

Foundation models: Building a future with less human input

The field of machine learning is constantly evolving and the emergence of foundation models is transforming the way we approach the development and deployment of machine learning models.

Temps de lecture : minute

8 June 2023

Unlike traditional AI models, foundation models use massive amounts of unstructured data to learn and generate new information, allowing for more efficient and accurate training of machine learning models.

These foundation models could revolutionise the field of machine learning by lowering the bar to entry as they significantly reduce the need for large labeled datasets that often take hours of human input.

That's in contrast to traditional AI models that typically rely on handcrafted features and explicit programming to perform specific tasks. This requires significant domain expertise from the team working on it and human effort to design and develop. As a result, the performance is often limited by the quality of the features and the complexity of the rules that are programmed into them.

In this article, we will explore the rise of foundation models and their impact on the future of machine learning.

What are foundation models and how do they differ from other types of machine learning models?

Foundation models are a type of machine learning model that are trained on massive amounts of unstructured data, such as text, images, and audio. They are called "foundation" models because they serve as the building blocks for other problem-specific machine learning models, such as language models and image recognition models. Foundation models are typically trained on large-scale datasets, using unsupervised learning techniques that don't require labeled data.

In contrast to traditional machine learning models, which rely heavily on labeled data, foundation models can learn from unstructured data, making them more adaptable to a wide range of tasks. Foundation models also have the ability to generate new information, such as text or images, based on the data they have learned from. This means that they can be used in a variety of applications.

The ability of foundation models to learn from unstructured data and generate new information makes them a powerful tool in the field of machine learning. One of the most widely known uses of foundation models at the moment is OpenAI's GPT-4.. It is used for a wide range of applications, including chatbots, writing assistants, and language translation.

Other examples include OpenAI's DALL-E, a foundation model for image generation and Google's BERT, which is used for natural language processing.

Driving the future with fewer labels

One of the major benefits of foundation models is their ability to reduce the need for labeled data. Traditional machine learning models require large amounts of labeled data to produce accurate results. However, labeled data can be expensive and time-consuming to obtain, particularly for tasks that require a large amount of data, high-expertise domains or for which there are few existing labeled datasets.

For example, in the field of natural language processing, creating labeled datasets for sentiment analysis or named entity recognition can require a significant amount of manual annotation, which can be time-consuming and labor-intensive.

On the other hand, even when labeled datasets are available, they may not be representative of the real-world data that the machine learning model will encounter. This can result in the model being biased or inaccurate when applied to new data.

As a result, there has been increasing interest in developing machine learning models that can learn from unstructured or semi-structured data, such as images, text, or audio, without relying on large amounts of labeled data.

This is where foundation models come in. Once a foundation model has been trained on a large dataset, it can be fine-tuned on a smaller, labeled dataset to produce accurate results for a specific task.

Currently, there is a lot of work being done around fine-tuning. Known as Parameter-Efficient Fine-Tuning (PEFT), this technique uses a combination of knowledge distillation and progressive shrinking of the model to reduce the number of parameters in the fine-tuned model while maintaining its accuracy.

It involves training a smaller student model on the target task, guided by a pre-trained teacher model. The student model learns to mimic the teacher's predictions through knowledge distillation. The student model is then further trained with a smaller learning rate and progressively reduced number of parameters through weight pruning, repeated until the desired level of parameter efficiency is achieved.

PEFT has been shown to significantly reduce the number of parameters in the fine-tuned model without sacrificing accuracy, making it a particularly useful technique for low-resource environments where computational power is limited.

Take the development of a self-driving car, for example. Typically, this would require a traditional machine learning model with large amounts of labeled data to produce accurate results.

Self-driving cars rely on machine learning models to interpret data from sensors such as cameras, LiDARs, and radars, and make decisions about how to control the vehicle. To train these models, large amounts of labeled data are typically required. This labeled data consists of sensor data from a variety of driving scenarios, including different weather conditions, lighting conditions, and road types.

If there wasn't enough labeled data for the machine learning model, the car wouldn't be able to accurately interpret vital data, like sensors, that it needs to make appropriate decisions. This could result in unsafe driving behavior, putting passengers and other drivers at risk.

In this example, foundation models could help in the development of self-driving cars by reducing the amount of labeled data required to train the machine learning models.

For example, a foundation model trained on a large and diverse dataset of images could be used as a starting point for training a self-driving car's

machine learning models. By leveraging the pre-trained foundation model, the machine learning models can learn to recognize visual patterns and features that are relevant to driving scenarios.

This approach could significantly reduce the amount of labeled data required for training the machine learning models, which could save a significant amount of time and money for the developer.

Additionally, PEFT could be applied by using a pre-trained language model such as GPT-4 as a teacher model to guide the training of a smaller student model for a specific self-driving task, such as lane detection.

The ability of foundation models to reduce the need for labeled data has truly significant implications for the future of machine learning. It reduces the barriers to entry and opens up the possibilities for researchers and practitioners to more easily develop and launch machine learning models for a wide range of tasks, including those for which there is limited labeled data available.

These could include tasks like developing sentiment analysis models for specialized domains, such as medical or legal texts, where labeled data is scarce or building recommendation systems for niche products or services that might have a small number of user ratings or reviews.

Roadblocks and ethical considerations

While foundation models have the potential to transform the field of machine learning, there are also a number of challenges and limitations associated with their development and deployment.

One of the main challenges is the computational resources required to train foundation models and storage, which can be prohibitively expensive for smaller organizations or researchers with limited resources.

Another challenge that is more of an ethical one is the potential for bias in foundation models. Foundation models learn from the data they are trained on, which means that any biases in the data is likely to be reflected in the model's output. This can lead to unintended consequences, such as perpetuating existing biases or discrimination. Addressing these biases requires careful consideration of the training data and the development of algorithms that are designed to mitigate the effects of bias.

In fact, the use of foundation models raises a number of ethical considerations, particularly in terms of data privacy. One of the main concerns is the use of personal data to train foundation models. Foundation models are typically trained on massive amounts of data, which can include personal information such as names, addresses, and social security numbers. This raises concerns about data privacy and the potential for misuse of personal data.

Foundation models can also be difficult to interpret. Because the data they learn from is unstructured, it can be difficult to understand how the model arrived at its output. This can be problematic in cases where the output of the model has real-world consequences, such as in medical diagnosis or self-driving cars.

Of course, as with any powerful technology, there is also the risk that foundation models could be used for malicious purposes, such as creating deepfakes or spreading disinformation.

Addressing these ethical considerations requires careful consideration of the data used to train foundation models, as well as the development of algorithms that are designed to mitigate the effects of bias. It's also really important to ensure that foundation models are developed and deployed in a way that respects data privacy and security.

As the use of foundation models continues to grow, it is important for organizations to prioritize overcoming potential roadblocks and take ethical considerations and responsible development and deployment into account from the outset.

The future of the industry

The use of foundation models is already widespread in industry applications, particularly in the fields of natural language processing and computer vision. However, we can expect to see its uses continue to grow rapidly in the near future.

Outside of autonomous systems, like self-driving cars and drones, another area where foundation models are likely to have a significant impact is in the field of healthcare.

Foundation models can be used to assist surgeons in performing complex surgical procedures. For instance, when they are trained on large amounts of medical data, including 3D imaging and patient medical records, they can be used to create a detailed model of the patient's anatomy. During the surgery, this model can be overlaid onto the patient's actual anatomy, allowing the surgeon to navigate through complex structures with greater precision and accuracy.

Furthermore, foundation models can also be used in resource tracking. For example, in a hospital setting, foundation models can analyze data such as patient flow, bed availability, and staff scheduling to optimize resource allocation and improve patient outcomes. The model can identify trends and patterns in the data that may not be immediately apparent to human analysts, allowing hospital administrators to make more informed decisions about resource allocation and patient care.

As the use of foundation models becomes more widespread, we can also

expect to see increased emphasis on responsible development and deployment. This includes addressing issues related to bias and fairness, as well as ensuring that foundation models are developed and deployed in a way that respects data privacy and security.

Encord's approach to foundation models

Encord is a leading provider of machine learning solutions, and our approach to the development and implementation of foundation models is focused on responsible innovation. Our mission is to enable every company to harness the power of artificial intelligence (AI).

Our team is dedicated to making AI practical by developing applications that facilitate the creation of active learning pipelines. This includes streamlining processes such as model training, diagnosis, validation, data annotation, management, and evaluation.

When it comes to how businesses can explore using foundation models in their own machine learning projects, these are our top tips:

Start with diverse and representative data: To mitigate the risk of bias, it is important to use diverse and representative data when training foundation models.

Address privacy and security concerns: Ensure that data privacy and security are a top priority throughout the development and deployment process by implementing data anonymization and encryption techniques, establishing clear policies for data access and sharing, and integrating privacy and security considerations throughout the model development lifecycle.

Use transparency and interpretability techniques: Foundation models can be difficult to interpret, so use visualization tools and explanation

methods to help users understand how the model arrived at its output. Transparency and interpretability techniques involve making the decision-making process of the foundation model more visible and understandable to users.

This can include the use of visualization tools, such as heatmaps, that highlight areas of the input data that were most influential in the model's output. It can also involve the use of explanation methods, such as generating textual or graphical explanations, to provide users with a better understanding of how the model arrived at its conclusions.

Monitor and evaluate performance: Continuously monitor and evaluate the performance of foundation models, and make adjustments as necessary to ensure that they are performing as expected.

Prioritize responsible innovation: Make ethical considerations a top priority throughout the development and deployment process, and strive to develop and deploy foundation models in a responsible and ethical manner.

The rise of foundation models represents an exciting development in the field of machine learning. By learning from unstructured data and generating new information, foundation models can reduce the reliance on labeled data, improve accuracy, and enable the development of more diverse and effective machine learning models.

However, as with any emerging technology, it is essential that ethical considerations are addressed to ensure the responsible development and deployment of foundation models.

With ongoing research and development, we can expect foundation models to continue to drive innovation and enable new applications across a wide range of industries and domains, ultimately leading to a more intelligent and efficient world.

Ulrik Stig Hansen is cofounder & President at Encord.

Article by Ulrik Stig Hansen