



ANDREW TRASK

Senior Research Scientist at DeepMind & OpenMined leader pioneering privacy-preserving Al

- Expert in privacy-preserving AI, federated learning, secure computation, and governance
- Founder & leader of OpenMined, global open-source community (~12,000+ members)
- Author of "Grokking Deep Learning" and educator making complex AI accessible
- Proven collaborator with finance, healthcare, government and academic sectors
- Visionary speaker who blends deep technical insight with actionable ethics and governance

Andrew leverages AI to mine databases of disparate disciplines to forge new and unexpected collaborations.

Andrew is a Senior Research Scientist at DeepMind studying Privacy and Al and the

Founder and Leader at OpenMined, an open-source community of over 12,000 researchers, engineers, and enthusiasts dedicated to making the concepts and tools necessary for privacy-preserving AI reach mainstream adoption.

Andrew has worked with hedge funds, investment banks, healthcare networks, and government intelligence clients on delivering Al solutions.

Andrew is also a passionate AI teacher, with a passion for making complex ideas easy to learn. He is the author of the book "Grokking Deep Learning", an instructor in Udacity's Deep Learning nano degree program, and the author of the popular deep learning blog "i am trask". He is also a member of the United Nations Privacy Task Force, raising awareness and lowering the barrier to entry for the use of privacy-preserving analytics within the public sector.

"In the 1960s, the U.S. government stepped in to create the ARPANET, democratizing access to the largest supercomputers in the country, because only a limited set of researchers had access to them. Now, the U.S. government steps in to create the NAIRR, democratizing access to the largest AI resources in the country, including the largest AI supercomputers, because only a limited set of researchers have access to them. And upon the promise of the NAIRR, the future of AI will be democratic, American innovation will flourish, and we're honored to be a part of it." — Andrew Trask, Executive Director of OpenMined.

In his talks, Andrew blends deep technical understanding with forward-looking perspectives on governance, ethics, and the social responsibility of AI. His audiences range from engineers and data scientists to policymakers, executives, and innovators seeking to understand how to build intelligent systems that respect privacy and human rights.

TEMAS

Andrew tailors each presentation to the needs of his audience and is not limited to the topics listed below. Please ask us about any subject that interests you:

- Artificial Intelligence
- Privacy, security, and governance in artificial intelligence
- Al and deep learning
- Safe, ethical, and privacy-enabled Al
- Governance, ethics, and the social responsibility of AI

PROGRAMAS

Building safe artificial intelligence

Andrew details the most important new techniques in secure, privacy-preserving, and multiowner governed artificial intelligence. Andrew begins with a sober, up-to-date view of the current state of AI safety, user privacy, and AI governance before introducing some of the fundamental tools of technical AI safety: homomorphic encryption, secure multiparty computation, federated learning, and differential privacy. He concludes with an exciting demo from the OpenMined open source project that illustrates how to train a deep neural network while both the training data and the model are in a safe, encrypted state during the entire process.

Privacy-preserving AI

Andrew touches on ideas such as Differential Privacy, and Secure Multi-Party Computation, and how these ideas come into play.

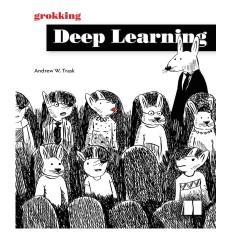
Many industries are limited by regulations of private data, and in response to the need for greater privacy in AI three cutting edge techniques have been developed that have huge potential for the future of machine learning in healthcare: federated learning; differential privacy; encrypted computation. These modern privacy techniques would allow us to train our models on encrypted data from multiple sources without sharing the data.

What is Meaningful Privacy?

Andrew covers how to solve the world's biggest privacy concerns using various state-of-theart privacy-preserving technologies, such as privacy-preserving machine learning.

PUBLICACIONES

Libros



GROKKING DEEP LEARNING

CONDICIONES

Travels from: Oxford, UKFee Range: Please Inquire