eBook

Deep Learning with MATLAB

What is Deep Learning?

Deep learning is a subset of machine learning, where artificial neural networks—algorithms modeled to execute categorization tasks. A deep learning model identifies objects or scenes in images, then processes and classifies extracted information to make decisions as part of a larger system. The term 'deep' refers to the number of layers in the network; the more layers, the deeper the network. Traditional neural networks contain two or three layers, while deep networks contain hundreds of layers.

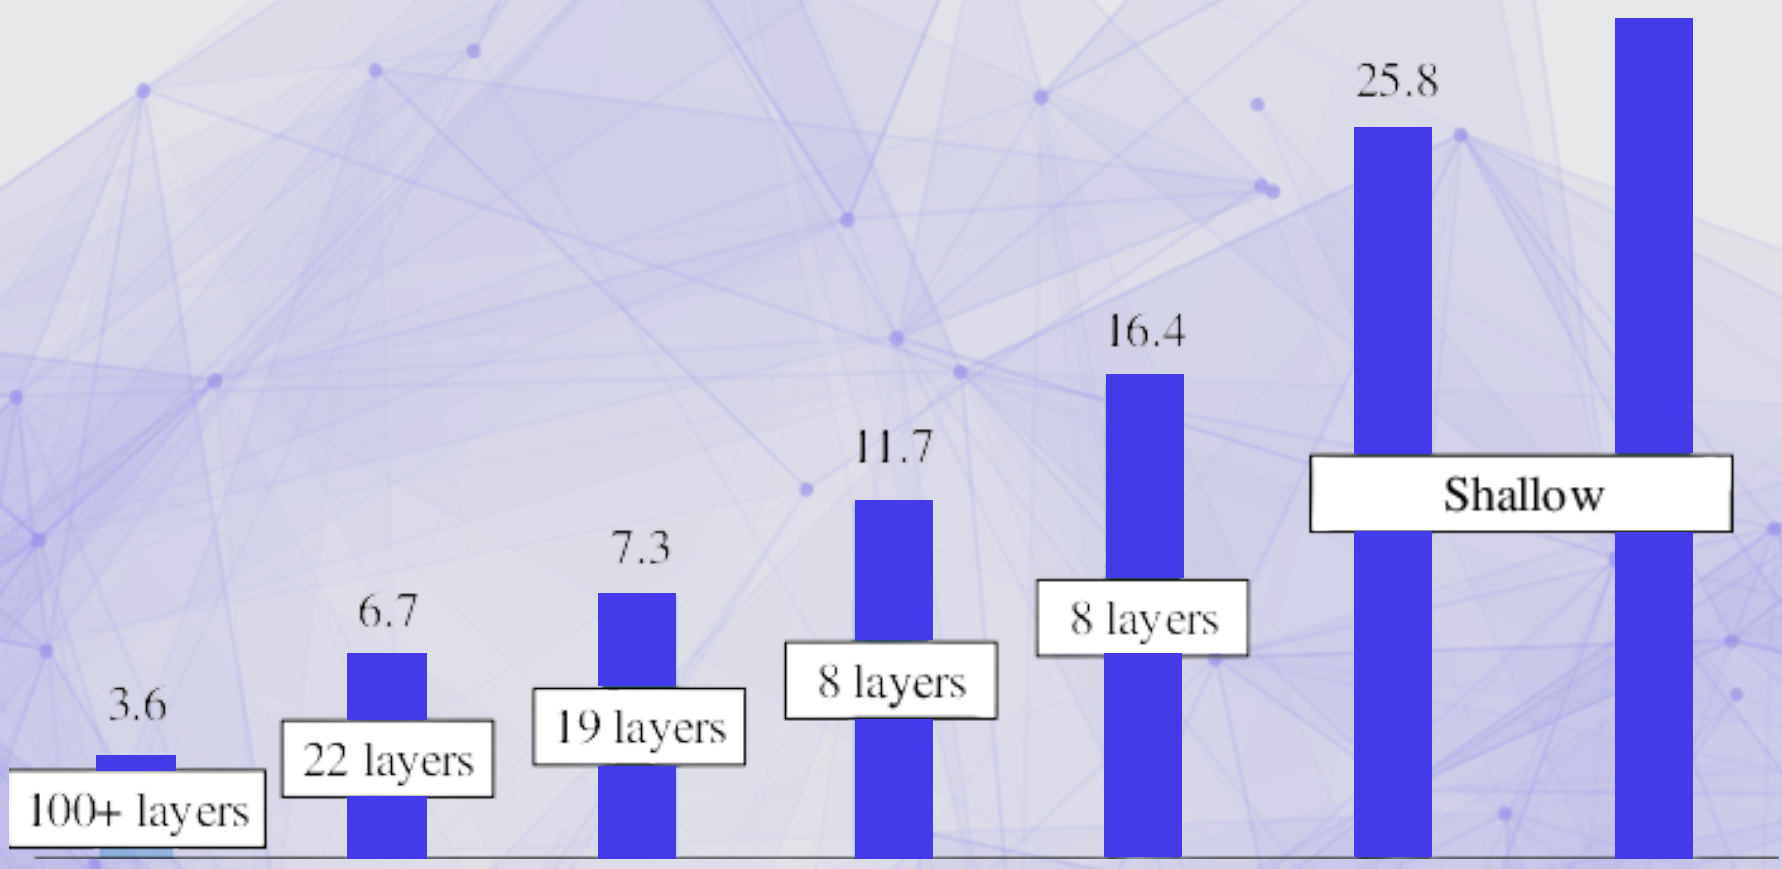

Deep Learning Applications

We would never have dreamed that deep learning applications would one day lead to the development of self-driving cars and virtual assistants such as Alexa, Siri, and Google Assistant. But that is precisely what is happening, and now we live in a world where these innovations are part of our everyday life.

While deep learning continues to pique our curiosity with its seemingly limitless applications, it is particularly useful for identification applications such as facial recognition, speech recognition, driving assistance systems, and so on.

What Makes Deep Learning State-of-the-Art?

While deep learning was first theorized in the 1980s, there are two major reasons it has gained momentum only recently:

- Deep learning requires big amounts of labeled data.
- Deep learning requires substantial computing power like high-performance GPUs with a parallel architecture combined with clusters or cloud computing.

When trained with a large amount of data, deep learning easily outperform humans when it comes to accuracy. Three technology enablers make this degree of accuracy possible:

Easy access to massive sets of labeled data

Data sets such as ImageNet and PASCAL VoC are freely accessible and valuable for training on numerous object kinds.

Increased computing power

High-performance GPUs expedite the training of the enormous volumes of data required for deep learning, decreasing training time from weeks to hours.

Pretrained models built by experts

Transfer learning allows models such us AlexNet to be retrained to solve new recognition tasks.

Deep Neural Network

A deep neural network consists of an input layer, several hidden layers, and an output layer. It is inspired by biological nervous systems; the layers of a deep neural network are interconnected via nodes, or neurons, with each hidden layer using the output of the previous layer as its input. These networks transform data – like the pixels in an image or the words in a document – until they can classify it as an output, such as naming an object in an image or tagging unstructured text data.

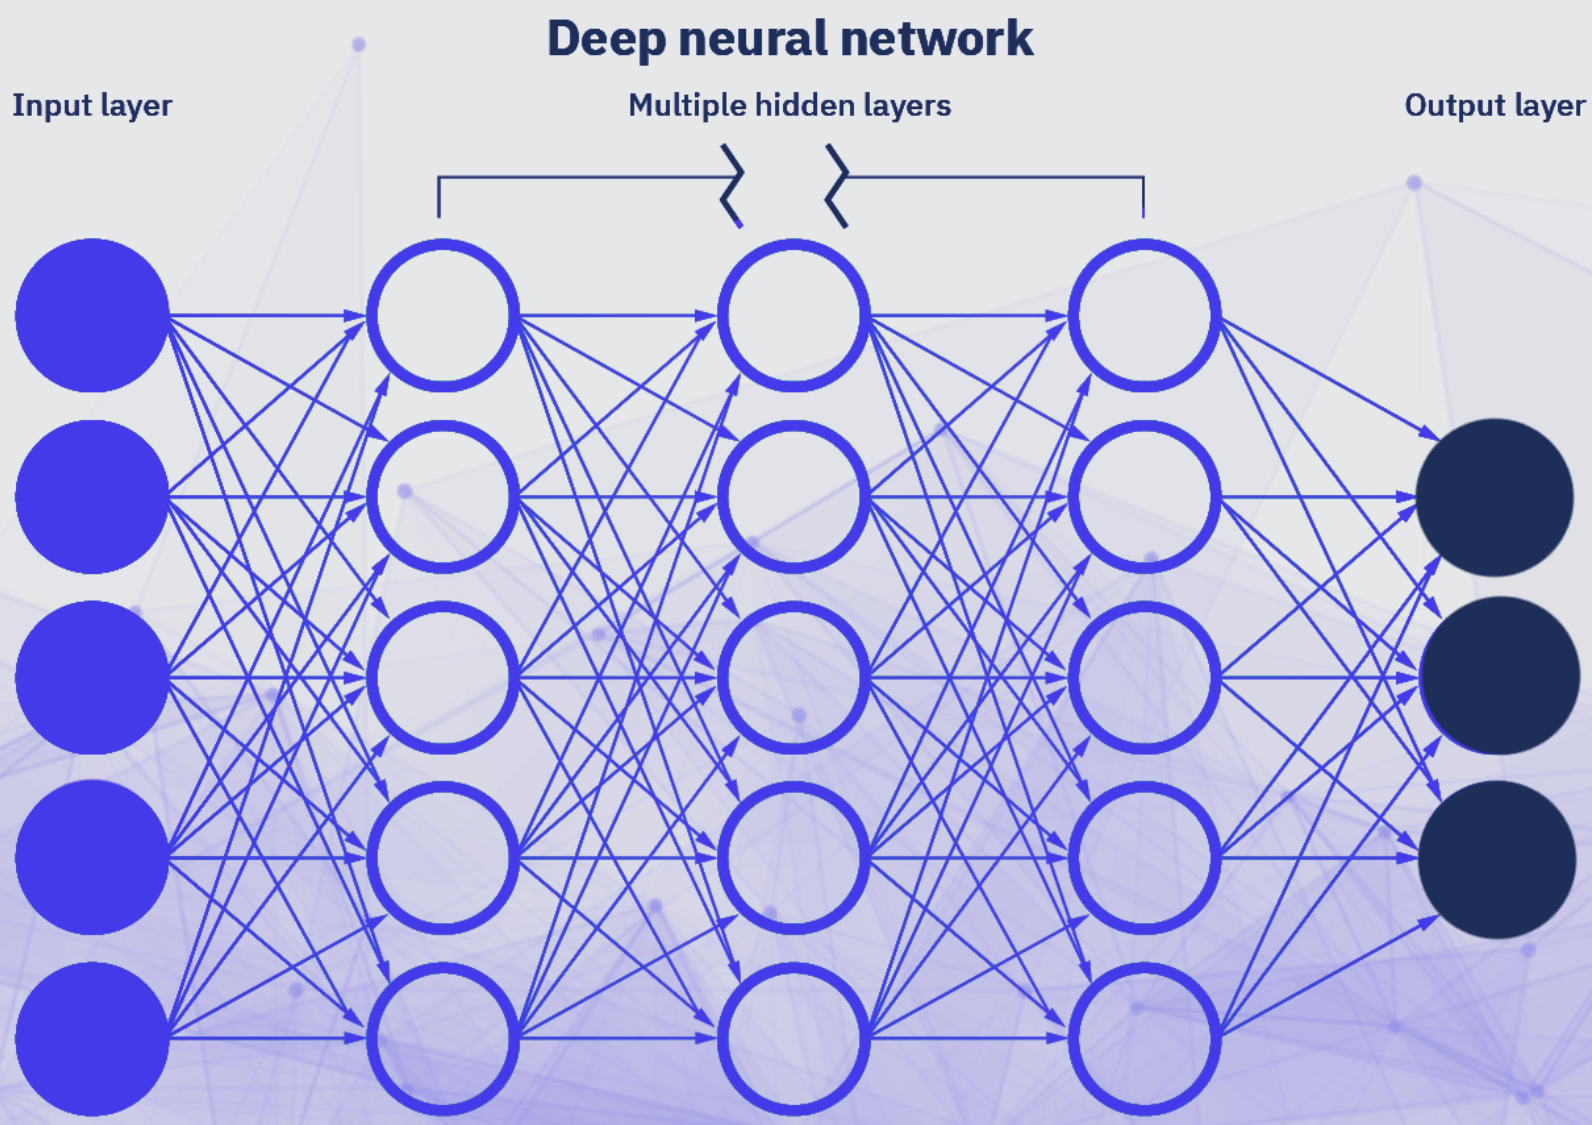

Deep Learning vs. Machine Learning

Deep learning is a subtype of machine learning. Both deep learning and machine learning offer ways to train models and classify data. When using machine learning, you need to manually extract data from an image to train the machine learning model. When using deep learning, you skip the manual step and feed images directly into a deep neural network that learns the features automatically. Deep learning requires access to vast amounts of data and a high-performance GPU, so the model spends less time analyzing those images.

Machine Learning	Deep Learning
+ Good results with small	+ Good results with small
data sets	data sets
– Requires very large data sets	– Requires very large data sets
+ Quick to train a model – Computationally intensive	+ Quick to train a model – Computationally intensive

Deep Learning vs. Reinforcement Learning

Reinforcement Learning is the learning of mapping from situation to action to maximize a scalar reward or reinforcement signal. In contrast to Machine Learning, the learner is not told which action to take but must instead discover which actions yield the highest reward by trying them. In the most exciting cases, actions may affect not only the immediate reward, but also the next situation, and through that, all subsequent rewards. Trial-and-error search and delayed reward – are the two most significant distinguishing properties of Reinforcement Learning.

Reinforcement learning and deep learning are not mutually exclusive. Complex reinforcement learning problems often rely on deep neural networks, a field known as deep reinforcement learning.

Why Use MATLAB for Deep Learning?

MATLAB makes it easy to incorporate deep learning into various applications, whether it's labeling data or designing algorithms.

Simulate Data

Test deep learning models by including them into system-level Simulink simulations. Test edge-case scenarios that are difficult to test on hardware. Understand how your deep learning models impact the performance of the overall system.

Preprocess Data

Use interactive apps to label, crop, and identify important features, and built-in algorithms to help automate the process of labeling.

Train and Evaluate Models

Start with a complete set of algorithms and prebuilt models, then create and modify deep learning models using the Deep Network Designer app.

Deploy Trained Networks

Deploy your trained model on embedded systems, enterprise systems, FPGA devices, or the cloud. Generate code from Intel®, NVIDIA®, and ARM® libraries to create deployable models with high-performance inference speed.

sciengineer