

Hyperproductive Machine Learning with Transformers and Hugging Face

Julien Simon

Chief Evangelist, Hugging Face

According to the latest State of AI report, "transformers have emerged as a general-purpose architecture for ML. Not just for Natural Language Processing, but also Speech, Computer Vision or even protein structure prediction." Indeed, the Transformer architecture has proven very efficient on a wide variety of Machine Learning tasks. But how can we keep up with the frantic pace of innovation? Do we really need expert skills to leverage these state-of-the-art models? Or is there a shorter path to creating business value in less time? In this code-level talk, we'll show you how to quickly build and deploy machine learning applications based on state-of-the-art Transformers models. Along the way, you'll learn about the portfolio of open source and commercial Hugging Face solutions, and how they can help you deliver high-quality machine learning solutions faster than ever before.