intetics

Where software concepts come alive™

ML Models:

Exclusive AI/ML Knowledge for Business and Engineers

White Paper

Table of Content

intetics
Where software concepts come alive™

1. Objective

Businesses keep investing in technology to have a competitive edge and remain resilient. These goals are now being fulfilled through data manipulation, the scope of which is growing daily. As a result, efficient data analysis, processing, and usage lead to enhanced agility, more accurate decision-making, a superior user experience, and optimized costs. To sift through and work with the extraordinary amount of data that is constantly being generated, businesses use machine learning (ML), the benefits of which will be examined in this White Paper.

Through automated ML algorithms, you can quickly and accurately process massive volumes of data, as well as apply the output to address various business challenges. These benefits contribute to ML adoption: according to the 2022 NewVantage Partners Survey, 9 out of 10 businesses continuously invest in ML and AI, and 91.7% of them plan to increase their investments.

Table of Content

intetics
Where software concepts come alive™

2. Main Problems to Be Solved/Needs Addressed

Although ML use cases are becoming more varied, customer-centric applications remain common. Here are some of the top AI/ML use cases worldwide, according to Statista findings:

Improving Customer Experience

By collecting and analyzing customer data, including demographics, preferences, online behavior, etc., businesses can better understand their users and provide more relevant offerings. As a result, businesses can increase long-term customer engagement, loyalty, and retention, often through reduced customer churn.

Generating Customer Insights and Intelligence

ML algorithms can process both structured and unstructured data to identify customer preferences. Companies will gain market insights to foster more accurate decision-making, with little to no guesswork.

Interacting with Customers

Natural language processing (NLP) is a technology at the core of virtual assistants that can provide tailored customer support around the clock. They can be contextual chatbots, keyword recognition-based chatbots, voice-enabled chatbots, and others.

Detecting Fraud

ML models are trained to identify suspicious activity, data anomalies, unaccepted content, and fraud, which allows businesses to address them in time, minimizing financial and reputational risks.

Acquiring New Customers

By analyzing customer segments, businesses can identify their preferences and come up with personalized marketing initiatives, which optimize and speed up customer acquisition.

Table of Content

Building Brand Awareness

NLP models allow for deciphering user sentiments, while other ML models and bots can monitor and respond to reviews, as well as communicate with the audience.

Recommendation Systems

ML models can be trained to analyze user behavior to provide personalized offerings, as well as to come up with a range of associated products and services, which increases sales.

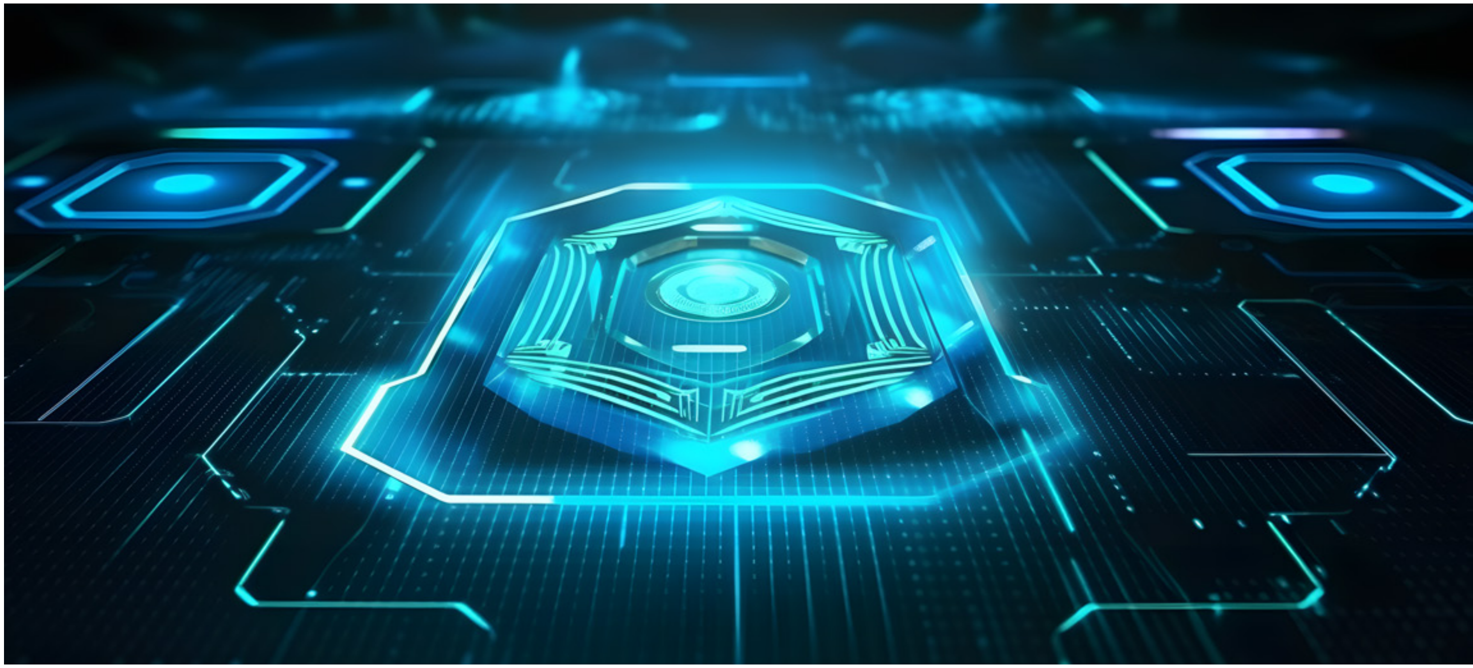

Keep in mind, though, that even with all the tangible and intangible benefits of ML deployment, there are still associated challenges when deploying and using the tech. According to Statista, the major barrier to overcome is IT governance, security, and auditability requirements. Besides this, businesses often struggle with programming languages and framework support, model performance monitoring, regular updates, cost management, and effort duplication across the organization.

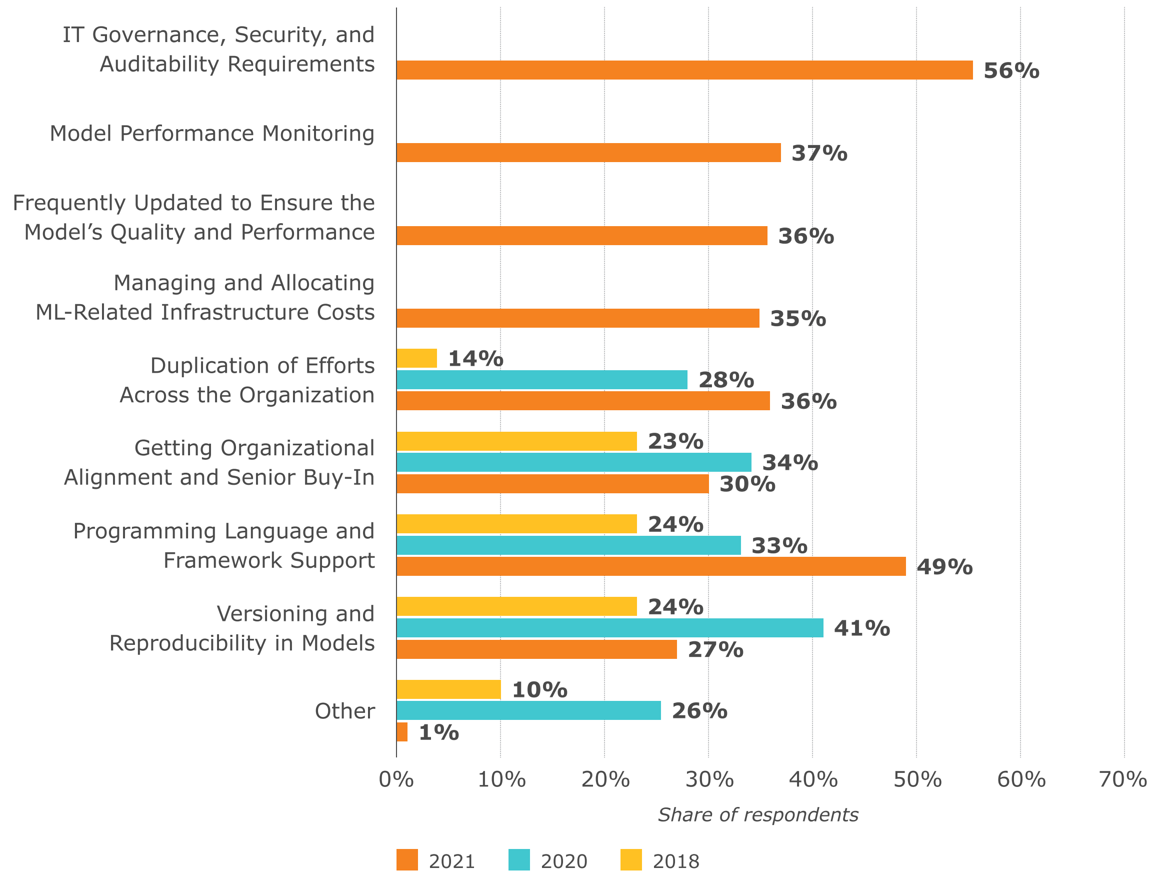

1. IT Governance, Security, and Auditability Requirements

When running ML models, IT often finds it difficult to ensure proper data governance and security, which may pose significant risks to the whole infrastructure. According to the O'Reilly report on enterprise machine learning adoption, ML-mature businesses tend to devote attention to compliance and privacy when building a model.

The security issue can be partially addressed through regular audits. They can be performed in various depths, which often require different technical competence and access policies.

intetics
Where software concepts come alive™

2. Model Performance Monitoring

ML models are dynamic and sensitive to data changes in real time, which leads to their degradation in terms of performance as soon as they are deployed. To ensure the results meet the objective, they need to be constantly monitored and trained to properly process the input data. So, businesses often struggle to evaluate the quality of the constantly generated data and to efficiently assess output quality. In fact, only a quarter of those deploying ML models check for possible bias in models, according to the O'Reilly report.

3. Frequently Updated to Ensure the Model's Quality and Performance

Since new piles of data are being continuously generated, businesses need to always update the model with new data and train upon them, if the data type is different from the initial dataset. Otherwise, the model's performance will deteriorate, and the output may not solve the initial challenge.

4. Managing and Allocating ML-Related Infrastructure Costs

ML models scale up as the number of users and, therefore, data, increases. This often requires businesses to increase data storage, improve processing capabilities, and even use parallel models to adhere to initial quality and performance benchmarks.

5. Duplication of Efforts Across the Organization

Managing ML models often requires combined efforts of different teams—for example, ML and DevOps developers that keep monitoring, updating, and training the models. With poor responsibility allocation, teams often fail to set up efficient collaboration, which reduces productivity.

6. Getting Organizational Alignment and Senior Buy-In

High-level management doesn't always realize the potential benefits and profits of ML, so they might not be willing to invest in the technology or scale it up to experience significant outcomes. One of the potential reasons is that it

Table of Content

may be difficult to showcase how ML algorithms solve vague business objectives.

Such challenges may arise because the management of ML teams differs from traditional software development, where core decisions are typically made on the top management level. According to the O'Reilly report, product managers often determine key metrics and decide team priorities.

7. Programming Language and Framework Support

Often, when building ML models, you may need to change a programming language to foster performance or align the models with the specific tech stack. Besides this, a chosen framework may not be supported any longer, prompting you to switch to another one. This is time-consuming and may undermine efficiency and performance.

8. Versioning and Reproducibility in Models

In order to achieve the desired outcome or improve the model performance, IT specialists often turn to model versioning to later reproduce the successful one. Many, however, find it difficult to maintain proper version control and lose productivity and accuracy because of ad-hoc scripts and processes.

Table of Content

intetics
Where software concepts come alive™

3. Technology Overview

Machine learning (ML) is often seen as a part of Artificial Intelligence (AI) and is a field of inquiry focusing on understanding and building ML methods that are trained and can "learn." This allows for leveraging data to improve performance on some sets of tasks.

We have compiled a list of the most commonly used terms regarding ML, as well as key contributors.

Main Terminology

Supervised Learning

An ML approach in which a program is trained on a predefined dataset with known input and output data. It can then predict future outcomes when new data is provided. For example, a sentiment analysis classifier is trained based on tagged positive, neutral, and negative comments.

Unsupervised Learning

An ML approach that implies determining hidden patterns or intrinsic structures in input data for improved decision-making. For example, a recommendation system identifies patterns and structures to provide relevant product recommendations.

Semi-Supervised Learning

An ML approach that combines both supervised and unsupervised learning.

Classification

A subcategory of supervised learning that involves a process of assigning a label to some input. The predictions are usually of a discrete or "yes and no" nature.

Neural Networks

A method that teaches computers to process data in a way similar to the human brain.

Table of Content

Regression

A subcategory of supervised learning similar to classification. It's applied when the "class" to be predicted is made up of continuous numerical values. Typically, it generates numerical predictions, answering the questions "How much?" and "How many?"

Clustering

A form of unsupervised learning that groups unlabeled examples. The process is based on the concept of maximizing intraclass similarities while minimizing interclass ones. In other words, it groups very similar instances and ungroups instances that are much less similar to one another.

Decision Trees

They are top-down, recursive, divide-and-conquer classifiers, often composed of 2 main tasks: tree induction (assigning attributes to pre-classified instances, splitting the database, and recursing the split datasets until training instances are categorized) and tree pruning (removing excessive structures from a tree to improve efficiency, readability, and accuracy).

Support Vector Machines (SVMs)

Supervised learning models with associated learning algorithms to analyze data for classification and regression analysis.

Natural Language Processing (NLP)

It’s a subfield of computer science, linguistics, and AI that focuses on algorithm processing and analyzing large amounts of natural language data.

Computer Vision

Implies methods for acquiring, processing, analyzing, and understanding digital images to produce symbolic or numerical information in the form of decisions.

Deep Learning

A category of ML algorithms often based on Artificial Neural Networks to generate models with multiple hidden layers in order to solve problems.

Table of Content

intetics

Where software concepts come alive™

Transfer Learning

An ML method in which a pre-trained model is reused for a similar model to facilitate the training process.

Main Players

IBM

IBM is among the key contributors in the spheres of AI and Analytics, whose proprietary software was ranked by the G2 2023 Best Software Awards. The company has developed IBM Watson Studio — a software platform for data science with multiple collaboration opportunities and open-source tools, including RStudio, Spark, and Python.

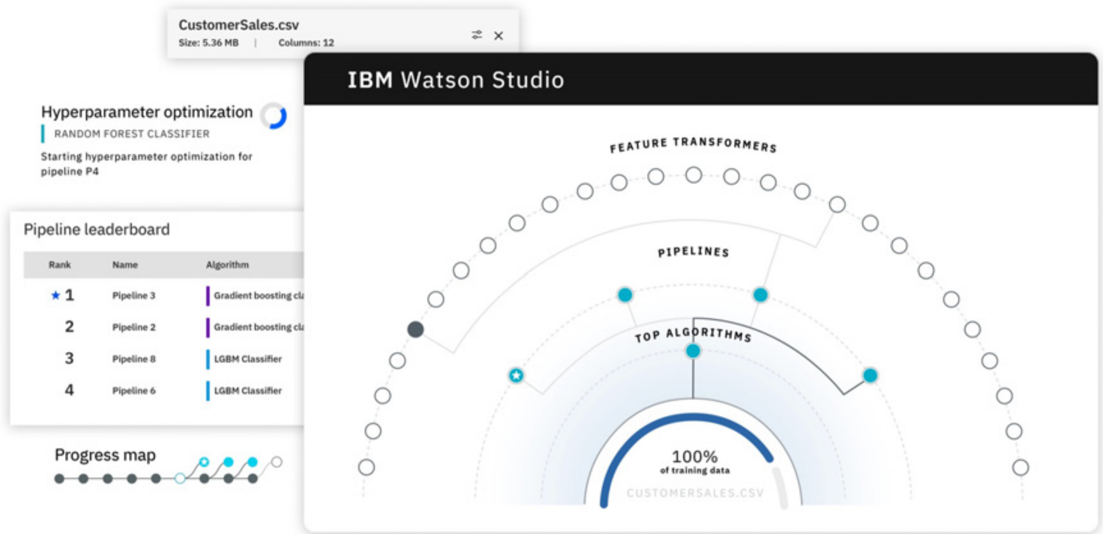

Source: ibm.com

The platform can be deployed both on-premise and in various clouds, allowing data scientists, developers, and analysts to build, run, and manage AI models.

Besides, there is a range of supported ecosystem tools for code-based and visual data science. This allows you to quickly bring AI models to production and scale them up across any cloud, while implementing MLOps and trusted AI algorithms and automating their life cycles.

Besides this, IBM also developed IBM SPSS Modeler — a data mining and text analytics software for building predictive models and carrying out other analytics tasks without programming. It's widely used for customer analytics, fraud detection and prevention, risk management, healthcare quality improvement, demand forecasts, etc.

SAP Data Intelligence Tool combines traditional ETL (extract, transform, and load) tools together with machine learning capabilities that can be applied on top of data. The tool is aimed to connect, discover, enrich, and orchestrate disjoined data assets into actionable business insights.

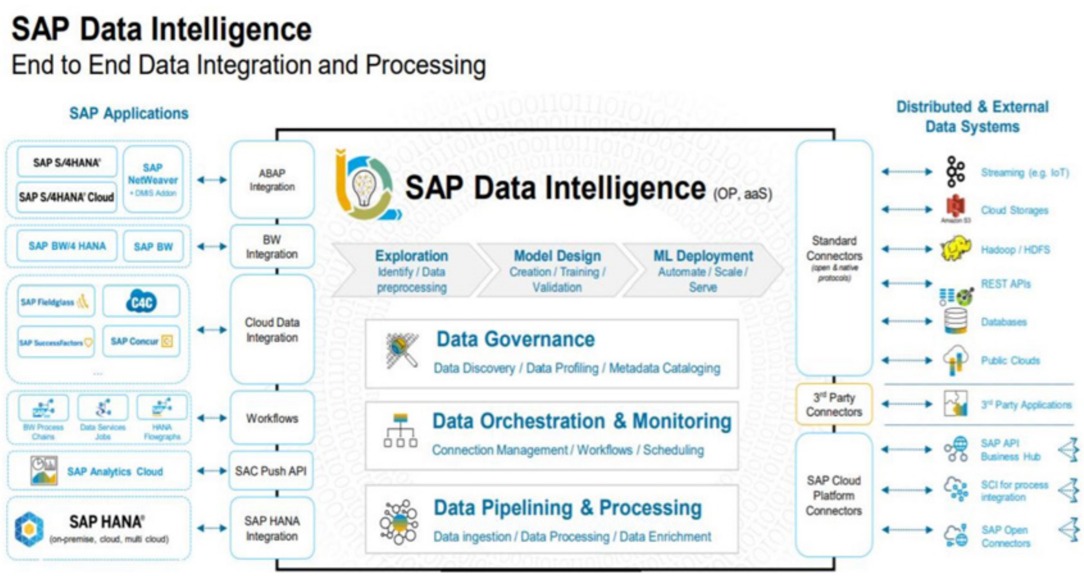

Source: sap.com

Besides this, there is a range of solutions for data and analytics that has been recognized by TrustRadius, an independent peer review platform for B2B software, with multiple Winter 2023 "Best of" awards. For example, SAP HANA Cloud, SAP Data Warehouse Cloud, SAP Master Data Governance, and SAP Data Intelligence Cloud received "Best Feature Set" and "Best Value for Price" awards.

ORACLE

Oracle has embedded data mining functionality in its Oracle Database, facilitating data mining and data analysis through classification, prediction, regression, associations, feature selection, anomaly detection, feature extraction, and specialized analytics. As a result, it allows you to create, manage, and deploy data mining models across the database environment.

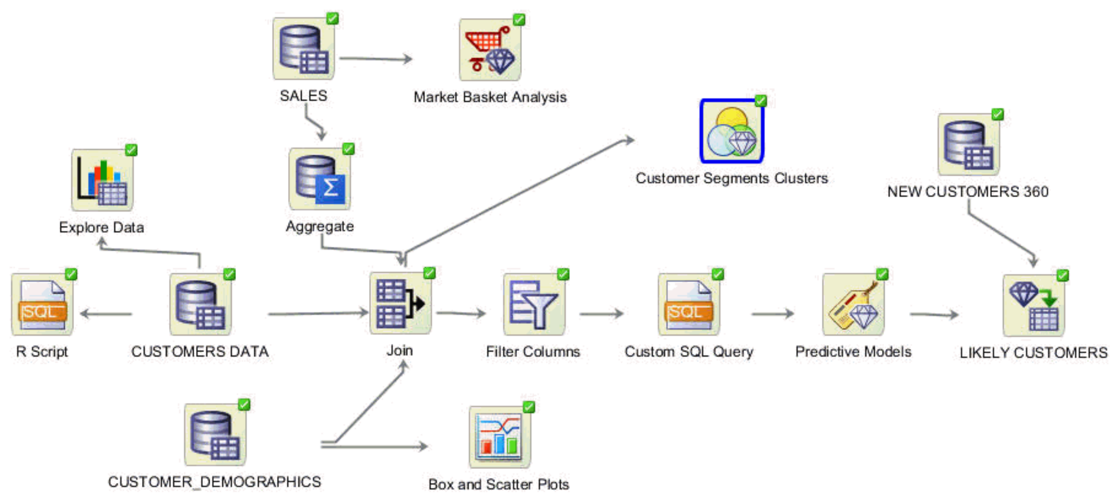

Oracle Data Mining was first introduced in 2002, being a successor to the Darwin data mining toolset produced by Thinking Machines, which was acquired in 1999. The solution was rewritten from scratch though and has passed several evolution stages until it began to offer a wide range of ML capabilities for data mining and analytics.

Microsoft

Microsoft Azure, a cloud computing platform released in 2008, provides an array of data management services and tools, including:

- Azure Data Explorer for Big Data analytics and data exploration
- Azure SQL Database to create, scale, and extend apps into the cloud with the Microsoft SQL server
- Azure Data Factory as a data integration service to create data-driven workflows in the cloud and automate data movement and data transformation
- Azure Synapse Analytics as a fully managed cloud data warehouse
- Azure Stream Analytics as a serverless scalable event processing engine to develop and run real-time analytics on multiple streams

These and other solutions form the basis of the Azure Machine Learning service to quickly deploy and manage ML models with governance, security, and compliance in mind.

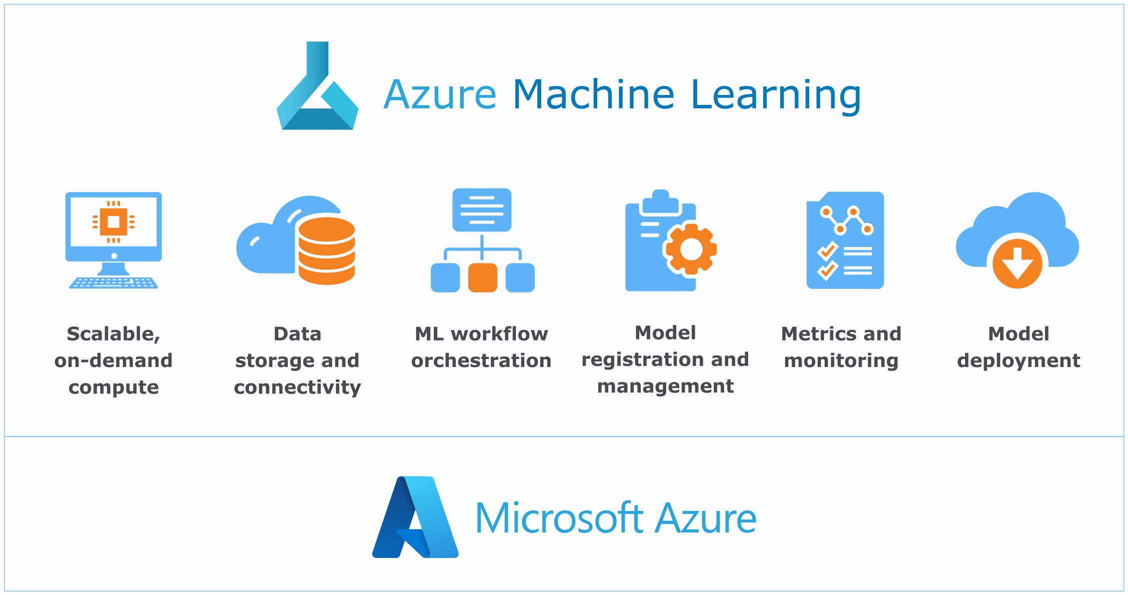

Table of Content

intetics
Where software concepts come alive™

4. Brief History of Development

Machine Learning has roots that go back to the mathematical modeling of neural networks. The first paper attempting to mathematically map out its thought processes and decision-making was published in 1943. A bit later, in 1950, Alan Turning started to classify machines as intelligent and unintelligent with the Turing Test. If a machine is classified as intelligent, it means it was able to convince a human being that the machine is also a human being.

From this point, there was a rise of intelligent machine learning algorithms and computer programs solving a variety of tasks:

1952 Arthur Samuel builds a program for an IBM computer to play checkers, and it improves with each game.

1956 The birthplace of AI is Dartmouth College, where John McCarthy invites well-known mathematicians, researchers, and scientists to brainstorm on thinking machines.

1963 Donald Michle develops MENACE (Matchbox Educable Noughts and Crosses Engine) that can learn how to perfectly play tic-tac-toe.

1967 The invention of the nearest neighbor algorithm gives computers basic pattern recognition. Used to plan routes for traveling salesmen.

1970 Seppo Linnainnmaa, a Finnish mathematician and computer scientist, publishes a paper about automatic differentiation (AD) of discrete connected networks of nested differentiable functions.

1973 AI winter, in which the British government cuts funding for AI research in all universities except for three.

Table of Content

1979 Kunihiko Fukushima's neocognitron is used for pattern recognition. Researchers from Standford University release the Cart, a robot that can navigate obstacles within a room.

1981 Gerald Dejong introduces explanation-based learning, allowing a computer to learn to analyze data and create a general rule for which to discard information it deems unimportant.

1982 Popularization of the Hopfield network — a type of recurrent artificial neural network, first described in 1974.

1985 The creation of NetTalk — a network that learns how to pronounce words, similarly to a baby.

1986 A paper on the parallel distributed processing framework, featuring neural network models for ML, is published.

1989 Development of Q-Learning — a model-free reinforcement algorithm that seeks the best action to take in any current state. Axcelis Inc. releases Evolver — the first commercial genetic algorithm package for PCs.

1992 Gerald Tesauro invents a program that can play backgammon and rivals the skill level of top-tier players.

1995 Tin Kam Ho, a computer scientist from IBM, publishes a paper on random decision forests (ensemble learning method). AT&T Bell Labs researchers release a paper on SVMs.

1996 An IBM chess-playing program defeats the reigning world champion of chess at the time, Garry Kasparov.

Table of Content

intetics
Where software concepts come alive™

1997 A paper on long short-term memory (LSTM) architecture (type of recurrent neural network) is published.

1998 The MNIST (Modified National Institute of Standards and Technology) database is released. Later, it is widely adopted as a handwriting recognition evaluation benchmark.

2002 The first open-source software library for ML, Torch, is released.

2006 Geoffrey Hinton, psychologist and computer scientist, coins the term "deep learning." Netflix launches an ML competition to create an ML algorithm more accurate than its proprietary user recommendation software.

2009 Fei-Fei Li invents a large database to aid visual object recognition development efforts — ImageNet.

2010 Microsoft releases the Kinect motion-sensing input device for the Xbox gaming console that tracks 20 various human features, 30 times per second.

2011 IBM Watson beats 2 champions on the Jeopardy game show. Google launches Brain AI.

2012 A presentation is given on the ImageNet classification — a model that significantly improves accuracy in ML and AI image recognition. Google releases a neural network able to teach itself to recognize cats through YouTube videos.

2014 Facebook releases the DeepFace algorithm to identify individuals in photos as accurately as a human can. Google unveils Sibyl, a large-scale ML project for predictive user recommendations.

2016 Google's AlphaGo beats a professional Go player, becoming the first program to do so.

Table of Content

intetics
Where software concepts come alive™

While the timeline covers key ML achievements and benchmarks, allowing you to notice the evolution of the technology, it's not comprehensive. With the rise of 4IR (Fourth Industrial Revolution), there have been numerous advancements in ML, data analytics, and science.

5. Total Market Volume

In 2021, the Global Machine Learning Market Size reached almost $14.91 billion, according to data published by GlobeNewsWire. The analysts admit fairly rapid market expansion, which is driven by technological improvements that result in enhanced system accuracy and ML integration in robots.

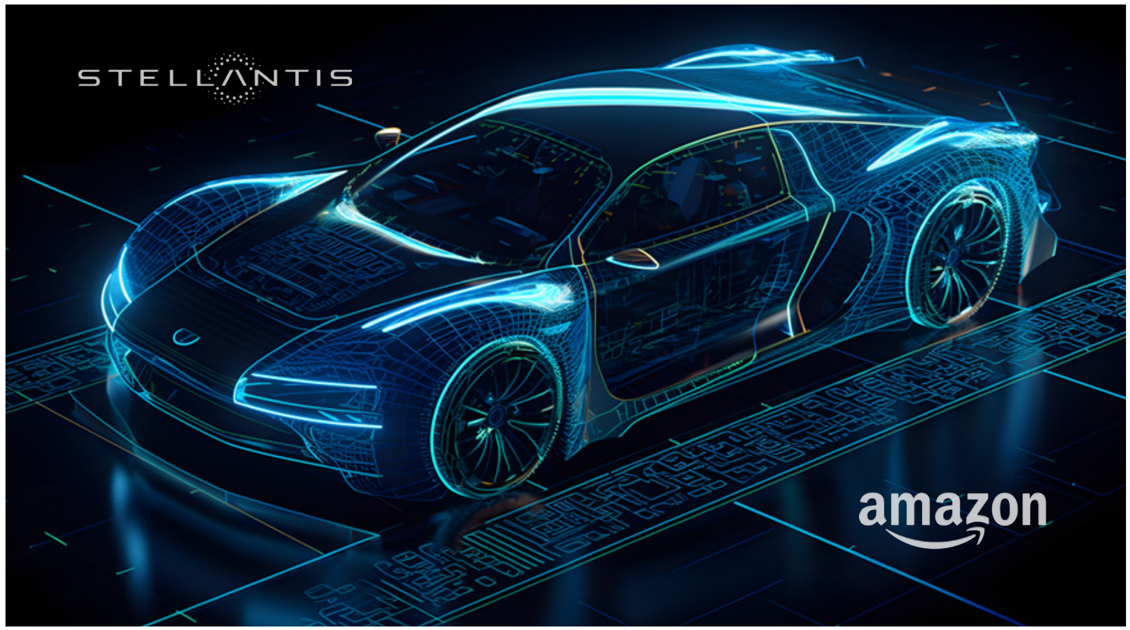

ML drives impact across a variety of industries, allowing for a more customer-centric approach. For example, analysts emphasize the Amazon-Stellantis collaboration as one of the recent achievements: the companies worked on a customer-focused connected experience across millions of vehicles to hasten the software transition to Stellantis.

Table of Content

Key ML market players are IBM, SAP, Oracle, Hewlett Packard, Microsoft, Amazon, Intel, Fair Isaac Corporation, SAS Institute, and BigML.

Forecasted Growth

The Global Machine Learning market is expected to grow at a CAGR of 38% by 2030, when it is anticipated to surpass $302 billion. Although a significant increase is forecasted in the services segment, the Machine Learning Platforms market is also predicted to grow at a CAGR of 33.6% between 2022 and 2028, according to Business Research Insights.

Speaking of the services segment, more focus may be devoted to automating the process of turning data into insights. This will allow businesses to get a better understanding of their customers and benefit from predictive modeling for anticipating customer behavior. The last one, the predictive maintenance segment, is anticipated to experience consistent revenue growth.

According to a recent report on MLOps (Machine Learning Operations), US-based ML decision-makers plan to increase their MLOps investments in 2023:

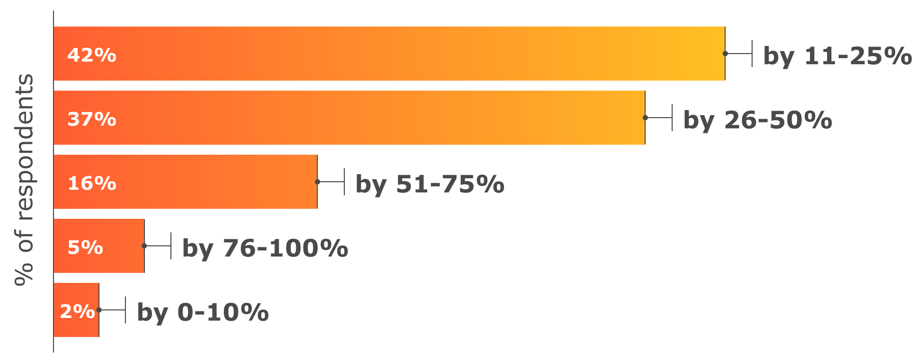

Table of Content

intetics
Where software concepts come alive™

6. How Do ML Algorithms Work?

Primarily, machine learning is set to work as a human brain to analyze data. An algorithm receives and analyzes input data to predict output values within a predefined range or dataset. The output is a trained ML model that can further apply the learned patterns to analyze similar datasets, developing intelligence over time. There are few to no limitations when it comes to the input data's format: it can be figures, text (powered by text recognition), graphics (image recognition), or voice (speech recognition).

The graph below illustrates the five core steps of training ML algorithms:

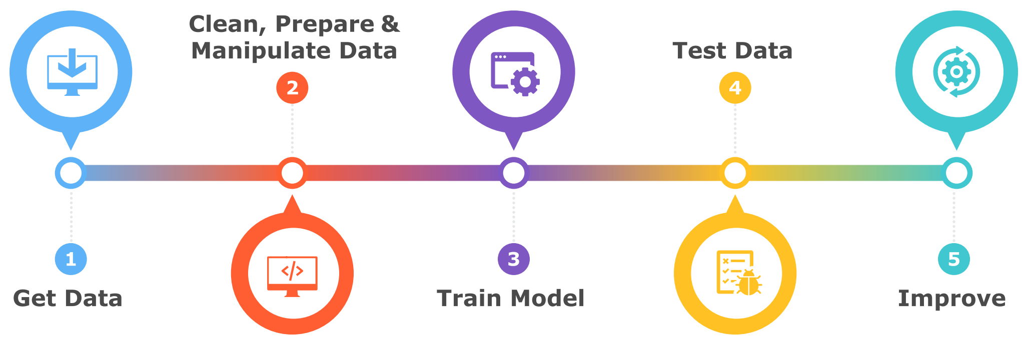

Sourcing Raw Data

It can be real-time data from IoT devices (for example, wearables in Healthcare or IoT devices embedded in autonomous vehicles), data from online repositories like Kaggle, etc.

Preprocessing

Since raw data is often unorganized, comes in different formats, is missing components, and has noisy elements, it needs to be cleaned and further manipulated to be organized in a clean dataset. This is often addressed through feature engineering — dataset manipulation through addition, deletion, combination and mutation — which allows you to improve the accuracy of the

intetics
Where software concepts come alive™

model. Then, it has to be converted into valid formats for the chosen ML platform.

Training

An ML algorithm receives input data from the training dataset and applies sophisticated mathematical modeling to learn and produce an output.

Model Validation

The model is applied to test data (preprocessed in step 2) to evaluate its performance and accuracy. If the results are satisfactory, the model can be applied to real cases. If there are some discrepancies, it's further trained or retrained.

Improvement

With a lot of new data being added to the dataset, the model loses its performance and accuracy. To ensure the same level of output, it needs to be constantly improved and trained.

intetics
Where software concepts come alive™

Although the process of training ML algorithms is quite standard, there are four main algorithm types:

1. Supervised Learning

The approach implies teaching the algorithm by example. A dataset includes desired input and output values, while the algorithm must discover a method to get them. In other words, the algorithm identifies patterns in data, learns from observations, and makes predictions. Since operators know the correct answers, they correct the predictions until they meet the required level of accuracy.

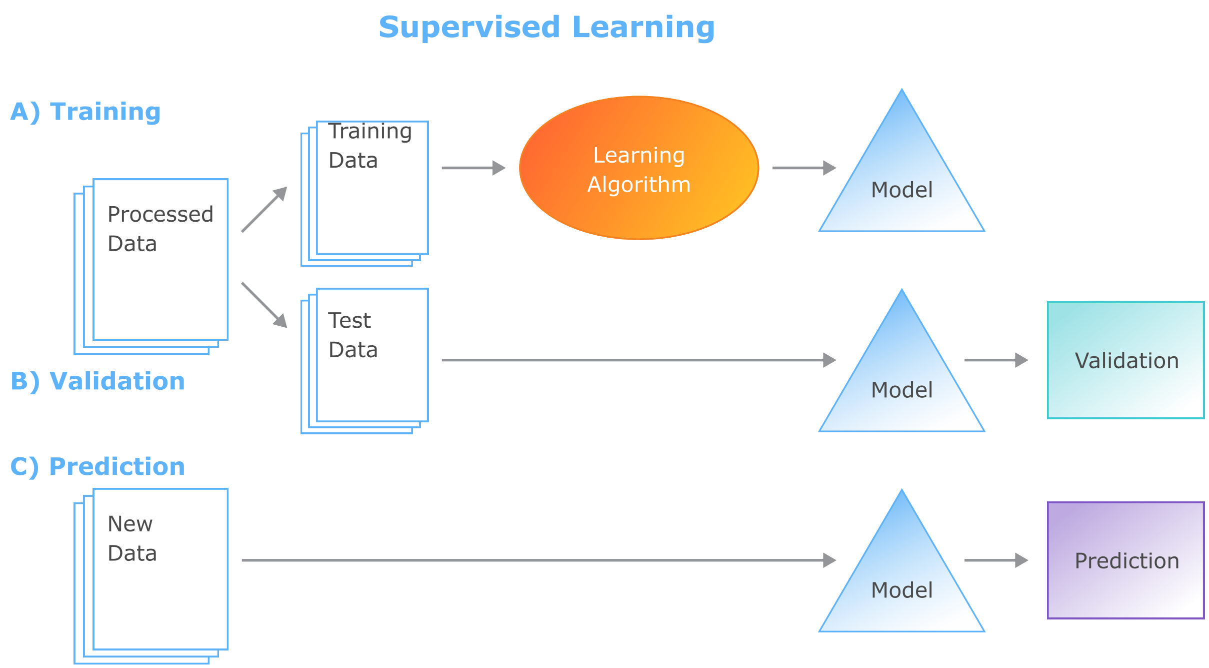

Here are the two main algorithms applied in supervised learning:

Classification

An ML algorithm concludes to which category a value belongs based on some similarities or differences—for example, if an email is spam or not.

Table of Content

intetics
Where software concepts come alive™

Regression

An ML algorithm estimates and understands the relationships between variables, where there is one dependent variable and other changing ones. This is often used for predictions and forecasting, for instance, when analyzing future trends.

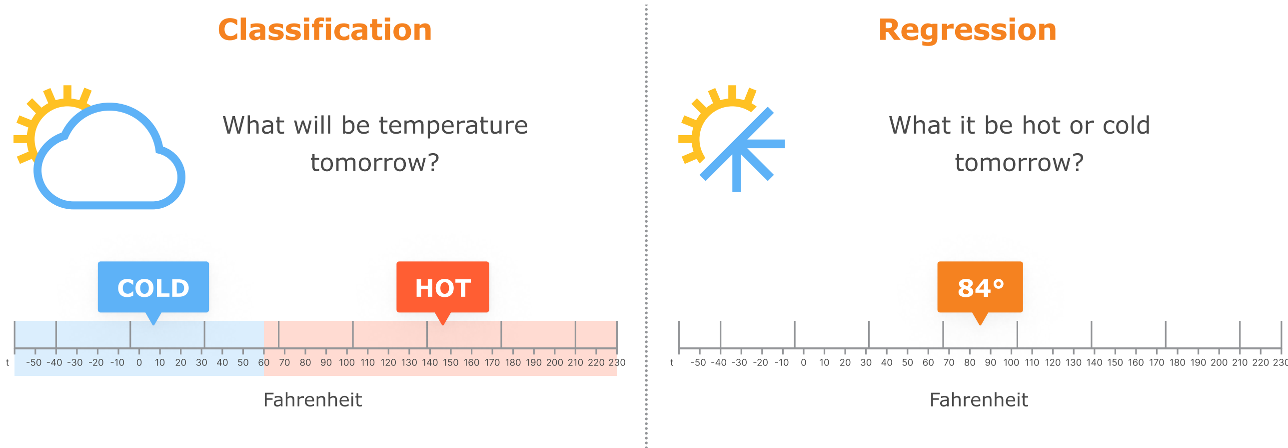

Data Labeling

Supervised learning is enabled through data labeling—a process of identifying raw data, be it an image, text file, speech, or video, and adding at least one informative label for an ML algorithm to have context to learn from. Some examples of labels are:

- **Healthcare**: whether an X-Ray contains a tumor
- **Chat-bots**: uttered words in an audio recording to determine the context and provide a relevant reply
- **E-commerce**: if a photo contains clothes or food so a relevant suggestion can be provided

The labeling process primarily involves a human who labels a given dataset. It's known as the tagging process, which can be down to yes/no, as well as identifying specific pixels of the image. The ML algorithm then learns based on the human-provided labels—which is model training.

2. Unsupervised Learning

The approach implies that an algorithm studies data to identify patterns. It doesn't involve data labeling. Instead, an ML algorithm analyzes a dataset to determine correlations and relationships, as well as to organize the data and describe its structure. The more data is assessed, the more refined the model is.

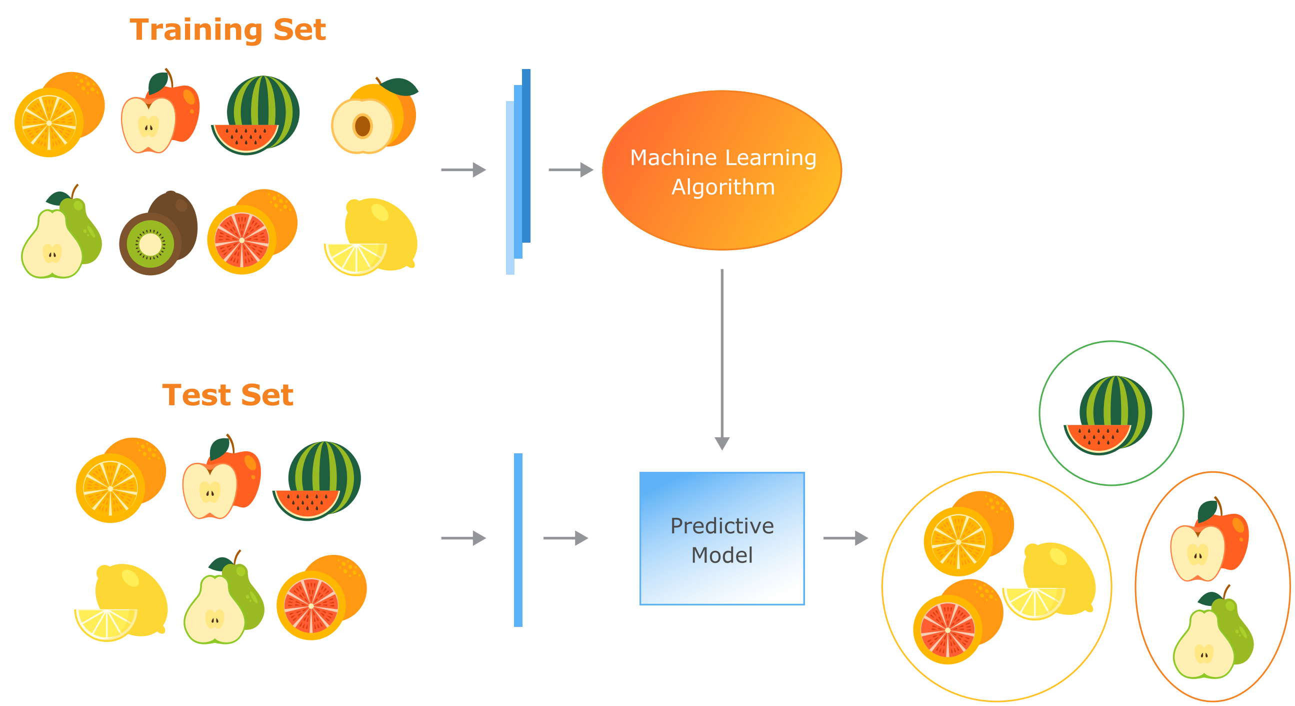

Two major ML algorithm types are:

- **Clustering:** based on the defined criteria, sets of similar data are grouped. This allows data to be segmented into several groups and for a deeper analysis to be performed on each to find patterns.
- **Dimensional Reduction:** it implies reducing the number of considered variables to determine the required output.

Table of Content

3. Semi-Supervised Learning

This approach combines both supervised and unsupervised learning—hence, an ML algorithm analyzes labeled and unlabelled data. Since labeled data has informative tags, the model uses them as a basis to learn to label unlabelled data.

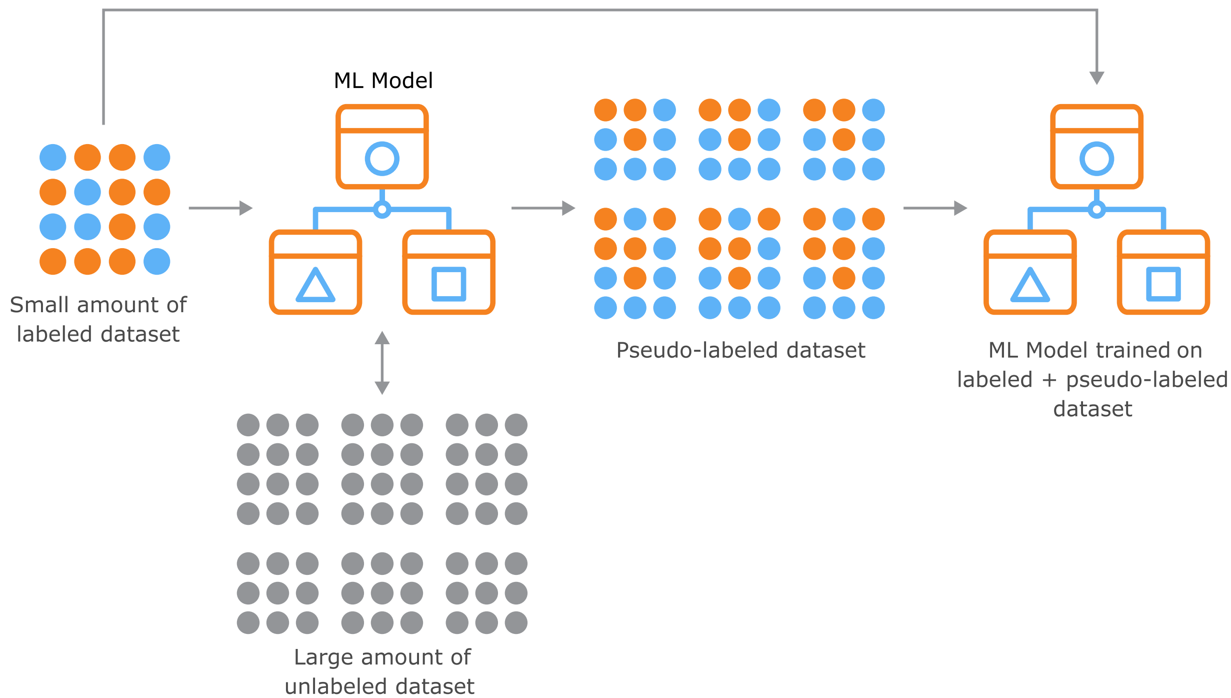

Since unlabelled data is abundant, easy to get, and affordable, semi-supervised learning allows you to quickly and rather effortlessly label it, without performance and accuracy losses.

Table of Content

4. Reinforcement Learning

Reinforcement learning focuses on regimented learning processes. An ML algorithm is provided with a set of actions, parameters, and an end value to learn trial and error. The algorithm tries to explore various options while monitoring and assessing different results, in order to determine the most optimal one. The model learns from past experience to propose the most optimal solution in the current case.

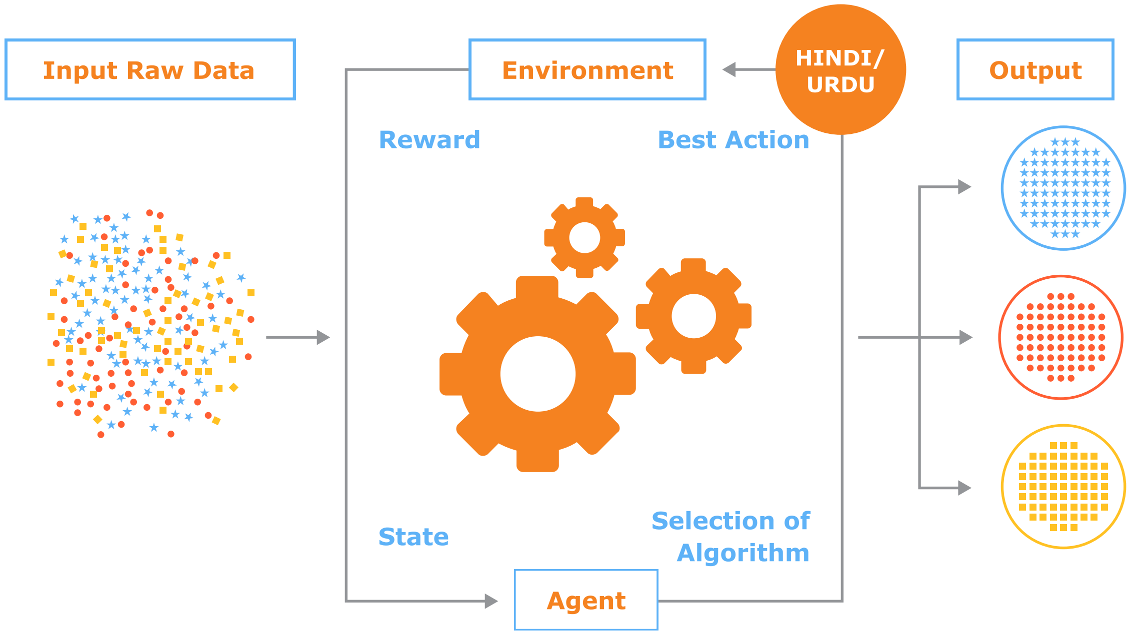

Key Challenges When Training ML Algorithms

Poor raw data quality

Raw data is often unclear and noisy, which requires a long preprocessing stage to bring it to a common state. Outliers and unwanted features must be removed, and missing values must be filtered. The process is time-consuming.

Table of Content

Expensive data labeling in specific niches

Considering data specifics in healthcare or finance industries, data labeling can be properly done by specialized experts.

Overfitting/underfitting

If underfitting refers to a model that can neither model training data nor generalize to new data (which often requires experimenting with algorithms), overfitting results in a model that studies the training data too well. As developers aim to get the maximum out of the training set, they may provide too much noisy detailed data. An algorithm then picks up noise and fluctuations and learns them as concepts. They, however, can't be applied to new data, so the model fails to properly generalize.

Poor performance when testing and deploying domains are different

Often, an ML model is trained on one domain and then applied to another, for example, the model can analyze LiDAR drone data of the surface in summer but fails to do so in winter. Training a model on a new domain from scratch requires a large testing dataset, which is time-consuming.

Accuracy drops

As new data is flowing into the model on the production, the model loses accuracy. As a result, it needs to be constantly trained or retrained.

Some challenges, like resource-intensive data preprocessing, issues when deploying a model on another domain, and long model training can be addressed through Transfer Learning. It's a process of applying elements of a pre-trained model to a new ML model, aimed to deliver similar outputs.

The method is based on generalization to adapt to new, unseen data. That's why only the knowledge that can be used by another model in different scenarios is transferred. Such models can be used in changing conditions and with different datasets, which speeds up the development and often optimizes the costs.

Table of Content

intetics
Where software concepts come alive™

7. Main ML Tech Architectures, Tools, Stacks

Although ML applications significantly vary across industries and use cases, the major infrastructure layers are rather uniform. Here are 8 key layers:

1. Data Layer

The data layer involves data warehouses, which are integrated with the ML stack. Some examples are:

- **Cloud-based**
  SnowFlake
- **AWS databases**
  RDS, Redshift, Aurora, S3-based data lake

AWS databases are often a good option for ML use cases because of their enhanced scalability in terms of query patterns and dataset sizes, compared to traditional databases.

2. Compute Layer

Large-scale computing implies setting up and operating a cluster to handle workloads. They are often rather diverse and numerous, as new data flows are constantly generated. Today, this level is often implemented with cloud-based auto-scaling systems like AWS Batch, which allow you to handle high workloads.

For data-intensive applications with diverse needs, it's best to use general-purpose container orchestrations like Kubernetes. It allows you to configure a scalable batch compute layer. However, its complexity is a significant drawback.

Table of Content

intetics
Where software concepts come alive™

3. Orchestration Layer

At this stage, given a workflow or DAG (Directed Acyclic Graph) definition, the workflow orchestrator executes the tasks defined by the graph in order using the compute layer. To ensure reliable production workflows execution, a system should be scalable and available:

- **Airflow**
(open-source workflow orchestrator)
- **Argo**
(performant on Kubernetes)
- **Google Cloud Composer**
(managed solution)
- **AWS Step Functions**
(managed solution)

4. Versioning Layer

A strong versioning layer allows you to manage the dynamic environment in which an ML app exists. It’s similar to taking snapshots of immutable points in time of models/data/code. Here are some tools:

- **Git**
(open-source distributed version control system)
- **GitHub**
(hosting service for software development and version control through Git)
- **Metaflow**
(Python library to develop, deploy, and operate data-intensive apps)
- **MLFlow**
(a framework supporting ML lifecycle)

Git and GitHub are good for managing the versioning of code and usual workflows of software development. If you need to track experiments, models, and data, it's best to use frameworks. Both Metaflow and MLFlow have custom versioning solutions.

5. Software Architecture Level

It’s often implemented with Python, which allows non-trivial business logic and mathematical concepts to be converted into an executable form. Data-centric programming usually implies Python code written in Jupyter notebooks in order to immediately receive the output, which is beneficial for hypothesis testing and fast debugging.

Table of Content

6. Model Operations Layer

The layer comprises tools for several purposes:

- **Model deployment**
  Seldon for scaling to thousands of production ML models and having advanced ML capabilities out of the box
- **Model monitoring**
  Weights and Biases for experiment tracking, dataset versioning, and model management
- **Model explainability**
  TruEra for ML testing and evaluation

7. Feature Engineering Layer

The layer is relevant for supervised and semi-supervised learning algorithms that require labeled data. Feature engineering involves:

- Scaling and normalization to simplify learning through adjusting the range and center of data
- Filling in missing values based on knowledge, ML techniques, and heuristics
- Feature selection to remove redundant or unimportant features
- Feature coding to define different categories by choosing a set of symbolic values
- Feature construction to create new features from existing ones
- Feature extraction to provide higher-level features useful for learning

This can be addressed through feature stores (example, Tecton) and labeling solutions (example: Scale, Snorkel). Feature stores allow you to connect similar data transformations across a team, while labeling solutions help overcome hand labeling challenges.

Table of Content

8. Model Development

This layer deals with mathematical modeling to determine a modeling technique, suitable model architecture, and model parameters. There are off-the-shelf libraries like PyTorch and sclkit-learn that address the needs.

9. Neural Network Architectures

If neural networks were initially used for simple classification problems, then, today, they are applied in such domains as chatbots, recommendation engines, as well as medicine. The neural network architecture is built from neurons—individual units that mimic the biological behavior of a human brain. The components of a neuron are:

- Input
- Weight
- Transfer function
- Activation function
- Bias

Neuron

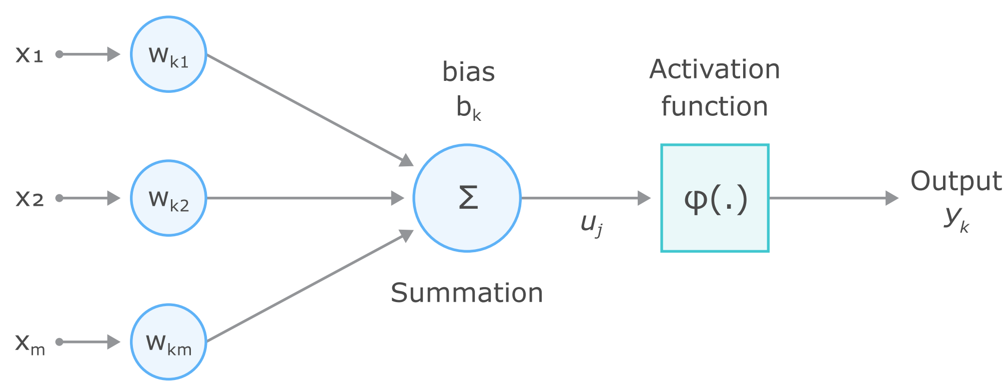

Table of Content

When several neurons are stacked together, they form a layer. Several layers piled onto each other comprises a multi-layer neural network. Its structure is as follows:

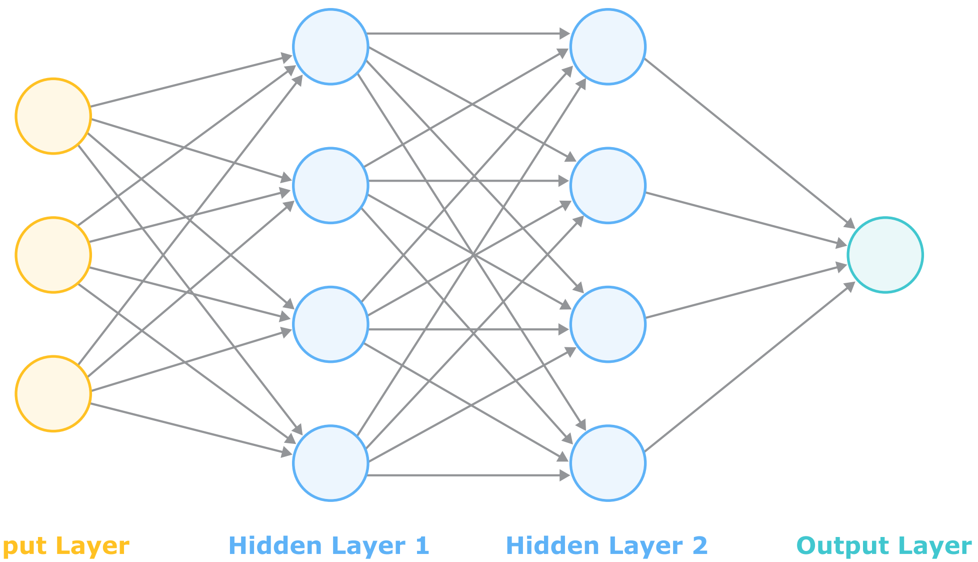

- **Input layer**
  It's a visible layer that consists of data that was fed from external sources (for example, a web service or a CSV file)
- **Hidden layers**
  These are intermediary layers for computation and feature extraction. They can be interconnected to search for various hidden features in data. Later, hidden layers perform more complicated tasks: for example, in image processing—complete object identification, while the first hidden layers would determine shapes, boundaries, and edges
- **Output layer**
  It's the most important layer that makes a prediction based on the model's learnings

Table of Content

Neural network architectures vary depending on the use case:

- U-Net for segmentation

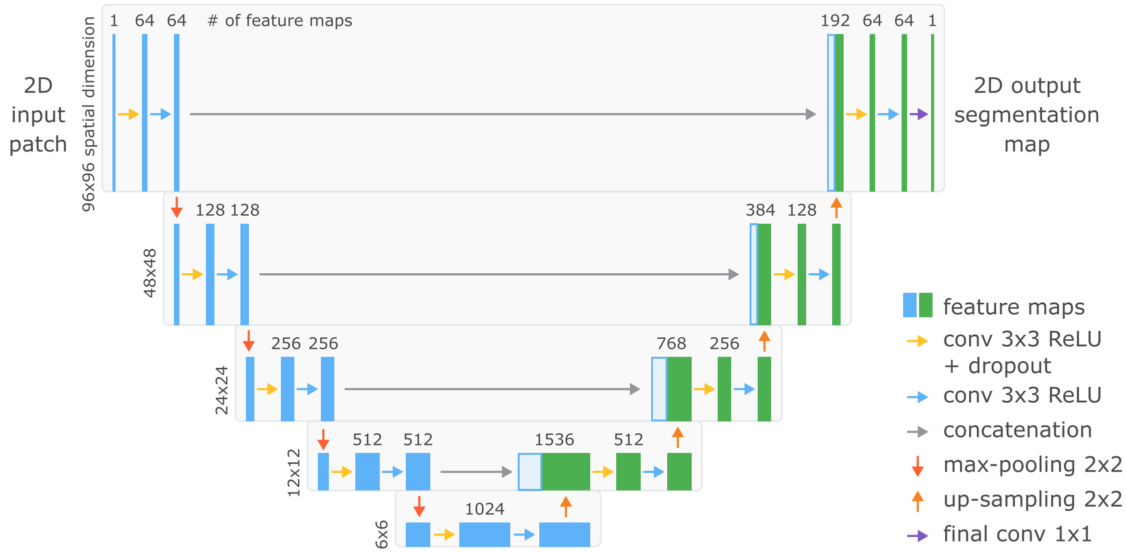

- ResNet for classification

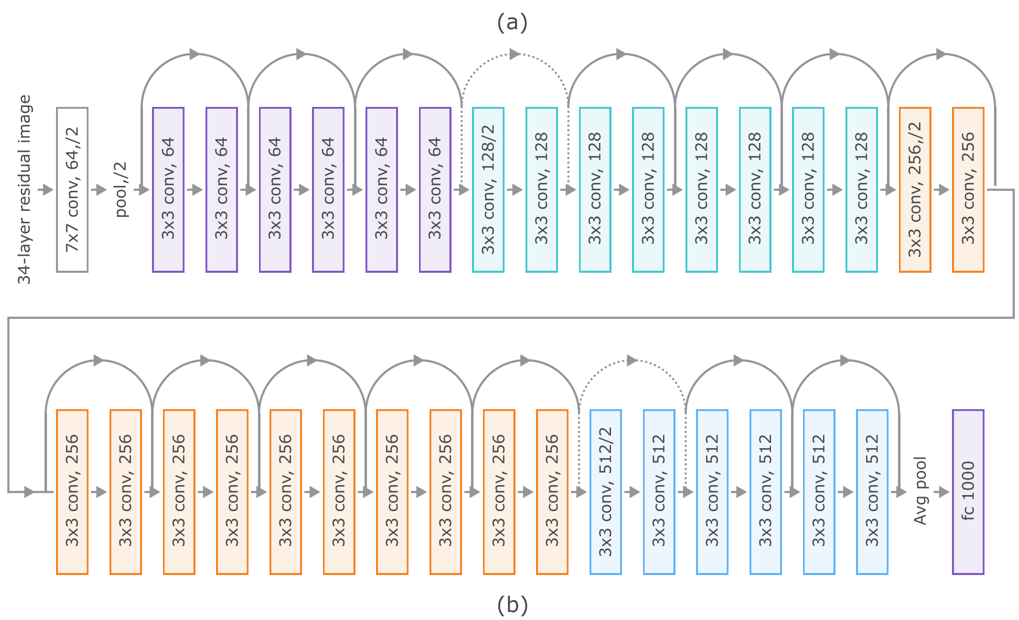

Table of Content

intetics
Where software concepts come alive™

8. Main Applications and Impact Delivered

Machine learning drives impact across virtually any industry that collects data and aims to benefit from its advanced processing. Here is a list of some industries that experienced significant disruption as a result of applying ML, with major use cases and delivered impact:

Automotive

Machine learning fosters the development and adoption of smart vehicles, equipped with advanced onboard computers. By collecting and analyzing data on its surroundings— like passing-by vehicles, buildings, pedestrians, lane markings, road signs, traffic lights, etc.— and correlating it with traffic rules and your destination, you can enable your vehicle with autonomous functions.

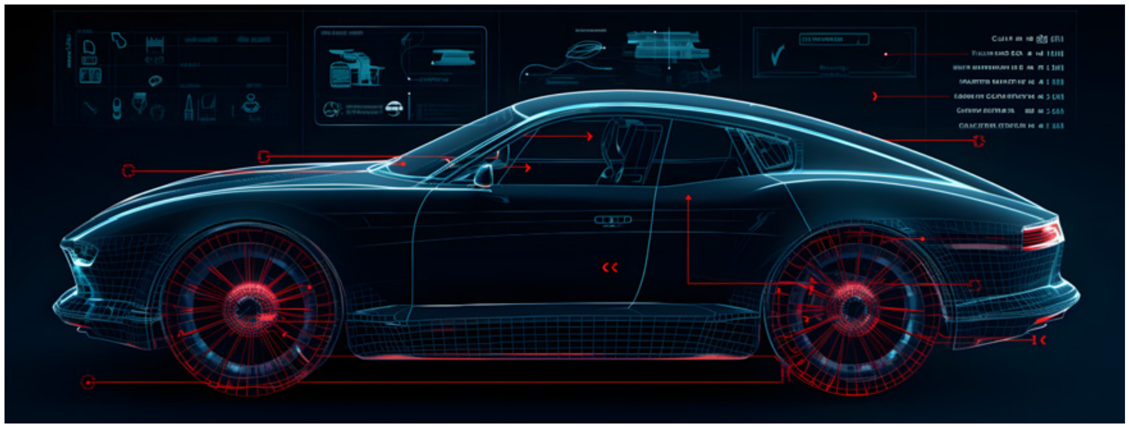

The Society of Automotive Engineers developed 6 Levels of driving automation as part of industry standards:

- **Level 0-2: Driver support features**
  They range from providing warnings and momentary assistance (example: automatic emergency braking, lane departure warning, blind spot warning) to steering and brake/acceleration support (example: adaptive cruise control, lane centering)

Table of Content

- **Level 3-5: Automated driving features**
  They imply that you aren't driving, unless the features can drive in specific circumstances. At level 5, there are no condition limits. Some examples are traffic jam chauffeurs (Level 3), local driverless taxis with or without pedals or a steering wheel (Level 4,5)

Self-driving (SD) cars are considered the next step, with some companies like Tesla already providing the functionality. Its autopilot can handle driving, including lane changing, autosteering, traffic-aware cruise control, and auto parking. This is possible through computer vision and data processing algorithms that allow self-driving cars to:

- Identify and classify the detected objects, as well as measure their direction, speed, and acceleration to respond accordingly. For example, Waymo software collects and analyzes data from radars and sensors to determine if there’s a green traffic light to start moving, as well as whether the lane is blocked.
- Predict the behavior of obstacles based on the training models to distinguish pedestrians from cyclists and predict their possible speed and direction. Waymo Driver’s ML model, for example, is based on 20+ million miles of real-world driving, as well as on 20+ billion miles in simulation. Relying on the data, it predicts numerous paths other road users may take.

Table of Content

intetics
Where software concepts come alive™

Case Study 1

Model Development for Automatic Parking Sign Recognition — 90% Accuracy Including Recognizing Unstandardized Signs

Business Domain
Parking services

Project Type
Machine Learning Services

Client

The Client delivers innovative products for the automotive and transportation industries, such as real-time parking and traffic information and solutions that facilitate autonomous vehicles' safety testing and deployment. They also provide new insights to various other industries to make better business decisions by understanding how people move throughout the day.

Project

Building an ML pipeline for parking sign recognition to detect a sign's location and recognize the restrictions and time bounds for parking with no less than 90% accuracy, including recognizing unstandardized parking signs with texts.

The algorithms should recognize and classify the input images from mobile mapping platforms installed on the cars.

Wrapping ML model into a stand-alone web service and its integration with a customer analytics platform.

Objective

Developing of a web service to automate the collection and entry of parking data.

Team Reinforcement

The Client does not have Data Science experts but wanted to test how machine learning could improve and speed up data processing workflows to provide actual data updates for the biggest clients. Intetics software engineers offer the required expertise.

Challenge

The project had several challenges: data analysis from several on-street images providers, development of an image download and pre-processing workflow, building a solution for parking sign detection and classification, extraction of text information about on-street parking from cropped images, integration of the end solution into Clients' cloud infrastructure.

Table of Content

intetics
Where software concepts come alive™

The Client approached Intetics with a need for accurate detection of parking locations, restrictions and time limits. The detection was based on roadside imagery covering the area of a city. No less than 90% of the signs had to be identified correctly. Parking signs are not standardized. This makes the recognition more complicated because the process is based not on the sign in general but on the actual text on the sign. To solve this issue, object detection supported by image classification was used.

Due to the method of roadside data collection, some signs were present on several photos. That required the additional task of grouping data for one parking sign and determining the exact location. The algorithm also had to deal with poor image quality, different types of weather, light conditions, time of the day, seasons, fog, and other distortions.

Quick Facts

- Semi-automatic data labelling
- Recognition of unstandardized parking signs with texts
- 90%+ accuracy of recognition

Technologies

Python / Flask / Tensorflow / Keras / OpenCV / Scikit-learn / Tesseract / AWS

Solution

01

To identify a specific parking sign, several steps were involved: identification of an image with a sign on it, detection of the parking sign on the image, and recognition of the text within the OCR component using basic NLP. Machine learning based on the TensorFlow framework and Keras library.

02

The project started with the analysis of data and sources and resulted in an estimate for the data processing phase.

03

The next stage included data labeling and data transformation. Semi-supervised active learning techniques were used and saved up to 70% of the time spent on manual data labeling. To eliminate errors associated with various distortions, all types were identically distributed over labelled datasets: training, development, and testing. A script for automatic data transformation was also developed during this stage to reduce the data to a single format (color, rotation, tilt, etc.)

04

The next stage covered iterative development and initial training of the Machine Learning model. Here it was possible to give a precise prediction of the accuracy that could be achieved with sign recognition and time for training the model.

05

Within the postprocessing stage, sign location clarification and the same sign association problems were solved.

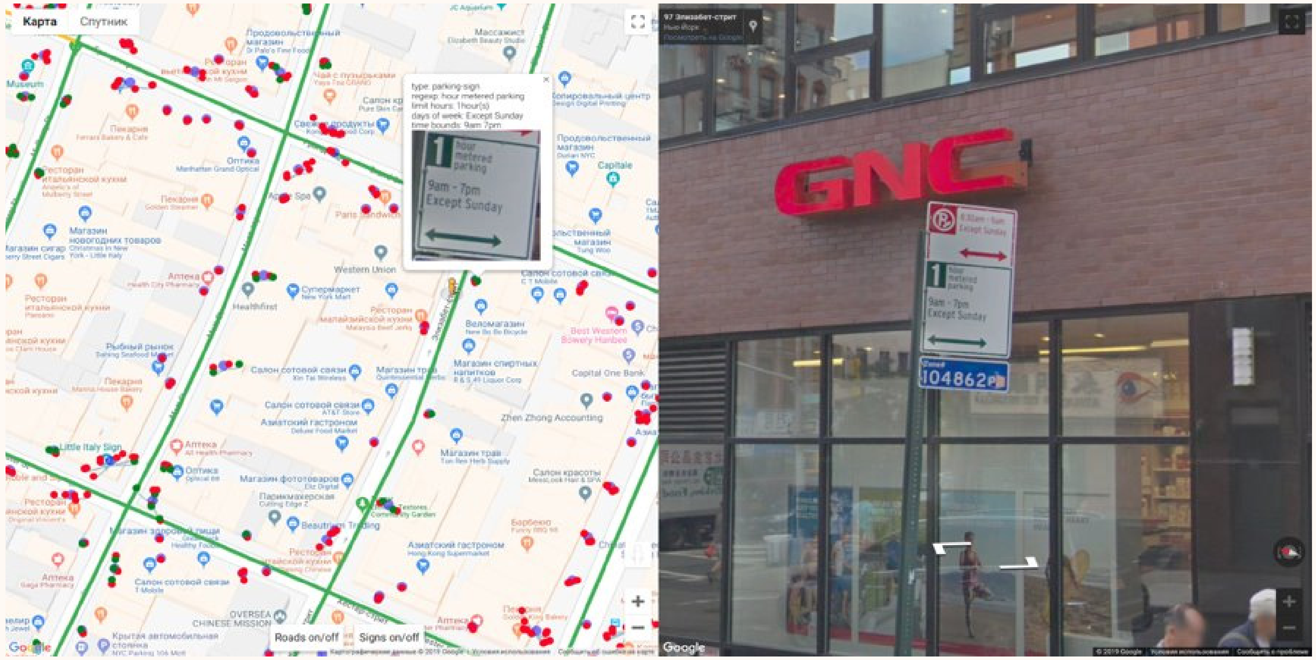

Client Reference

Due to the predictable development approach, Intetics provided a detailed analysis and estimates for every step of the project and ensured all expectations would be met.

Benefits and Results

- The algorithm was successful in identifying 85%+ of signs on the streets in the city. 90%+ of the signs were identified correctly.
- All the ML and data processing algorithms were implemented in a single web application and supported with a detailed description of the model and project documentation.
- The Client received a solution that automated parking sign detection and reduced the workload significantly.

Team: 6

Project Manager, Data Analysts, Data Scientists, Python developer

More information on the use case can be found here.

Table of Content

intetics
Where software concepts come alive™

Table of Content

GIS/Geospatial

Advanced geospatial data processing drives impact in a number of sectors and can be applied to improved urban planning, weather forecasting, mapping, site selection, visit attribution, network planning, environmental risk assessment, etc. However, there are several data collection and processing challenges that are overcome, thanks to ML. Among them can be found such barriers as:

- A large scope of constantly generated data
- Hard-to-reach and extreme environments
- Large-scale territories to analyze
- Poor data quality, often gathered during automatic feature extraction, and data fusion from different sensors

Trained ML algorithms allow for enhanced UAV drone data processing and LiDAR data processing:

- DSM and DTM extraction allow the classification of vector features of natural terrain like rivers and ridges.
- 3D models and 3D measurements allow for remote LiDAR data analyses and classification.
- Land use and topographical maps instead of photos or LiDAR data from drones allow for accurate data assessment and analyses.

Besides this, LiDAR processing can align strips and calibrate boresights, smoothen and clean the point cloud, and digitally elevate models. All this makes POI geofencing accurate and relevant to your business needs.

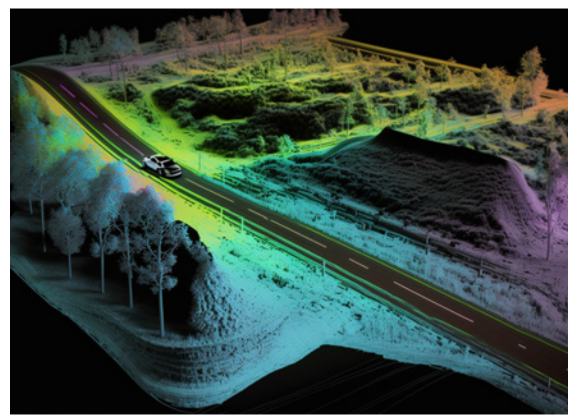

Table of Content

intetics
Where software concepts come alive™

A Predictive Algorithm for Upcoming Wildland Fire Conditions and Improved Current Fire Perimeter Map With 84% Predictive Accuracy for a US Tech Company

Business Domain
GIS and Geospatial solutions, AI

Project Type
Web app

Client

WTVIII, Inc. is an innovative leader in system integration, consulting, and software development, with 20+ years of experience helping customers streamline operations through automation.

Project

Creating a model to predict how the fire will spread further in fire-affected areas with Computer Vision. Refactoring and enhancing a web portal that provides critical wildland fire information to public consumers.

Objective

The Client needed to enhance interactive fire and smoke maps of current wildland fire conditions using content from Google, USGS, NIFS, as well as promote a nation-wide data exchange to enable interoperable wildland fire operations and assess the fire danger level throughout the USA. Besides, a history dataset for all reported wildland fires in the USA and 1-9 day forecasts for fire incidents were to be established.

Team Reinforcement

The Client required a highly experienced Remote In-Sourcing® Team to develop industry-specific ML algorithms to solve the Objectives.

Table of Content

intetics
Where software concepts come alive™

Challenge

WTVIII, Inc. is an innovative leader in system integration, consulting, and software development that delivers world-class solutions to public and private clients with 95%+ Dunn and Bradstreet client satisfaction scores.

The project's technical objectives were:

- To store public data feeds with current conditions of fire and smoke spreading
- To assess fire perimeter risk
- To predict smoke conditions, intensity and direction
- To deliver website integration and component prototype.

The Client didn't have the required in-house expertise, so they looked for a reliable partner to implement the solution at reduced costs, without risk of failure, and in time.

Quick Facts

- 84% predictive accuracy
- 30% budget savings during development
- Delivered the project in 3 months instead of 5

Technologies

Web Server: Nginx 1.18 / Language: PHP 8 / Framework: Laravel tabase: PostgeSQL 12 / Gis extention: PostGIS 3 / CRON Jobs / TypeScripr / ReactJS / react-leaflet / NPM 8.15 / React 18.2 / Leaflet 1.8 / Python

Solution

★ 01

Fire forecasts decrease negative affects on human health through smoke emissions and safety risks, as well as allow for better forest resource management and planning for wildland fire evacuations.

★ 02

The ML model forecasts fire and smoke spreading with an industry-leading accuracy of 84%.

★ 03

AI allows users to click anywhere on the map to get a prediction on how many days it will take to extinguish a fire, if one takes place at the chosen area. History data is also accessible.

★ 04

The data was collected from various sources and visualized for the end user on a map. It is presented in GeoJSON format and raster images.

★ 05

Delegating the task to the Remote In-Sourcing® team allowed WTVIII, Inc. to complete the project 2 month faster than required and to reduce the development costs by 30%.

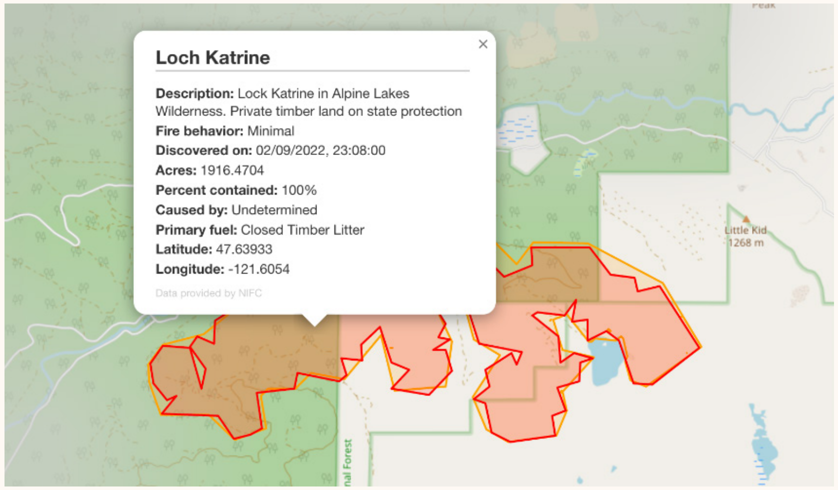

Client Reference

We would like to highlight the professional work of the team, which not only efficiently completed the short-term initial tasks before the specified deadline but also expanded the project by proactivity, ideas, self-organization, and predictions. Throughout the project journey, they proposed new solutions, which were successfully implemented in the deliverables. It allows us to form our business plan, and the work on the project continues.

Benefits and Results

- The out-of-the-box solutions with thought-through technical details allowed for improved fire and smoke maps with enhanced predictive algorithm.
- The Client successfully implements ready-made project solutions on their side and clearly understands what needs the product should cover.
- The AI/ML models are being enhanced by the best practices within the domain and are becoming more sophisticated every day, facing the most creative challenges.
- 80% of deliverables were ready in 3 months instead of the projected 5 months. The Remote In-Sourcing® team proposed several ways to continue developing the Client's business, so the cooperation is still ongoing.

Team: 5

Project Manager, ML engineer, Frontend engineer, Backend engineer, QA

More information on the use case can be found here.

By processing geospatial data, ML algorithms also foster impact across the transportation industry. Relevant traffic alerts and navigation enable precise routing based on the number of current drivers, their routes, and historical data, thereby avoiding traffic jams. This is used not only by navigation software but also by taxi, location-sharing, and delivery apps. For example, Uber implements its own DeepETA model to enable enhanced delivery and pick-up.

Moreover, ML allows for improved parking solutions, which results in optimized parking slots and an improved driver experience. For example, based on geo data, combined with a trained ML model taking into account weather conditions, holidays, traffic, and history data, users can identify vacant parking slots at the time of their arrival and navigate to them.

Healthcare

Machine learning enables impact across healthcare in various directions. This includes predictive treatment, drug development and discovery, virtual communication with patients, remote assistance, and optimized workloads of medical staff—often through medical records reading and analysis. Let's discover some most impactful cases.

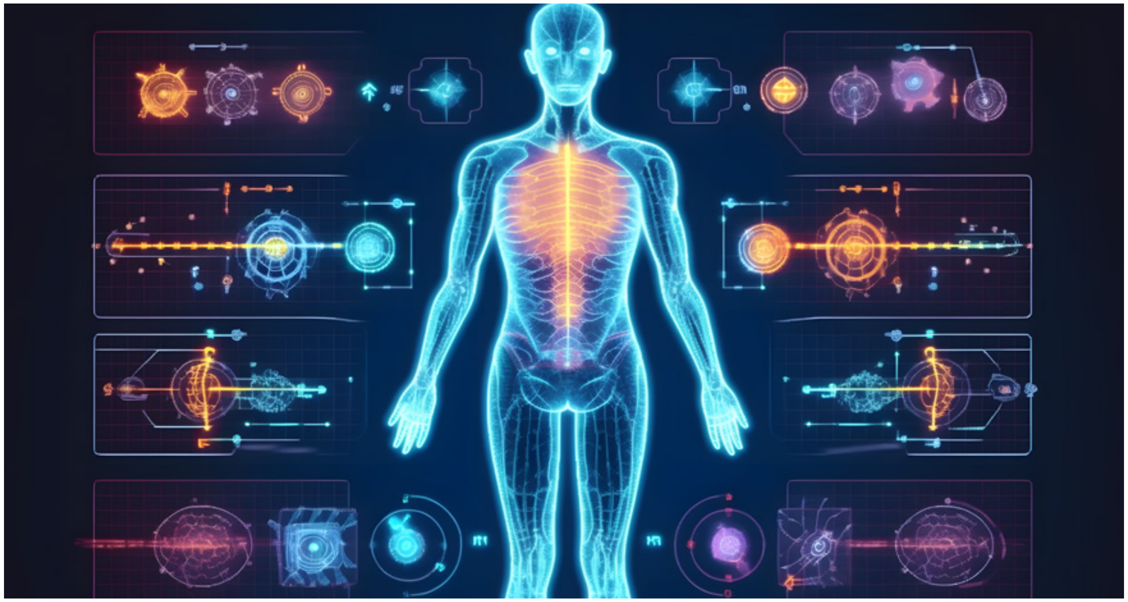

- **Remote treatment and diagnosis**
  By leveraging patients' historical data, ML models can assess a patient's symptoms, determine the possible causes, and provide recommendations, as well as enable virtual assistants like Symptomate to provide instant remote care.

  The technology can also be implemented in vending machines to instantly provide patients with treatment. For example, an automated lens scanner embedded in a vending machine allows patients to have their eyesight assessed on demand and instantly order new lenses. If there are no available options that meet the patient's requirements, custom lenses can be designed and delivered to the vending machine for pick up in a couple of days

- **More accurate and timely treatment through predictive analytics**
  By analyzing patient records, often together with real-time medical data collected from wearables, ML models can allow for predictive treatment. For example, they can identify patterns in a patient's cardiovascular history to predict potential heart failure. Such a solution is developed by KenSci. It can predict risks and provide improved care for chronic kidney disease, hospice patients, and those who are at risk of readmission and sepsis.

- **Improved patient service on site**
  By analyzing historical data like the number of visits at specific times of the day/week/month, available staff, emergency beds, and the layout of emergency rooms, an ML model can predict the required number of specialists and beds at a specific time/day.

- **Personalized treatment**
  Based on a deep analysis of patient records, ML models can recommend more tailored treatment, which improves overall care standards. Tempus, for example, delivers extensive cancer research in order to deliver more personalized treatments based on unique patient conditions.

- **Drug development and enhancement**
  By leveraging both machine learning and computational biology, drug development can be more efficient and cost-effective. For example, Insitro builds predictive models from massive biological datasets and then applies ML algorithms to sift through data in order to reveal trends like new disease subtypes. This allows the specialists to adjust medicine production, allowing them to tackle evolving diseases.

- **Surgeon support for increased accuracy and care**
  Through image recognition and machine learning, practitioners can have extensive support during surgeries. For example, Microsoft's Project InnerEye uses 3D radiological images to differentiate between healthy anatomy and tumors. This facilitates radiotherapy and surgical planning.

Case Study 3

Medical Device Company Received a New Application That Helps Surgeons Navigate Neural Implant Insertion

Business Domain
Healthcare and Life Sciences

Project Type
Desktop application for surgeons

Client

Established in 2000, BrainsGate is a medical device company committed to developing innovative therapies for patients suffering from Central Nervous System (CNS) diseases. BrainsGate's platform technology involves electrical stimulation of the Sphenopalatine Ganglion (SPG), a nervous center known to increase cerebral blood flow.

Project

Development of a desktop application executed on special computers that assist a surgeon during an operation by inserting an implant into the patient's brain. The application helps in navigation

Objective

The Client required a new application that would navigate a surgeon in inserting an implant into the patient's brain. For this purpose, MRI images of a patient's head are processed by the software; ML&AI algorithms are used to build a path for an implant, which is shown on the screen on top of an MRI image.

Team Reinforcement

BrainsGate doesn't have a software development team. The Client's team consists of doctors and scientists. Intetics' Remote In-Sourcing model was used to create a team of engineers that developed the application. The formed development team led the app development process while reporting to BrainsGate and the Client's Product Management

Table of Content

intetics
Where software concepts come alive™

Challenge

BrainsGate is exploring several applications for their technology and is currently focusing on the treatment of acute ischemic stroke. BrainsGate's technology is based on established scientific evidence that electrical stimulation of the Sphenopalatine Ganglion (SPG) increases cerebral blood flow.

The Client had many doubts about the technology selection. The application uses C++ for speed and control. On the other side, the application needed a friendly and functional user interface.

The Client doesn't have their own software developers and expected a vendor to provide a high level of service and develop the application that would be used during very sensitive surgery.

Quick Facts

- FDA Approved solution
- Designed 100% of Client's application and implemented it
- The Client is able to concentrate on medical solutions, not technical details

Technologies

C++ / ML&AI / Qt / Windows API / DirectX / OpenGL / Direct3D / DICOM / Multithreading / CT / MPR

Solution

★ 01
The Client received the application, which helps surgeons save many lives and improve a patient's condition after an ischemic stroke.

★ 02
The created software fully emulates and presents to a surgeon in 3D view everything that is happening while inserting the implant.

★ 03
The software complex includes a special camera and pointers that help in navigation.

★ 04
The application is a part of the hardware and software complex that consists of a special camera, implant injector, optical markers, and CT marker. This is a very complicated solution that shows the patient's head and the position of the implant in real-time.

★ 05
The Remote In-Sourcing Team® was formed from high-level professionals that have specific knowledge not only in programming but also in the medical domain. The programmers were required to work with CT and MRI images along with their 3D models.

Table of Content

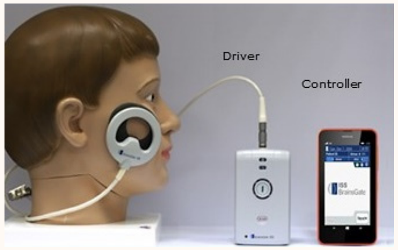

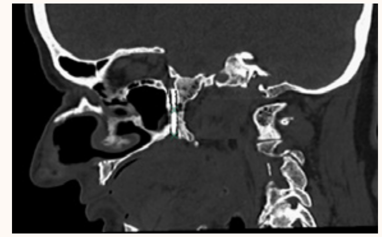

Client Reference

The Intetics team helped us with many directions, some of which we hadn't encountered before. Design, technology selection, implementation, and testing were made fully by Intetics.

Product owner

Benefits and Results

- Only the unique knowledge of the Intetics team helped the Client build the solution with FDA-required parameters.
- The newly created solution enables the Client to certify the whole hardware and software complex and start selling it in US clinics.
- BrainsGate is currently running the ImpACT-24b, a multi-national, pivotal study to assess the safety and efficacy of its treatment for stroke patients in a 24-hour window.
- The software solution was developed and implemented within 12 months of conception.

Team: 7

Project Manager, System Analyst, Software Developers, QA Engineers, Designer

More information on the use case can be found here.

intetics
Where software concepts come alive™

Digital

Machine learning allows you to improve user experience, increase online sales, and spot fraudulent activity in the digital landscape. Here are some common ML applications:

- **Virtual assistants**
  Chatbots, including voice-enabled ones, efficiently communicate with customers 24/7 to provide seamless customer support. Through keyword, text, voice, and image recognition, they can analyze text, images, and speech to maintain communication. Today, they are either integrated as part of customer support or launched as independent solutions. For example, Kindly is a virtual shopping assistant software that guides customers to the items that match their needs and, therefore, boosts online sales.

At the end of 2022, natural language processing and natural language generation tools gained specific attention and rather wide adoption. The interest was provoked by a high-performant AI-based bot, ChatGPT, based on the GPT-3 model. It's the most advanced model for now that can provide natural language texts very similar to human ones. A conversational bot can code, provide all types of text content, and serve as a search engine.

Table of Content

intetics
Where software concepts come alive™

- **Recommendation systems**

  They are often used in e-commerce across numerous sectors, from media distribution to retail, to provide users with tailored content/goods and improve their satisfaction and retention rates. This also reduces churn rates.

  For example, by collecting all possible data about users, from demography to online behavior, Netflix runs its recommendation system to suggest tailored video content. As a result, 75% of watched content is recommendations. The ML algorithm provides personalized recommendations based on user behavior and preferences. A similar ML model is used by Amazon: 35% of all purchases are based on algorithm-powered recommendations.

- **Fraud detection**

  By identifying data anomalies in datasets, ML models can help detect fraudulent activity. This is often applied in FinTech solutions, in which trained models classify an activity as fraudulent or suspicious based on predefined patterns. For example, AWS provides a complete infrastructure to run such a model. Besides this, fraud detection can be enabled by third-party tools like IBM Trusteer that identify new and existing customers across omnichannel customer journeys.

- **Price prediction and dynamic pricing**

  An ML model analyzes timing, demand, demographics, initial and competitive prices, and order sizes to come up with the most accurate price prediction or dynamic price. The latter allows you to dynamically ramp up the prices at peak times or festive seasons, as well as decrease them when the demand is lower. Airbnb offers such a functionality for its users, allowing them to optimize pricing without constantly monitoring it.

Table of Content

intetics
Where software concepts come alive™

Case Study 4

Intetics Created a Machine Learning Algorithm That Recognizes Human Emotions To Improve Wearables for Sports Fans

Business Domain
Media & Publishing

Project Type
Data Processing

Client

The Client is a UK-based wearables development company.

Project

Enriching the product with an emotional-response feature.

Objective

To develop a machine learning algorithm for emotion recognition and find correlations between biosensor data and emotional events.

Team Reinforcement

To successfully implement the feature the Client needed a team of experts to build the machine learning algorithm for emotions recognition. The Client approached Intetics with the request to fulfill the task within a strict time span.

Challenge

The Client is a UK-based wearables development company that was launching a wearable for sports fans.

Table of Content

intetics
Where software concepts come alive™

The Client decided to enrich the product with an emotional-response feature and needed a team of experts to develop a machine learning algorithm for emotion recognition.

The Client carried out research and gathered data from different biometric sensors, worn by fans during live events, and also came up with suggested emotional responses. They delegated Intetics to develop the ML algorithm and the feature.

Quick Facts

- Valence of an emotion and the arousal level are recognized
- The model is edge-computing ready and optimized for wearables
- As a result of Intetics' research, the hardware cost was optimized

Technologies

MATLAB, C# .NET

Solution

★ 01

The Client’s wearable device for sports fans proves that the recognition of emotions from the data of biosensors is possible and works. The consumers can now use the feature to share their emotional reactions to sports events.

★ 02

- Having biosensor data and time-coded reactions to the events, Intetics decided to apply a supervised machine learning approach when creating the algorithm.
- Before developing the algorithm, Intetics implemented a visual tool for biosensor data labeling. Using the correct labels is critical for supervising machine learning issues. In that context, labels corresponded with time intervals in which any intense emotions were expected. The biosensor data of each fan was labeled individually based on XML timecoded feeds, video recordings of games, and fans’ physical reactions. Using the visual tool, an operator was able to quickly and efficiently make the labels.

★ 03

The preprocessing of the data was the first step of the algorithm. This included data filtering and removal of some artifacts. Along with that, each filter introduced a group delay in the data. This means that the output filter data was a bit shifted against the input data (i.e., slightly delayed). Each filter had its own group delay. These delays need to be considered because misaligned time data may cause low results of recognition, even if the rest of the algorithm is perfect. After all of these actions, the output of each sensor was normalized and fit the range from 0 to 1.

Table of Content

intetics
Where software concepts come alive™

★ 04

The second step included extracting features and performing segmentation. Segmentation meant that all data from each sensor, processed at step one, was presented in short pieces of time, usually from 1 to 3 seconds. The feature extraction process included calculating the values for each time segment and sensor.

★ 05

The third step was about dimensionality reduction. After the first and second step, the Intetics team had a large volume of data.

★ 06

The fourth step was related to data splitting for training and testing purposes. That step required specific algorithms. The data was split into two parts. 70% of it was used for training while 30% was for testing.

★ 07

Training was the fifth step. To train the algorithm, the Intetics team used 70% of collected data and relevant labels to teach the algorithm which time segments were emotional and what type of emotion they related. Following this approach, the algorithm remembered segments that were characterized by emotions.

★ 08

The testing of classification was the last step. The trained algorithm used the remaining 30% of data. During the testing process, the algorithm estimated each segment and made decisions.

Often in the real world, algorithm training and all previous relevant steps are performed offline, before uploading firmware into the device. Classification and all relevant steps from the first one are performed in real-time on the device.

★ 09

During this research, Intetics found the most revealing metrics for emotion recognition. This allowed us to optimize the hardware cost by including only the necessary sensors.

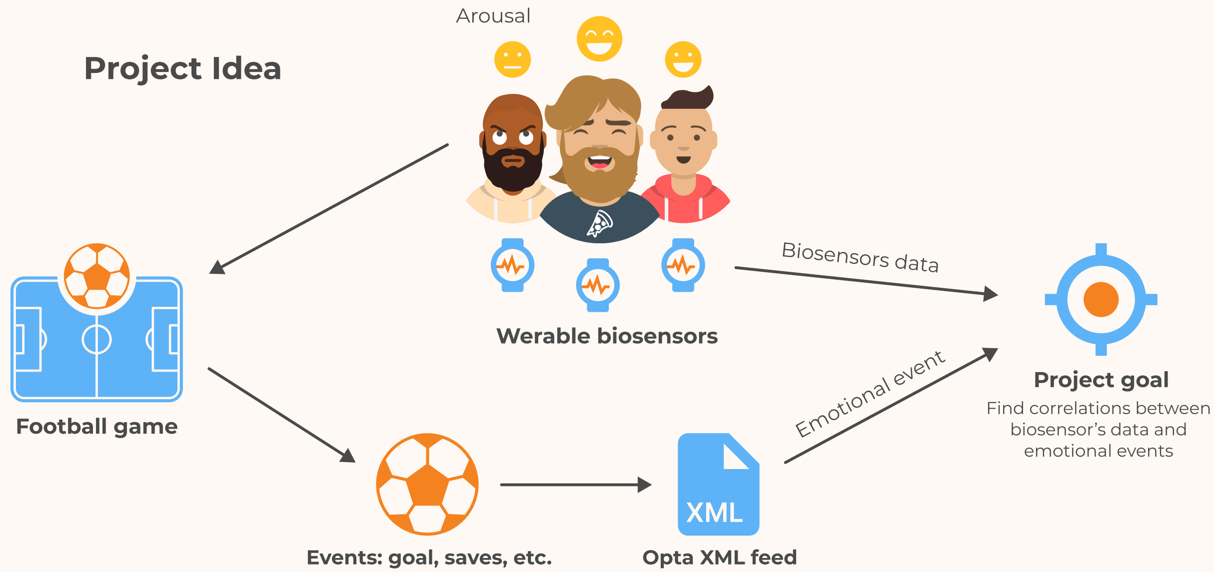

Table of Content

An example of classification results, time domain

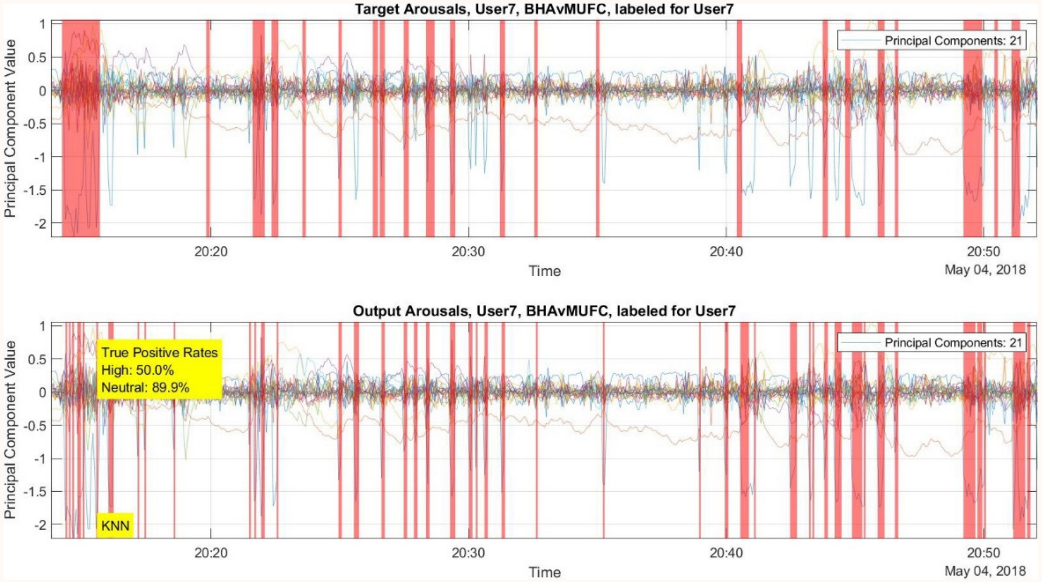

Client Reference

Together with Intetics, we've not only developed a wearable device with the embedded emotion recognition feature but also proved that such features is possible based on biosensor data. Our work may become a significant contribution in the niche and find more application across different niches.

CTO

Benefits and Results

- The Client managed to implement the algorithm that confirmed that the recognition of emotions from the data of biosensors is possible and works.
- Sports fans can now share their emotional reactions to sports events via wearables.
- The algorithm allowed to add the innovative feature to the product and thus boost customer loyalty to the product.
- Intetics delivered the solution in time, regardless of the short timeframe.

Team: 4

2 data scientists, 1 C# .NET developer, 1 data entry specialist

More information on the use case can be found here.

intetics
Where software concepts come alive™

Table of Content

intetics
Where software concepts come alive™

9. Industry Standards for Responsible and Ethical Machine Learning

Standards and best practices are important for ensuring that ML systems are developed, implemented, and used responsibly and ethically. They help to prevent unintended consequences, such as biased or discriminatory models, and promote transparency and accountability in ML development and deployment.

Some of the most important standards in use for ML are:

Open Neural Network Exchange (ONNX)

is an open standard for representing and sharing deep learning models. It enables interoperability between deep learning frameworks, allowing developers to switch between frameworks without having to rebuild their models.

The TensorFlow Model Garden

is a repository of pre-trained TensorFlow models that are available for use. The models in the Model Garden are implemented according to best practices and have been evaluated on various datasets.

Fairness, Accountability, Transparency, and Ethics (FATE)

is an open-source toolkit for secure and privacy-preserving ML. It provides a suite of tools for data processing, federated learning, and model training and evaluation. FATE is designed to ensure that ML models are fair, accountable, transparent, and ethical.

ISO/IEC JTC 1/SC 42

is a subcommittee of the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) that is responsible for developing standards for AI and ML.

It's also important to consider the Global Data Privacy Regulation (GDPR), which applies to all organizations that process the personal data of EU citizens, regardless of where the organization is located. The GDPR is an important standard for ML because it requires organizations to ensure that their ML models do not infringe on the privacy rights of individuals. In the USA, a similar role is fulfilled by California Consumer Privacy Act (CCPA).

In addition to industry standards, there are also best practices for machine learning (ML) engineering that are developed and promoted by leading companies in the field. For example, Google has published a set of best practices for ML engineering, which cover rules such as "Choose a simple, observable, and attributable metric for your first objective," "Reuse code between your training pipeline and your serving pipeline whenever possible," and more.

10. Industry Resources and Services for ML Practitioners and Professionals

For organizations looking to explore the potential of ML models, there are several resources available that can provide guidance and support in this area:

Future Tools

An aggregator of best AI tools designed for different purposes, including those for finance, content generation, translation, fun, self-improvement, productivity, inspiration, marketing, and more. The service offers 1440+ AI-based tools.

Top AI Tools

The service allows you to easily navigate top-performing AI-based tools to facilitate the search of the right one. You can easily submit your own or subscribe for a weekly digest to receive the handpicked top tools regularly.

Table of Content

intetics
Where software concepts come alive™

Data Science Central

This is an online community for data science professionals that features articles, tutorials, and discussions on ML. The portal also hosts webinars and online events for experts and practitioners to share their knowledge and insights.

Google AI

Google AI regularly publishes research papers, provides a range of tools and resources for ML professionals, such as pre-trained models and APIs, and offers the Google AI Residency Program. Google AI researchers are also responsible for the development of TensorFlow, one of the most widely used frameworks in the world.

Microsoft Research

Microsoft Research is known for its work in deep learning, reinforcement learning, and natural language processing. Its researchers have developed the Microsoft Cognitive Toolkit, a powerful open-source framework with extensive documentation and tutorials to help developers get started with using it.

The Association for Computing Machinery (ACM)

Founded in 1947, the ACM is the world's largest educational and scientific computing society, with over 100,000 members from around the world that provides resources and research on ML and related topics.

IEEE Computer Society

The organization offers a range of training programs to help members stay up-to-date on the latest developments and best practices in this field. IEEE hosts conferences and workshops, with the International Conference on Machine Learning and Applications (ICMLA) being one of the notable events.

The AI Podcast

This is a podcast produced by NVIDIA, featuring interviews with experts in the field of AI and ML, as well as updates on the latest developments.

Table of Content

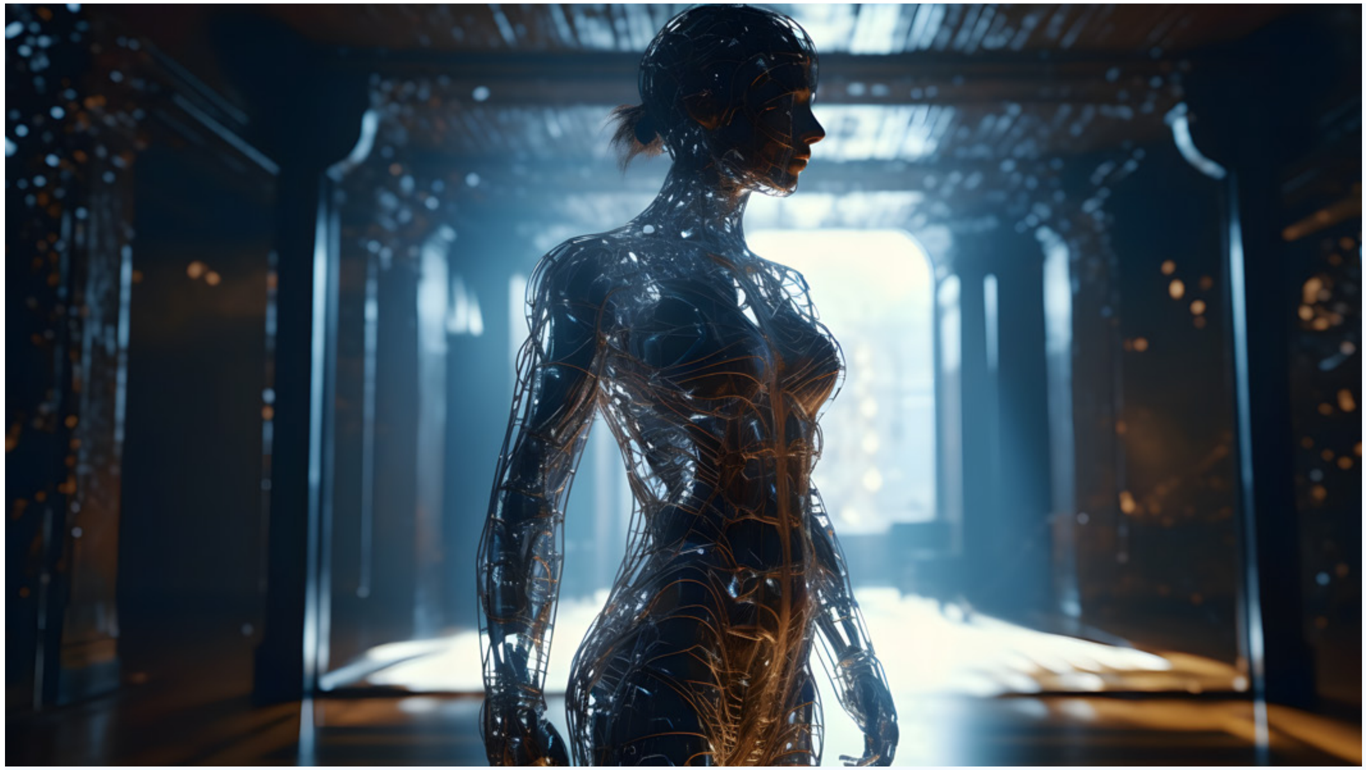

MIT Technology Review

The magazine features news, analysis, and commentary on the latest developments in the field of AI and ML, as well as in other technology-related fields.

KDnuggets Data Science & Machine Learning on LinkedIn

This is a group for machine learning professionals to share information, ask questions, and network with others in the field.

Table of Content

intetics
Where software concepts come alive™

These are cutting-edge tools and services that help users stay ahead of the curve and tackle complex problems:

- **Midjourney** is an emerging text-to-image AI tool accessible through a Discord bot. The laboratory behind it aims to elevate human creative potential, exploring the intersections of design, infrastructure, and AI.
- **Adobe Firefly** is Adobe's family of creative generative AI models that allows users to generate text and images when working across the Adobe ecosystem.
- **Bing Chat** is an AI feature embedded in the Bing search engine that allows users to carry out online search in a conversational manner, similar to ChatGPT.
- **Google Bard** is a conversational chat bot that performs text-based tasks like answering users's questions, providing text summaries, or creating text content.
- **Notion AI** provides invaluable assistance to writers, translators, and anyone working with text through advanced algorithms and natural language processing capabilities. For now, you can join the waiting list.
- **Adobe Podcast** is a free service that helps users enhance their audio recordings to a studio-level standard.
- This **ChatGPT bot** is powered by OpenAI and provides an unrestricted and highly efficient platform for seamless conversations.
- **Smartwriter** is becoming a must-have tool for writing compelling cold emails and LinkedIn introductions, especially in the field of B2B marketing.
- **Grammarly** is an AI-powered writing assistant that provides users with real-time suggestions for improving their writing (grammar, punctuation, spelling, and style).
- **DeepL Translate** is known for its exceptional translation quality and accuracy. This was reportedly achieved through the use of deep neural networks and a large dataset of human-translated texts.

intetics
Where software concepts come alive™

- **Erase.bg** is a game-changing tool for removing image backgrounds. It operates without the need for registration or any technical expertise.
- **Bigjpg** is an online tool for improving the quality and size of images with a variety of quick and efficient solutions.
- **vidIQ** is a powerful analytics platform that helps manage and optimize YouTube channels through competitor analysis, SEO optimization, and other features.
- **LogoAI** offers an extensive database of logo design templates and AI-generated designs.
- **Analisa.io** is a social analytics tool that can analyze any public profile or hashtag on Instagram and TikTok, offering a wealth of data and insights to help you optimize your social media presence.
- **Smartly.io** is a cross-platform management tool that streamlines advertising channels for businesses. However, it's worth noting that its prices are rather high.
- **Lumen5** is an all-in-one tool with valuable resources and templates for creating media content, including video and blog post generation.
- **Lexica** is a comprehensive platform for generating graphics. It has over 5 million AI-generated images and an extensive database of design resources.
- **Runway** is an impressive AI montage tool that can help you remove unwanted objects from videos, apply motion tracking, collaborate with others, and more.
- **Texts** combines all messengers into one easy-to-use platform (iMessage, WhatsApp, Telegram, Signal, Messenger, Twitter, Instagram, LinkedIn, Reddit, Google Chat, Slack, and Discord DMs).
- **Ad Creative** is a crucial asset for businesses seeking to manage advertising channels. It offers an expansive collection of creatives, collaboration tools, and analytics for optimizing ad strategies.

Table of Content

intetics
Where software concepts come alive™

- **Vidyo** is a video generation platform that specializes in creating content for TikTok, Instagram Reels, and YouTube Shorts.
- With **Tome**, the user simply needs to enter a prompt, and the tool will generate an entire narrative. Additionally, users can utilize Tome's DALL·E 2 tile for bespoke images.
- **Creatosaurus** is a collaboration platform for social media content creation that provides post planners, AI writers, and a large gallery of finalized content.
- **Beautiful.ai** is a sleek, intuitive presentation tool that features a unique voice comment feature and a design bot.
- **Otter.ai** transcribes audio and video conversations in real time. It’s a game-changing tool for anyone who struggles with the influx of information during virtual meetings.
- **Rewind.ai** is a productivity tool that is exclusive to Mac users. It uses advanced AI to understand user behavior and help users find the information they need quickly and efficiently.

With the rise of AI-generated content, you may want to discover if the content was produced by humans or AI in order not to face copyright infringement penalties, as well as ensure the content is authentic, secure, and ethical. Consider using AI detectors like OpenAI's AI Classifier, GPT-2 Output Detector, and GPTZero.

Table of Content

intetics
Where software concepts come alive™

11. Authorities and Regulators in the Field

Here are some of the main figures and organizations that have made significant contributions to the field and organizations that are focused on promoting ethical and responsible AI development.

Google Brain

Google Brain is a research team that is focused on advancing the field of machine learning and artificial intelligence. The team is known for developing a range of innovative ML applications and has also published numerous influential papers, including the famous AlphaGo paper.

Meta AI Research

The Meta AI team is a group of researchers and engineers at Facebook. They have published numerous influential papers in top-tier conferences, such as NeurIPS, ICML, and CVPR, and have made important contributions to open-source projects like PyTorch.

DeepMind

The DeepMind team is a group of researchers and engineers at DeepMind Technologies, a London-based AI and ML research lab. They frequently publish their research findings in top academic conferences and journals and make their code and datasets publicly available, enabling other researchers to build on their work.

Open AI

Through their work on TensorFlow and PyTorch, OpenAI has helped democratize access to ML technology and make it more accessible to developers and researchers. Additionally, the company has made contributions to developing reinforcement learning, a key area of ML research. OpenAI has also been vocal in advocating for the responsible use of AI.

Table of Content

Partnership on AI

The Partnership on AI is a non-profit organization that is focused on advancing responsible AI and ML development and promoting best practices in the field. The organization is backed by a number of major tech companies, including Amazon, Google, and Microsoft.

AI Now Institute

The AI Now Institute is a research organization that is focused on studying the social implications of AI and ML and developing best practices for responsible development. The organization is affiliated with New York University.

Olivio Sarikas

Olivio Sarikas is an influencer specializing in AI art tutorials that cover a wide range of techniques, including Midjourney, Stable Diffusion, Automatic 1111, InvokeAI, and more. Olivio shares his secrets during live streams and tutorials, helping artists of all skill levels discover the world of AI and sharing how it’s transforming the art landscape.

Matt Wolfe

Matt Wolfe is a well-known content creator who shares his expertise on various digital topics, including AI, No-Code, Tech, Futurism, Digital Marketing, and Productivity. His videos focus on digital life hacks and cutting-edge technologies, with tutorials, news, and long-form essays.

Paul Roetzer

Paul Roetzer is a Founder and CEO at Marketing AI Institute, as well as a co-author of Marketing Artificial Intelligence: AI, Marketing and the Future of Business. In his LinkedIn profile, Mr. Roetzer regulalrly publishes recent industry news, as well as his opinion on what impact the advancements may provoke.

Table of Content

intetics
Where software concepts come alive™

12. Machine Learning Certifications

Machine learning certifications are becoming increasingly important for practitioners and organizations alike. Here are some of the available certifications:

Certified Machine Learning Engineer (CMLE)

This certification from Google Cloud validates the skills and knowledge required to design, build, and deploy machine learning models. The certification exam covers topics such as data preparation and feature engineering, model training and tuning, and deploying and monitoring models.

Microsoft Certified: Azure AI Engineer Associate

This one demonstrates the ability to design and implement AI solutions using Microsoft Azure tools and technologies. The exam covers topics such as machine learning, natural language processing, and computer vision.

Table of Content

intetics

Where software concepts come alive™

AWS Certified Machine Learning

This certification from Amazon Web Services validates the ability to design, implement, deploy, and maintain machine learning solutions using AWS services. The exam includes data engineering, feature engineering, model selection and training, and deployment and monitoring.

Databricks Certified Associate Developer for Apache Spark

It validates the skills required to develop machine learning models using Apache Spark. The certification exam covers topics such as data preprocessing, model selection and training, and model deployment and monitoring.

IBM Data Science Professional Certificate

This is a comprehensive program that covers the essential concepts and techniques of data science, including machine learning. It requires completing a series of courses and hands-on projects on topics such as Python programming, data visualization, and machine learning with scikit-learn.

Certified Machine Learning Professional

It is offered by the International Association of Business Analytics Certifications (IABAC) and is designed to test the candidate's ability to apply concepts and techniques to real-world scenarios.

For organizations, there are also certifications available to demonstrate their commitment to ethical and responsible machine learning practices. One example is the Machine Learning Ethics Certification from the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. This certification program is designed to help organizations demonstrate their commitment to ethical machine learning practices and ensure that their systems are transparent, explainable, and inclusive.

intetics
Where software concepts come alive™

13. Assessing an Organization's Readiness for ML Implementation

It is important to assess the organization's readiness and need for this technology before implementing it. Here are some key questions to ask during a health check:

Do you really need ML?

As mentioned, organizations need to determine whether ML is the best solution for a particular problem or task. Some tasks may be better solved algorithmically, and implementing ML without a clear need can be a waste of resources.

Do you understand the potential impact of introducing ML?

Managing expectations is crucial when implementing ML. Clearly articulate what the model will do and how to evaluate its effectiveness.

Do you have access to sufficient data?

ML models require a lot of data for training and validation. If the necessary data is not available, it may not be possible to implement ML effectively.

Do you have access to data labeling experts?

Some ML tasks, such as image or speech recognition in specific areas, such as medicine, require expert labeling of the training data. If this expertise is not available in-house, it may be necessary to outsource or hire additional staff.

Do you have the resources to train models?

Training ML models can require significant computational resources. Consider the cost of these resources and whether they are available within the organization.

Do you have the resources to maintain the model in production?

Once a model is in production, it requires ongoing maintenance, including data collection, processing, and retraining. Consider the cost and availability of these resources as well.

14. Additional Resources

If you're looking to delve deeper into the world of machine learning, there are a plethora of resources available that can help you gain valuable insights and knowledge. Here are some recommendations:

Books

- "Deep Learning (Adaptive Computation and ML Series)" by Ian Goodfellow, Yoshua Bengio, and Aaron Courville
- "Pattern Recognition and Machine Learning (Information Science and Statistics)" by Christopher M. Bishop
- "Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series)" by Kevin P. Murphy
- "Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems" by Aurélien Géron
- "Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies" by John D. Kelleher, Brian Mac Namee, and Aoife D’Arcy
- "Programming Collective Intelligence" by Toby Segaran

Guides

Professional ML Engineer Exam Guide

The guide covers a range of topics relevant to building, deploying, and managing machine learning models on the Google Cloud Platform.

People + AI Guidebook

It contains a wealth of information, including methods, best practices, and real-world examples of how to incorporate AI into your designs effectively. Over a hundred Googlers, industry professionals, and academic researchers contributed their data and knowledge to this guide.

Table of Content

Websites

- **TensorFlow**: tutorials, guides, and documentation for using TensorFlow to build ML models
- **PyTorch**: tutorials, documentation, and a community forum for users to get help and share knowledge
- **Kaggle**: competitions, datasets, and forums for data scientists and machine learning practitioners
- **GitHub**: many repositories of open-source ML projects and libraries

15. Miscellaneous

While machine learning is a serious field with important applications, there are a few funny anecdotes and jokes that have emerged around it.

One anecdote involves a team of researchers who used an ML algorithm to generate new names for paint colors. They fed the algorithm a list of existing paint color names and asked it to generate new names based on patterns it detected. The results were a mix of hilarious and bizarre names, such as Electric Lavender and Toxic Bubblegum.

16. Summary and Conclusions

Machine learning is a technology that can drive impact and innovation across all possible industries. By analyzing large piles of data, it can come up with relevant and accurate predictions that may facilitate decision-making, enrich your product with valuable insights, or help you provide a more personalized user experience.

Recent advancements in technology, like transfer learning, are likely to boost its adoption and deployment, as they will allow you to deploy ML models faster and reuse them in similar scenarios. This reduces the training and development time, therefore, optimizes costs and allows for quicker returns.

intetics
Where software concepts come alive™

INTETICS MEANS YOUR SUCCESS

Toll Free: +1 (877) SOFTDEV
US: +1 (239) 217-4907
DE: +49 (211) 3878-9350
UK: +44 (20) 3514-1416

www.intetics.com