

Data Science White Paper

How Deep Learning Facilitates Automation & Innovation and When to Use It

In our whitepaper, Mosaic explains when deep learning or machine learning algorithms are right for the problem using practical project examples.



Machine Learning

Mosaic data scientists collaborate with customers, digging deep into the data to inform design and deployment of custom ML tools that make a difference.



Artificial Intelligence

Mosaic integrates powerful AI tools into clients' existing technology stack to solve complex business challenges



Business Analytics

Mosaic helps corporations of all shapes and sizes take advantage of their data, transforming their decision-making processes.

DEFINING DEEP LEARNING

In today's technological age, deep learning is at the forefront of innovation in the workplace. Studies indicate that [48% of businesses](#)¹ use deep learning effectively. [Another survey from Modzy](#)² found that out of the 821 technology decision-makers surveyed from companies with more than 1,000 employees, 80% use deep learning AI techniques.

To understand the appeal of deep learning, one must first wrap their head around the concept of machine learning. Machine learning is a broader term that involves the use of algorithms to make projections learned from data using supervised, unsupervised, or reinforcement schemes. Neural networks are a subset of machine learning algorithms that artificially mimic the structure of the human brain. Finally, deep learning, often called “artificial intelligence,” extends the concept of neural networks from simple structures to complex networks with many layers.

Contrary to classic, rule-based approaches, deep learning algorithms develop their behavior by digesting large volumes of annotated examples, a process called “training.” Because of this in-depth learning ability, organizations that rely on deep learning models can achieve state-of-the-art accuracy, often exceeding human-level performance.


For example, [Google's Deep Learning ML](#)³ machine learning engine is 89% accurate in detecting breast cancer, a higher efficiency compared to pathologists (74%). Deep learning techniques account for about 40% of the annual value potentially created by all analytics techniques.

TYPES OF DEEP LEARNING MODELS

Most deep learning methods use neural network architectures, so deep learning models are often referred to as deep neural networks. The term “deep” refers to the number of hidden layers in the neural network. A popular example of this is an Artificial Neural Network (ANN), a computational network inspired by the biological neural network of the human brain.

Deep learning models can remove the human process of feature engineering, algorithm selection, and validation. These deep learning models are trained by using large sets of labeled data and neural network architectures that learn features directly from the data without manual feature extraction. Popular algorithms include CNNs, LSTMs, RNNs, GANs, RBFNs, MLPs, SOMs, DBNs, RBMs, Autoencoders, and Transformers.

Accurate automation can save any organization on operational costs. This flavor of machine learning promises automation not previously available in traditional machine learning, but there is still the expertise needed of when, how, and what model type to use based on explainability, performance, input data, and outputs desired.



Sample architecture diagram of a powerful NLP technique; transformer models.⁵ Input embeddings are passed to an attention layer which are then passed to a feedforward layer. The output of the encoder is passed to the decoder, which also includes an encoder-decoder attention layer.




TRADITIONAL MACHINE LEARNING VS. DEEP LEARNING

When it comes to machine learning, maturity matters. Instead of asking what happened and why, ML serves to understand what will happen and determine optimized actions. The level of complexity associated with deep learning steers us toward specific use cases that cannot otherwise be accomplished with less involved machine learning techniques. By the same token, there are certainly instances where simpler, albeit less accurate models are preferred, making traditional machine learning algorithms the right path forward.

Typically, traditional machine learning algorithms are used when the decisions being modeled are not overly

complex or need substantial subject matter expertise. In other words, explainability is weighted over performance. Organizations can benefit from machine learning when updated predictions are promptly required and when algorithm training time is constrained. In addition, if your organization is working with smaller data sizes, machine learning can typically provide you with the output you need.

With deep learning techniques, organizations that are modeling decisions that are highly complex can benefit if there is no need to sacrifice performance for explainability in modeling decisions. Additionally, deep learning can prove to be useful if algorithm training time is not constrained. One should note, however that deep learning models typically require more training data than classic machine learning algorithms, unless pre-trained models can be leveraged.



POPULAR DEEP LEARNING APPLICATIONS

There are several cross-industry uses for deep learning algorithms. This technique can save your organization time and resources while helping you stand out from the competition with more accurate and efficient decision-making processes.

Computer Vision

[Computer vision](#)⁴ is a field of artificial intelligence that trains a computer to see, identify and process images or video more efficiently than human vision. Deep learning is typically the engine computer vision runs on. The ability to interpret raw photos and videos has been applied to projects in retail, manufacturing, transportation, and energy, to name a few. If you haven't embraced the power of classifying images, your competition probably is. Deep learning models are used in applications such as:

- Facial recognition
- Automated surveillance
- Driverless vehicle transportation
- Reverse image searches
- Surgical operating quality assurance
- Automated inspections
- Quality management

Natural Language Processing

If you were to ask five different people what Natural Language Processing (NLP) is, you would likely get five very different answers. The rise of personal assistants and chatbots has helped spread this technology into our everyday lives. Still, most businesses are barely scratching the surface of what is possible with these algorithms.

There are plenty of use cases for NLP and deep learning, including:

- Document Classification
- Contextual Search Engine Development
- Natural Language Generation
- Quantifying Customer Experience
- Speech Recognition

Structured outputs of NLP models can even be used as inputs to downstream predictive machine learning models. Mosaic can help you find the needle in the unstructured haystack, pulling insights from text, voice, audio, image, and speech to inform operational & strategic decision making across any business unit.

Internet of Things (IoT)

The massive number of IoT sensors already deployed in common settings produce gigantic amounts of data for facilitating a wide range of applications. Deep learning would undoubtedly play a role in generating valuable inferences from this massive volume of data and hence will assist in creating smarter IoT.



The list of use cases for deep learning in the IoT space is extensive, proving that there is ample opportunity for organizations to better optimize IoT data for efficiency using deep learning.


- Smart Homes (behavior monitoring, human activity recognition, posture detection, etc.)
- Smart Healthcare (disease risk prediction and detection, health monitoring, missing clinical data prediction, mortality prediction, etc.)
- Smart Agriculture (plant classification, plant disease detection, crop yield prediction, fruit counting, weed detection, etc.)
- Smart Transportation (traffic congestion management, traffic speed prediction, obstacle detection, accident detection, car park occupancy prediction, etc.)
- Real-time Predictive Maintenance
- Real-time GIS Intelligence

Extracting actionable insights from raw IoT data is an arduous task that goes beyond the proficiency of traditional data analytical paradigms. Deep learning presents a perfect solution for various classification and prediction tasks in the IoT as it can learn hierarchical representations from the input data.

CASE STUDIES: CUSTOMER CHURN WITH CLASSIC ML VS DEEP LEARNING


In this project, Mosaic built a churn model based on historical data to provide a churn score for a retail energy company. After an extensive exploratory data analysis phase, it was decided that classic ML algorithms were best suited for the job. Mosaic's data science team underwent an exhaustive feature engineering phase to standardize the different paths a customer might take in their customer journey.

Classic ML Customer Churn Prediction




Our data scientists trained three algorithms: decision tree, random forest, and logistic regression to evaluate which model met the acceptance criteria. The logit model performed best in part because it enables more fine-grained predictions than tree-based algorithms, in other words there is much less chance two customers have the same score. The logit model was able to predict churn weeks out at 80% accuracy.


This chart shows the churn score vs. customer spend; as spend increases, there is a higher churn rate. Customers in pink are on auto-delivery and are generally less likely to leave than those on will-call delivery.



Deep Learning Customer Churn Prediction



In a different project, Mosaic was tasked with integrating customer call center sentiment into churn predictions. The first step was to automate insight extraction from unstructured text call logs in a scalable, repeatable way. BERT proved incredibly useful in extracting deeper layers of sentiment. Mosaic then used topic modeling to identify common topics present in transcripts as early indicators of churn. To automate the churn prediction, Mosaic blended the outputs from the topic models and the BERT embeddings into a CNN architecture to identify a customer likely to churn. The deep learning workflow was able to identify 23% more cases of churn than the classic ML approach.



CONCLUSION

Deep learning is a critical technology quickly becoming adopted by an array of businesses across multiple industries. No matter the industry, chances are your organization can benefit from the applications and use cases of machine learning within some aspect of your business. The benefits are invaluable, including increased efficiency, lower costs, and happier customers.

Those looking to reap the full benefits of AI can only do so when the appropriate competency levels are developed in advanced data technologies such as deep learning for extracting reliable business insights. For this, it is best to reach out to an expert in the field: a data scientist.

Mosaic has the deep machine learning expertise necessary to take advantage of these powerful algorithms quickly. We can help you develop and implement ML or DL models that aim to solve problems within your organization and contribute to overall business growth. Our unique blend of data engineering, machine/deep learning, business acumen, and custom deployment experience lends itself towards a powerful consultative approach to building solutions that drive powerful results for our clients.

Endnotes

1. <https://explodingtopics.com/blog/ai-statistics>
2. <https://techrepublic.com/article/companies-are-using-ai-to-hit-business-goals-even-though-they-cant-explain-how-it-works/>
3. <https://research.aimultiple.com/ml-stats/#:~:text=Google%E2%80%99s%20Deep%20Learning%20machine%20learning%20program%20is%20accurate,compare%20to%20pathologists%20at%2074%25.%20%28Google%20AI%20Blog%29>
4. <https://mosaicdatascience.com/2019/11/19/machine-vision-introduction/>
5. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In Advances in Neural Information Processing Systems, pages 6000–6010.

