Machine Learning Applications in Today's Business Environments

By Samuel (ShengEr) Zhou

Author
Samuel (ShengEr) Zhou

Machine learning (ML), as a branch of artificial intelligence (AI), is a topic of increasing focus given its prevalence and potential integration into wider AI discussions. This essay seeks to explain what ML is and how we can apply ML models to create value for businesses.

What Is Machine Learning?

ML is basically a branch of AI that is capable of building a system or a model that can make decisions based on past data. ML algorithms use patterns in past data to make predictions and help businesses improve their performances. In order to deploy the model, the data has to be broken into a training portion and a testing portion. The training portion of the data is used to train the ML model so that it is able to make predictions, while a validation portion is used to tune the hyper parameters so that the model is capable of making better predictions. The test data is used to evaluate the model's performance before it can be deployed to make predictions on new data sets. The general flow of the ML project involves data collection and data preprocessing, feature engineering, model selection, training, hyper parameters tuning, evaluation, deployment and monitoring *(Chart 1)*.

CHART 1
Demonstration of overall ML steps, by training, testing & deploying a model

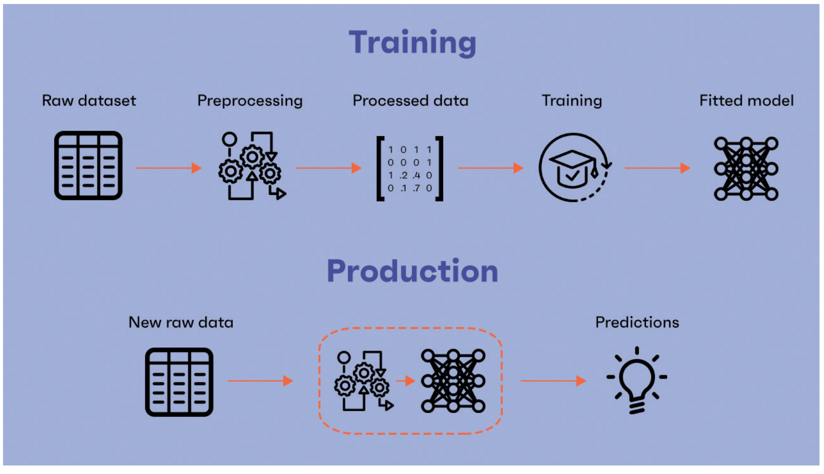

Source: https://www.clearbox.ai/blog/2021-02-10-automating-data-preparation-and-preprocessing-in-production-ready-ml-models

How Can ML Help Businesses?

Businesses create and capture value by two means – increasing revenue and reducing costs – through enhancing customer experience, through personalization, and reducing waste in efficiency. ML can help achieve this. ML models such as K-Nearest Neighbors (KNN) and the Support Vector Machine (SVD) can help businesses create a recommendation model that can better predict customers' needs and wants by learning through their past purchases. Application number one below, through music recommendations, will demonstrate such a use-case.

In terms of reducing inefficiencies and waste, it is notable that ML is widely used in the banking industry where default risks can be managed through creating a fraud label and training a model with the highest accuracy, precision and recall to deploy it. ML models can also be used in fraud detection in that industry, through models trained on past data, and they can be deployed to new data sets to predict the likelihood of fraud in advance, potentially saving businesses tremendous extra expenses or capital deployed to manually detect fraud or predict default risks. Other applications in the business arena include supply chain management where predictive analytics methods could be employed to better forecast demand and optimize inventory, thereby reducing costs.

It should also be noted that in order for ML models to be deployed and make predictions on future data, high quality data is needed, and therefore high quality data collection methods should be leveraged. Other complementary items would include the appropriate infrastructure as well as methodologies for ethical data collection practices without infringement of individual privacy over how the data should be used.

I will introduce three different use-cases, though there are thousands of different ones, and hope this article inspires readers to explore how ML can be integrated into their own businesses to create value for them and their customers.

ML in Music Recommendations

Streaming businesses are in high demand all over the world, and one of the ways in which ML can be deployed is through personalization in music recommendations. Prominently, there are two different types of algorithms that can be leveraged to build such an ML model – KNN and SVD. Item-based KNN basically finds the most similar items to those that a certain user has already interacted with, therefore predicting future purchases. This approach allows users to interact with items that are similar to ones they have already been introduced to, and can be a better predictor than user-based recommendations, whereby a purchase item is recommended to the customer based on the likes of other customers. This approach is also more data friendly, as instead of collecting data from other users, KNN will be able to operate on the one user and other music pieces alone, potentially making this way more scalable and interpretable. In the world of music discovery, KNN is therefore one of the more reliable algorithms for a user to fully engage in other music that an individual would like to listen to.

On the other hand, SVD, a matrix factorization method, essentially translates such matrices into lower-dimensional matrices, uncovering the latent factors within the data. The general flow of such an algorithm includes data collection, matrix construction, SVD application, prediction and recommendation *(Chart 2)*.

The larger the pool of music that a streaming service offers, the more variety will lead to a better trained model and more accuracy for such a deployment. Personalization is nowadays a critical element of customer service and is a key in driving the willingness to pay for streaming businesses, whether music streaming or video streaming platforms. Again, either user-based or item-based algorithms could be deployed and trained to make an ML model for deployment on new data. Normally, a train/test split would be 80/20, to ensure that enough data is being exposed to ensure the minimization of error terms, which famously gave ML a reputation for being able to learn from the past and correct itself.

For example, a prominent player in the music industry, Spotify, utilizes a "Daily Mix" playlist for the sake of gathering insights into the listening habits of its users. The ML model is being applied in a way that clusters similar songs that a person listens to and extracts the audio features from them, which in essence allows the company to generate a variety of different playlists that cater to the moods and tastes of different genres that person enjoys. It also processes natural language in a way that can track the blogs and articles a person writes linked to their Spotify account and will use those insights to find new releases that align with that person's tastes in music given their documented interest in those blogs.

In contrast, music platform Pandora breaks down music into 450 different attributes such as rhythm, key tonality, and even vocal harmonies. Through this process, it matches users with similar attributes that they liked, thereby increasing the accuracy of their predictions. Even more advanced is YouTube's ML models that uses deep learning models to gather insights from your watch histories and items of search to update you on music that will suit your moods, noting that it uses data from other sources that you like, not from music itself, and from other search information.

YouTube also employs ML algorithms toward using other contextual information such as the

CHART 2

Fundamentals of the proposed approach, including matrix factorization in SVD

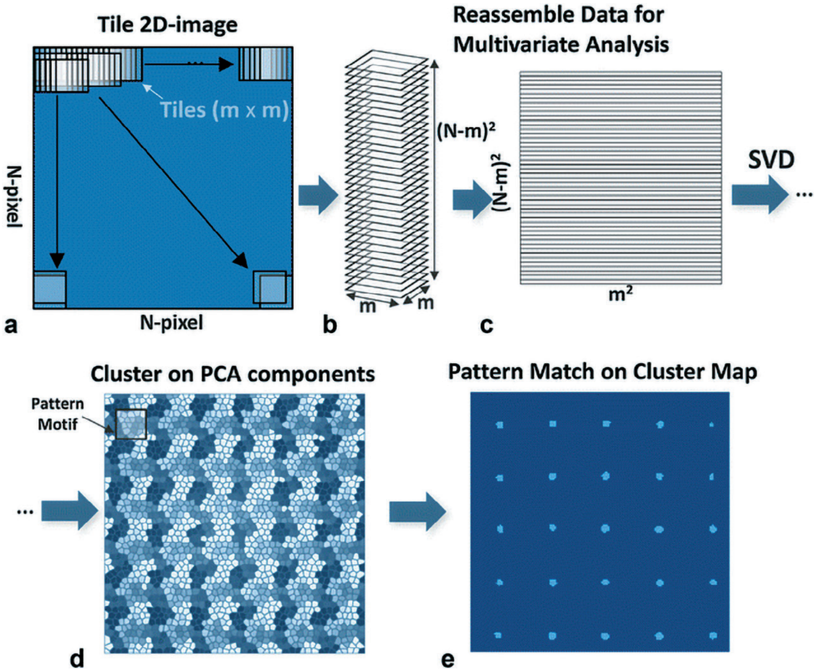

Source: https://www.researchgate.net/figure/Schematic-illustrating-the-fundamentals-of-the-singular-value-decomposition-SVD-based_fig1_323496715

CHART 3

Flowchart of music recommendation use case

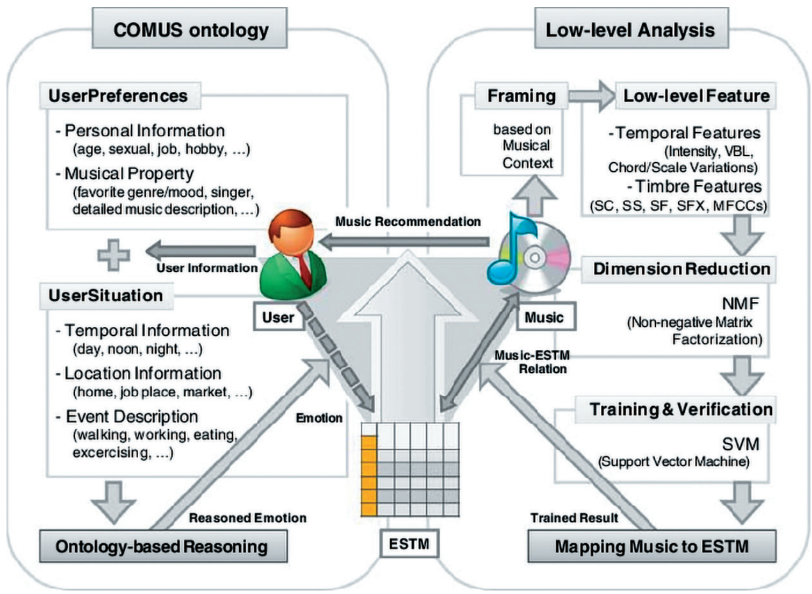

Source: *https://www.researchgate.net/figure/Architecture-of-our-music-recommendation-system_fig2_229031003*

time of day, a person's specific location, and recent activities to recommend particular tracks. It is more probable therefore that YouTube will recommend a more upbeat song in the morning hours and a more laid-back tune during Sunday afternoon, all by using advanced ML algorithms. YouTube also uses feedback, such as user likes and dislikes, to serve as an improving tool to increase its accuracy.

One last ML-driven feature includes voice-activated assistants like that of Amazon's Alexa which uses a sound recognition system and natural language processing to providing recommendations that suit a person's needs on a real-time basis. As Alexa can be integrated into different streaming services such as Spotify, Pandora and Apple Music, a broader range of information and user-based recommendations can enhance such contexts, thereby increasing the personalization. Given the fact that ML algorithms can be combined and used in an integrated fashion, in essence we can combine all of these different features to cater to those who want different music genres at any given time *(Chart 3)*.

ML in Fraud Detections

The use-cases described above show how businesses could use ML to increase their customers' willingness to pay for a service, potentially increasing the revenue side of the profit equation. This use-case seeks to explore how ML algorithms and models can also be deployed to save businesses costs, especially in the banking industry, and how to predict a customer's likelihood of defaulting and fraudulent transactions. It should be noted that this application spans a field much wider than the banking industry, as all industries can leverage it to detect fraud and manage risks.

In order to dig deeper into this topic, we need to understand one of the most important concepts – evaluation metrics. Evaluation metrics consist of true positives, true negatives, false positives, and false negatives. In the context of fraud detections, a true positive is a transaction that correctly identifies itself as fraudulent. A false positive, on the other hand, is when a legitimate transaction is labeled as fraud incorrectly. There is going to be a trade-off here between customer experience and security. As ML algorithms can't have 100% accuracy, precision, and recall, understanding of false

TABLE

Confusion matrix with classification metrics

		Predicted Class: Positive	Predicted Class: Negative	
Actual Class	Positive	True Positive (TP)	False Negative (FN) **Type II Error**	**Sensitivity** $\frac{TP}{(TP+FN)}$
	Negative	False Positive (FP) **Type I Error**	True Negative (TN)	**Specificity** $\frac{TN}{(TN+FP)}$
		Precision $\frac{TP}{(TP+FP)}$	**Negative Predictive Value** $\frac{TN}{(TN+FN)}$	**Accuracy** $\frac{TP+TN}{(TP+TN+FP+FN)}$

Source: *https://medium.com/@_SSP/confusion-matrix-f7ff01c5bbb6* positive tolerance is essential for maintaining a good customer experience so that not too many transactions are falsely identified as fraud. But to invest in a low false positive rate would also entail investing more resources and therefore increasing costs for businesses *(Table)*.

Three different models can be trained for this fraud detection process: logistic regression, random forests, and gradient boosting. By examining metrics such as accuracy, precision and recall, as well as a receiver operating characteristic (ROC) curve, extreme gradient boosting (XGBoost) seems to be the most effective model – though a business should always compare the effectiveness of different models and evaluate the most effective one before its deployment. Throughout this process, business analysts can assess what variables are the most important in predicting fraudulent transactions or default risks, and what features deserve the most attention and the most capital to monitor, while the least important features can be given the least amount of attention.

Another example of ML application in this field is Visa's real-time fraud detection using ML algorithms like that of XGBoost and neural networks, which analyze all transactions in real-time by assessing important features, such as transaction amounts as well as the different types of merchants involved. When such a transaction is made, Visa uses its ML models to compare with the cardholder's historical spending habits or habits of a similar cluster of users to detect any anomalies when shopping. When such a transaction deviates from past habits in a significant way, it is going to be flagged for further review to minimize potential losses for the customer before more illegal activities can be undertaken.

Google also uses ML to detect fraud with models such as logistic regression and deep learning through analysis of clicking patterns, as well as user behavior, to prevent fraudulent clicks associated with activities aimed at driving up adversaries' advertising costs. The ML model is able to identify clicking patterns by analyzing a particularly high frequency in a short timeframe from a single IP address that does not present itself as a meaningful interaction. By protecting its advertisers, Google is able to continue to encourage ad spending on its platform to help drive up its revenue.

By applying all of these different ML models to detect fraudulent activities, businesses are able to safeguard financial security while building trust, which is essential in today's business landscape where such trust is seen as foundational for future exploration in business.

Sentiment Analysis for Natural Language Processing

ML can also be used for natural language processing. It is able to analyze texts and extract sentiments from them and give a score, whether positive or negative. These ML methods include the Bidirectional Encoder Representations from Transformers (BERT) and other ensemble methods. They are great for businesses trying to understand the views of their customers and gather business intelligence, such as by analyzing their tweets and other comments. Even at music festivals, individual songs sung or played can be analyzed in real-time to judge sentiments, thereby improving the marketing efforts of the artists. In the age of big data, this is a very popular method of gathering customer intelligence.

It can also be used by businesses to conduct customer research, such as popular shoe and clothing retailers like Zappos that aims to deliver exceptional customer service by collect user feedback. Zappos uses its trained ML models to scan customer complaint emails, feedback via Chatbots and on social media to assess common problems and identify strategies to address them, which increases customer loyalty and serves as a differentiator from its competitors.

More sentiment analysis involves popular brands such as Nike and Coca-Cola which aim to increase engagement by amplifying positive reviews while minimizing negative reviews by collecting them through a trained ML model.

ML in Business Applications

There are many more ML applications that this specific branch of AI can be applied to so as to create value for businesses. One of the central rules for innovation in business is the incorporation of new technologies and algorithms. Leveraging of data and analytics is no longer just another operational effectiveness booster, but a competitive advantage. By building models for better customer experience, businesses can learn much more about their customers, as well as learning from past data, and their ability to make better predictions can be used to prevent fraud and thereby reduce their costs. JS

Samuel Zhou is the co-founder of Shenger&Ames, a company dedicated to integrating AI into holiday experiences.