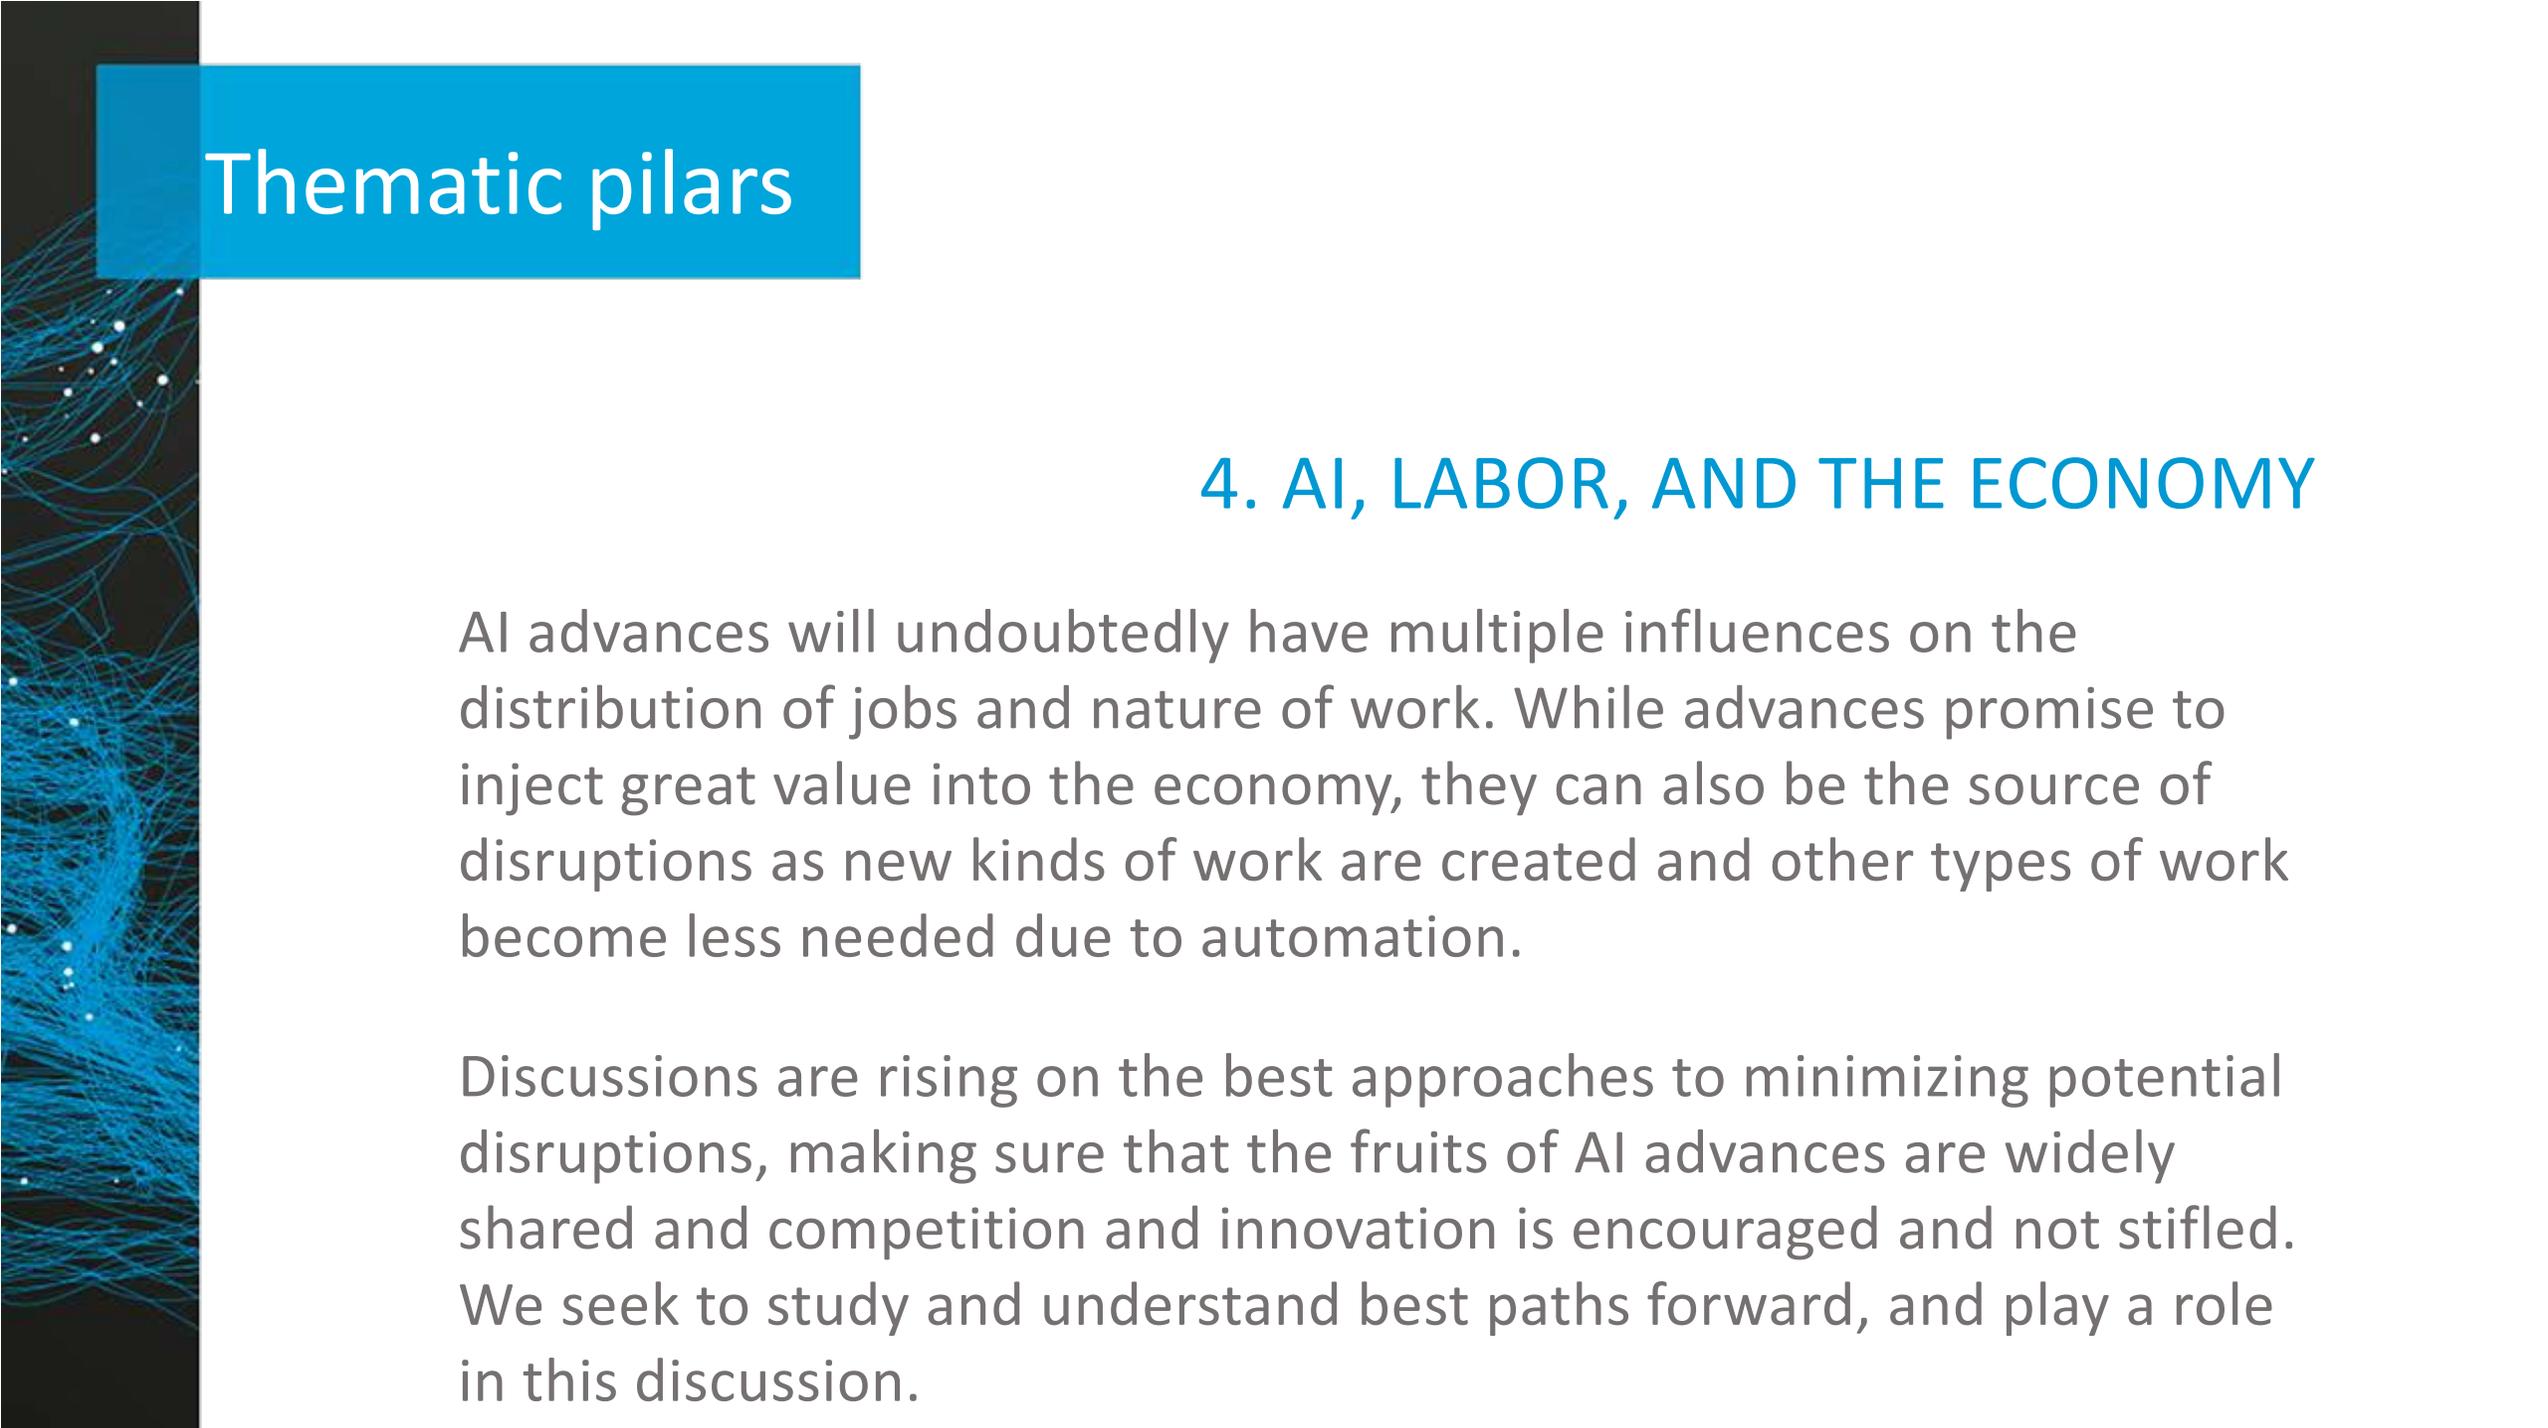
We will pursue opportunities to develop best practices around the development and fielding of fair, explainable, and accountable AI systems.

Thematic pillars

3. COLLABORATIONS BETWEEN PEOPLE AND AI SYSTEMS

A promising area of AI is the design of systems that augment the perception, cognition, and problem-solving abilities of people. Examples include the use of AI technologies to help physicians make more timely and accurate diagnoses and assistance provided to drivers of cars to help them to avoid dangerous situations and crashes.

Opportunities for R&D and for the development of best practices on AI-human collaboration include methods that provide people with clarity about the understandings and confidence that AI systems have about situations, means for coordinating human and AI contributions to problem solving, and enabling AI systems to work with people to resolve uncertainties about human goals.

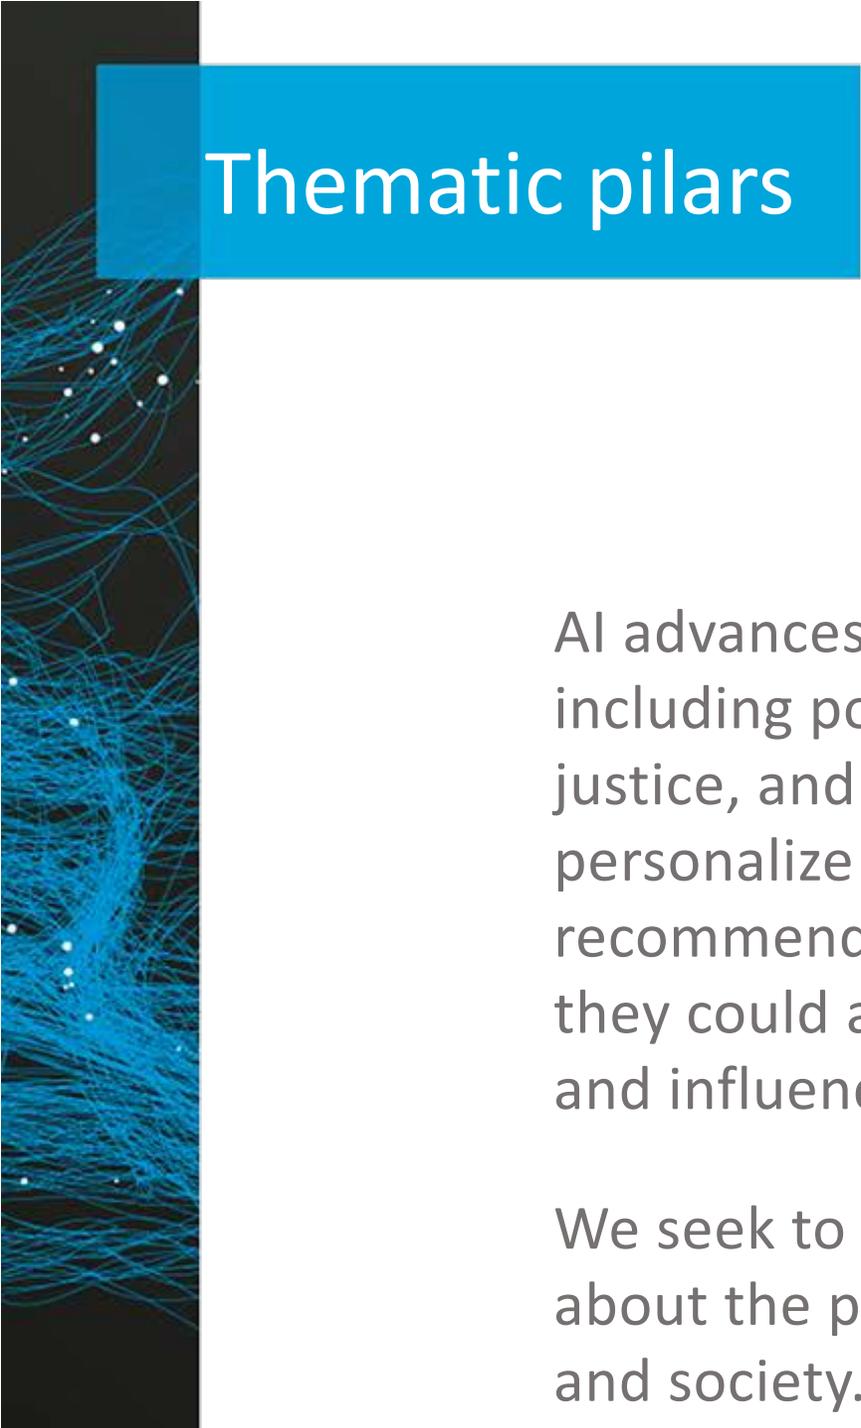


Thematic pillars

4. AI, LABOR, AND THE ECONOMY

AI advances will undoubtedly have multiple influences on the distribution of jobs and nature of work. While advances promise to inject great value into the economy, they can also be the source of disruptions as new kinds of work are created and other types of work become less needed due to automation.

Discussions are rising on the best approaches to minimizing potential disruptions, making sure that the fruits of AI advances are widely shared and competition and innovation is encouraged and not stifled. We seek to study and understand best paths forward, and play a role in this discussion.

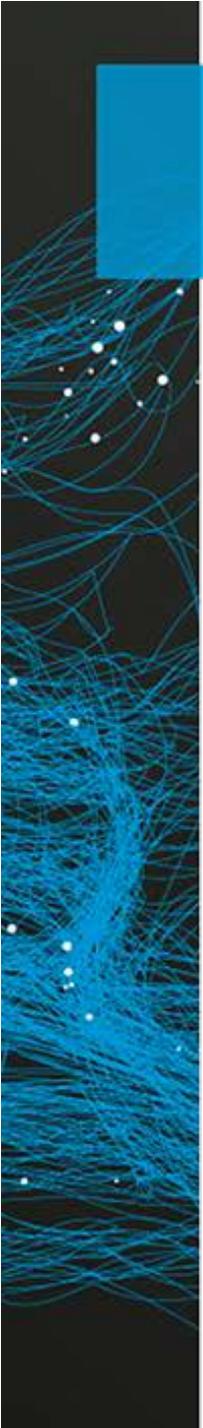
A decorative blue bar at the top left contains the text 'Thematic pillars'. To the left of the main content area is a vertical graphic of a network with blue lines and nodes on a dark background.

Thematic pillars

5. SOCIAL AND SOCIETAL INFLUENCES OF AI

AI advances will touch people and society in numerous ways, including potential influences on privacy, democracy, criminal justice, and human rights. For example, while technologies that personalize information and that assist people with recommendations can provide people with valuable assistance, they could also inadvertently or deliberately manipulate people and influence opinions.

We seek to promote thoughtful collaboration and open dialogue about the potential subtle and salient influences of AI on people and society.

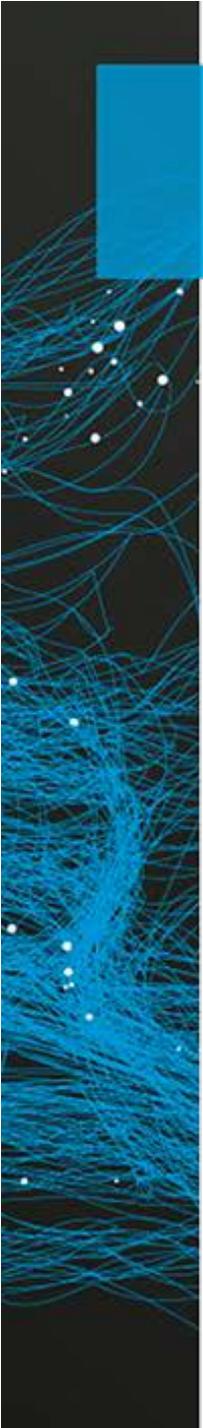


Thematic pillars

6. AI AND SOCIAL GOOD

AI offers great potential for promoting the public good, for example in the realms of education, housing, public health, and sustainability. We see great value in collaborating with public and private organizations, including academia, scientific societies, NGOs, social entrepreneurs, and interested private citizens to promote discussions and catalyze efforts to address society's most pressing challenges.

Some of these projects may address deep societal challenges and will be moonshots – ambitious big bets that could have far-reaching impacts. Others may be creative ideas that could quickly produce positive results by harnessing AI advances.



Thematic pillars

7. SPECIAL INITIATIVES

Beyond the specified thematic pillars, we also seek to convene and support projects that resonate with the tenets of our organization.

We are particularly interested in supporting people and organizations that can benefit from the Partnership's diverse range of stakeholders.

We are open-minded about the forms that these efforts will take.

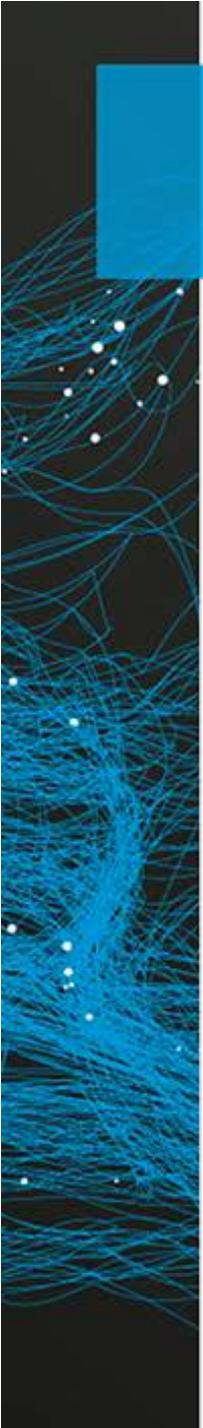
Partnership on AI

STRUCTURE

A Board of Directors will drive the activities of the Partnership. The board is comprised of six seats for representatives from each of the founding companies, and an equal number of seats for non-company representatives.

An Executive Steering Committee, composed of working group leaders and others invited by the Board of Directors, will guide, commission, and evaluate activities within the overall objectives and scope set by the Board.

Day-to-day operations is overseen by an executive director (Terah Lyons) who works closely with the Board of Directors. Conferences, meetings, panels, projects, and working groups will be commissioned by the board and conducted by the executive staff.



Partnership on AI

MEMBERS OF THE BOARD

Each founding company will have a representative on the board.

The current representatives are Ralf Herbrich from Amazon, Tom Gruber from Apple, Greg Corrado from DeepMind/Google, Yann LeCun from Facebook, Francesca Rossi from IBM, and Eric Horvitz from Microsoft.

In addition to these founding members in January 2017 we welcomed six new Independent Directors: Dario Amodei (OpenAI), Subbarao Kambhampati (AAAI/ASU), Deirdre Mulligan (UC Berkeley), Eric Sears (MacArthur Foundation), Carol Rose (ACLU) and Jason Furman (HKS).

Partnership on AI

PARTNERS*



(*) This list is not complete



Partnership on AI

HOW CAN I GET INVOLVED?

We are very excited to work with anyone who is interested in joining our effort.

This is a collaborative and multi-stakeholder organization and we want people with an interest in AI from across all fields to take part.

If you want to take part in some way, please email us at getintouch@partnershiponai.org.

The background of the slide is a dark blue/black space filled with a complex, abstract pattern of thin, light blue lines and small white dots. The lines are mostly horizontal and wavy, creating a sense of movement and data flow. The dots are scattered throughout, some appearing as bright points of light.

EUROPEAN

BIG DATA VALUE FORUM

VERSAILLES, NOV 21-23 2017

THANK YOU