

AI in a Nutshell

EVERYTHING YOU NEED TO KNOW ABOUT AI

Powered by:

 **travel.ai**



March 2024

CONTENTS



Introduction	03
Artificial intelligence	04
Big Data	12
Neural Networks	17
Machine Learning	22
Learning to Rank	28
Natural Language Processing	34
Computer Vision	44
Glossary	55





Introduction

Artificial Intelligence (AI) is transforming our world. To some, it may still evoke images of science fiction movies. However, AI has been part of our lives for quite some time. From recommendation systems that know us better than we do ourselves to facial recognition features on our phones, almost everyone uses this technology daily. Having at least a basic understanding of AI is essential. But how can one achieve this when the internet is awash with content filled with technical jargon that is unclear to beginners?

As a company with years of experience in this technology, we've decided to help everyone who wants to dive into the secrets of AI. To achieve this, we have created an ebook designed to explain this powerful technology in an accessible manner.

This publication not only elucidates complex concepts in an easy-to-understand way but also provides specific examples of how companies globally leverage AI to enhance their products and services for customers.


Our ebook is divided into two parts: The first focuses on offering the most accurate explanations of the functions, operations, and applications of various AI areas in life and business. In the second one, you will find a glossary encompassing all the most important terms related to AI.

We wish you an enjoyable reading experience!

What is Artificial Intelligence



Artificial Intelligence (AI) is a field that combines computer systems and datasets to simulate human intelligence. It involves mimicking tasks previously performed by humans and performing complex tasks more human-like than ever before. Artificial intelligence is based on modeling human thinking and cognitive processes. It can help solve problems, make decisions, and perform tasks that require speed, accuracy, and complex thinking. It includes, among other things, planning, problem-solving, language understanding, voice and image recognition, learning, and any other task considered “intelligent.”



Examples of using artificial intelligence in everyday life



1

Voice assistants

Programs like Siri, Alexa, or Google Assistant use AI to recognize speech and provide answers to questions asked by users.

2

Filters in social media

Apps like Instagram or Snapchat use AI algorithms to recognize faces and automatically apply real-time effects.

3

Shopping recommendations

Online stores, such as Amazon or Aliexpress, use AI algorithms to analyze customer behavior and offer personalized offers.

4

Personalized travel recommendations

Travel companies and websites use AI algorithms to analyze user preferences and, based on that, offer them travel packages that best match their interests and needs.

Examples of using artificial intelligence in everyday life



5

Product and content recommendations

Services like Netflix, YouTube, or Spotify use AI-based recommendation systems to offer users content tailored to their preferences based on their previous activity.

6

Translators

Automatic language translation systems such as Google Translate or DeepL use Natural Language Processing (NLP), a subset of AI, to translate text in real time.

7

Chatbots

Computer programs that automatically answer user questions and help solve problems, for example, in online banking or e-commerce, are another popular solution based on NLP.

8

Intelligent search engines

Some tools for searching information on the web or specific websites also use AI technologies, such as natural language processing and machine learning. Thanks to this, they can better understand and interpret user queries to provide them with more relevant and personalized results.

Benefits of using artificial intelligence



Improving business efficiency

According to Accenture, AI can increase business productivity by up to 40%. By automating time-consuming and repetitive tasks such as invoice generation, data processing, or customer service, companies can dedicate this saved time to other activities that lead to faster company growth.

Increasing customer satisfaction

AI enables the creation of more personalized and tailored products and services. An example is Spotify, which uses AI to recommend music tailored to the listener's taste.

Improving customer service

In today's world, customers expect immediate answers to questions and customer service available 24/7. To meet these requirements, more and more companies use AI-based chatbots. These intelligent computer programs can work continuously, providing customers with answers to frequently asked questions and even performing simple tasks such as sending invoices or booking accommodation.

Better data analysis

AI enables faster and more accurate analysis of large amounts of data, enabling a better understanding of the market and customers and making more accurate, data-driven business decisions.

Increasing security

AI can help detect and prevent fraud and other irregularities. An example is PayPal, which has used AI for years to detect suspicious transactions and fraud.

Improving medical processes

AI has many medical applications, such as therapy planning, improving clinical trials, or analyzing medical images. It can also help in the early diagnosis of many diseases, such as cancer.

AI Applications in Tourism



Chatbots

These virtual assistants play a crucial role in fast and efficient customer service. Companies that offer chatbots on their websites or apps can provide their customers with 24/7 support, which is extremely important today. Such automated assistants can help with site navigation, finding important information, performing simple tasks like canceling or changing reservations, generating invoices, and much more.

Humanoid Robots

Chatbots are a relatively common and inexpensive way to support customer service, but what about robots that can serve customers in places like hotels or airports? Such solutions are also gaining popularity, as they help avoid long queues at information desks or reception areas, thereby increasing customer satisfaction. With this technology, guests receive fast and efficient assistance, meeting their needs even during their ongoing journey.

AI Applications in Tourism



AI-Based Recommendations

AI-based recommendations have become essential to the tourism industry, allowing for better personalization of customer offers. Artificial intelligence can analyze vast amounts of data, such as preferences, interests, and demographic information, to provide perfectly matched suggestions.

A prime example that utilizes this type of solution is [TripAdvisor](#). Their system analyzes vast data from reviews, photos, and user comments to help travelers choose suitable tourist attractions, restaurants, or accommodations. Thanks to artificial intelligence, TripAdvisor can present the most accurate suggestions tailored to the individual needs of a specific user.

Using AI to develop recommendations also has advantages for tourism businesses, such as hotels, restaurants, and travel agencies. It allows them to tailor their offerings to meet customer requirements, increasing the chances of attracting interest in their services and building customer loyalty.

AI Applications in Tourism



Intelligent Search Engines

Nowadays, people are inundated with vast information from various sources. That's why contemporary customers expect a personalized approach to help them find offers tailored to their preferences.

The elements that undoubtedly require improvement on most travel websites are the search engines. Traditional search methods in tourism are inadequate for today's travelers' demands, especially when considering the expectations of younger generations.

Fortunately, AI also comes to the rescue here – specifically, intelligent search engines. With these advanced tools, travel companies can quickly and accurately provide customers relevant offers and information tailored to their needs. It leads to increased conversion rates, improved customer satisfaction, and enhanced efficiency of the entire booking process.

AI Applications in Tourism



An example of such a solution is Qtravel Search, developed by **Qtravel.ai**. This intelligent search engine analyzes vast amounts of data, considering customer preferences, such as location, duration of stay, budget, and additional amenities, to propose the best offers meeting specific requirements.

Moreover, the tool features a function that remembers users' choices, making the search process more precise and intuitive with each subsequent use. It benefits the current user and the next ones who utilize the search engine.

The screenshot displays the Qtravel.ai web interface. At the top, the logo 'Qtravel.ai' is on the left, and navigation buttons for 'PLAYGROUND', 'DASHBOARD', and 'DEMO_PL' are on the right. Below the navigation bar, there are four tabs labeled 'PRZYKŁAD 1', 'PRZYKŁAD 2', 'PRZYKŁAD 3', and 'PRZYKŁAD 4'. The main search area contains a text input field with the query: 'hotel z infrastrukturą dla niepełnosprawnych, jacuzzi i centrum fitness wylot w czerwcu lub wrześniu z Warszawy dla pary'. To the right of the input field is a magnifying glass icon. Below the search bar, a message states: 'Kluczowe informacje wykryte w Twoim zapytaniu dzięki narzędziom rozumienia języka naturalnego:'. Underneath this message, several tags are displayed: 'Warszawa' (with a location pin icon), '2024-06-01 → 2024-06-30' (with a calendar icon), '2024-09-01 → 2024-09-30' (with a calendar icon), and 'liczba osób dorosłych - 2' (with a person icon). Below these tags, there are two more tags: 'transport - samolot' (with an airplane icon) and 'zakwaterowanie - hotel' (with a bed icon). At the bottom, there are three dark blue buttons with white text: 'atrakcje w hotelu - centrum fitness', 'udogodnienia w hotelu - udogodnienia dla niepełnosprawnych', and 'udogodnienia w pokoju - jacuzzi'.

What is Big Data?



As the name suggests, the term Big Data refers to massive, complex data sets that are difficult or impossible to process using traditional methods due to their size, speed of generation, and diversity.

3V model

One of the more popular terms for Big Data is the definition created by Doug Laney, an industry analyst who, in the early 2000s, defined Big Data using the so-called “3V model.”

Volume (Volume): Technological advances have made it possible to collect and store data on a previously unattainable scale. The first V thus refers to volume as understood by huge volumes of data.

Speed (Velocity): In the age of the Internet of Things and instant communication, data is generated and transmitted incredibly fast. The second V, therefore, stands for the lightning-fast speed at which data must be processed and analyzed.

Variety (Variety): Refers to the variety of sources and forms from which data is collected. Data can include numbers, text, audio, video, graphics, and more. Each data type requires specific tools to collect, store, and analyze it.

How Does Big Data Work?



To achieve the benefits of Big Data, you need to prepare it properly. Working with big data involves several critical steps, each contributing to achieving this outcome. These steps may vary based on the specifics of the data and the practices of the company processing it. Below you will find the most common activities related to Big Data.



Jak działa Big Data?



1. Data Collection

The first step is the generation and/or collection of data. This process varies from organization to organization, as each has unique methods for data acquisition. It includes collecting unstructured (refer to glossary) and structured data from various sources, such as the cloud, mobile apps, IoT sensors, and more.

2. Data Organization

As mentioned above, Big Data encompasses structured and unstructured data. While the former is typically easier to organize, the latter must be prepared appropriately to yield correct information.

3. Data Cleaning

To obtain reliable results and enhance the quality of the data, it needs to be cleaned. This process involves removing duplicates and unnecessary data that could lead to distorted results and incorrect conclusions.

3. Data Analysis

Data on its own holds no value until it is analyzed and valid conclusions are drawn. To do this, you need to use advanced analytical techniques such as data mining, predictive analytics, and deep learning.

Key Benefits of Big Data



Customer Acquisition and Retention

Big Data allows companies to understand customer preferences, needs, and purchasing behaviors deeply. A notable example is Amazon, which uses big data to personalize shopping experiences and suggest products based on customers' past purchases. One look at any product recommendations when visiting the platform should illustrate this point perfectly 😊

Like Amazon, travel companies can use Big Data to personalize their services, recommending attractions or accommodations to customers based on their past choices or the preferences of others with similar interests.

Personalized Offers and Higher ROI

Analyzing Big Data enables companies to target specific groups with personalized offers, leading to more effective promotional campaigns and a higher ROI (Return on Investment).

Key Benefits of Big Data



Forecasting Trends

By analyzing large data sets, companies can forecast future trends, enabling them to better prepare for changes and stay ahead of the competition.

Identifying Potential Risks

Big Data is an excellent tool for risk management – regardless of the industry. In the case of tourism, it could involve monitoring weather conditions, local events, airline worker strikes, or other factors impacting travel.

Dynamic Pricing

Big data analytics is perfect for adjusting prices based on demand, availability, and other factors such as seasonality or weather. It can help optimize pricing strategies and increase profits.

What are Neural Networks?



Neural networks are another popular concept we have decided to address in our [articles on artificial intelligence terminology](#). In this post, you will find a comprehensive definition and learn about their history, how they function, and their various applications. We invite you to read on!

Neural networks (NN), or artificial neural networks (ANN), are loosely connected models with flexible structures and large parameter space.

Thanks to the inspiration from the biological mechanisms of brain function, neural networks can recognize patterns in data and solve complex tasks related to language understanding and computer vision, among other things.

3 types of neural network layers



Therefore, neural networks are composed of layers, of which we can distinguish three types:

Input layer

This is where the data is directed, which the network will process.

Hidden layers

These are located between the input and output layers. They are called “hidden” because their inputs and outputs are not directly visible.

In these layers, the model gains the ability to recognize patterns in the data, updating its parameters (weights) during training based on the input data.

Output layer

This layer returns the final result of the network's operation. It could be, for example, recognized word categories in text or objects detected in an image.

How do neural networks work?



Neural networks are inspired by their biological counterpart – the human brain. They consist of many closely interrelated nodes – processing units responsible for a different aspect. Each processing node contains its database, including all the previous rules it has been trained on.

Creating a neural network starts with the first nodes receiving raw input data. Then, each node receives the results from the previous one. It continues until the last node processes the information and generates the output data, the final result.

For this article, let's imagine that each node is a robot that analyzes certain information. Each robot that makes up the neural network receives information from the outside (input data) and focuses on a specific aspect of that information. For example, one might analyze shapes, and the second one – colors.

When a robot receives information, it weighs it according to its own “importance” scale, and some things may be more important to it than others.

Next, it checks if the sum of this information exceeds a certain threshold. If so, the robot activates and sends a signal to the next robot in its network. If not, the robot does not activate or send a signal.

During this process, all information passes through the network, from one robot to the next, until we finally get a result – for example, about what is in a given image.

It's worth adding that the mentioned weights and thresholds are initially set randomly, which means that the network initially operates with low effectiveness. However, by processing an increasing amount of data, the network adjusts its parameters, constantly increasing its effectiveness and precision.

Examples of Neural Networks Applications



1

Tourism

Neural networks learn from historical data, identifying specific patterns and relationships – for example, during which months there is greater interest in selected travel destinations. Then, they can, among other things, predict demand for specific travel offers, allowing travel companies to plan their activities more effectively.

2

Search Engines

Neural networks are also employed [in search engines](#) that analyze user queries and deliver relevant results. By learning from billions of queries, they can understand the context and nuances of language, providing answers to even the most complex queries.

3

Social Media

Neural networks analyze large amounts of data generated by users, which allows for a deeper understanding of their behaviors. This information can be utilized in various ways, including creating personalized advertisements.

Examples of Neural Networks Applications



4

Aircraft

Neural networks can be used for error diagnostics during aircraft production, crew training, and even in autopilot systems.

5

Meteorology

Neural networks can analyze data such as air humidity, temperature, solar radiation, and wind speed to provide the most accurate weather forecasts.

6

Defense

Neural networks are used for object location, analyzing armed attacks, controlling autonomous drones, and managing air and sea patrols.

7

Healthcare

Neural networks can accurately analyze images, such as X-rays or computer tomographies. Moreover, they're also used in research for discovering new medications.

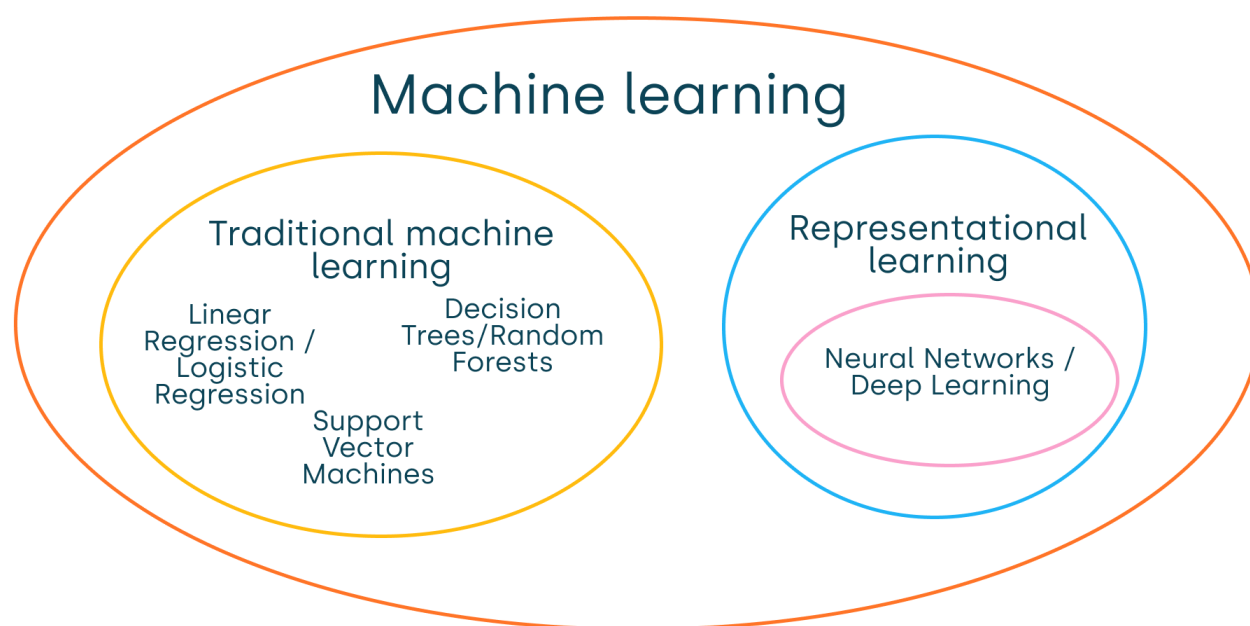
What is Machine Learning?




Machine Learning is a subset of Artificial Intelligence. As the name implies, it focuses on enabling machines—primarily computers—to learn specific tasks in a manner reminiscent of how humans learn.

This technology imparts knowledge to computers through data, observations, and interactions with their environment. Subsequently, they identify distinct patterns and insights. To achieve this, machine learning systems employ various algorithms, such as [neural networks](#) and clustering (check out: Glossary).

Crucially, machine learning algorithms hinge on a computer's capability to discern pertinent patterns and correlations within the data automatically. They then make optimal decisions based on these insights.



How machine learning work?



Input Data: The journey begins with data collection, encompassing various forms like photos, texts, sounds, and numerical figures.

Data Preparation: Once gathered, the data undergoes a cleaning process to ensure its accuracy. This phase involves rectifying inaccuracies, addressing errors, and filling in missing values. While this step is optional, it's highly recommended to achieve superior outcomes.

Data Analysis: At this juncture, the machine learning model delves into understanding the data's structure and discerning relationships between its various features.

Pattern Search: Leveraging insights from the analysis, the model unearths hidden patterns and connections within the data.

Prediction: Drawing on these identified patterns, the model is primed to forecast outcomes for new, unseen data.

Decision-making: Informed by the predictions, the system arrives at specific conclusions and takes corresponding actions. An example would be the system classifying an email as spam or legitimate; depending on this classification, it'll then allocate the email to the relevant folder.

Machine Learning Techniques



Supervised Learning

Supervised learning is a technique where a model learns from input data and its corresponding output, provided with known labels (usually given by a human). Take, for instance, a scenario where you aim to train a system to distinguish between an apple and a strawberry. Your training dataset would comprise images of the fruits (the input) and their respective identifications as an apple or a strawberry (the label). The system then processes this data, grasping the connection between the input and the output. As a result, it's equipped to identify these fruits in future instances independently.

Common supervised learning applications encompass tasks like spam email detection and the fruit identification example mentioned earlier.

Within supervised learning lies regression, a technique aimed at predicting numerical values. It seeks to establish a relationship between one or multiple independent variables and a dependent variable.

For example, regression might be used to anticipate a company's stock price or the selling price of items.

Machine Learning Techniques



Unsupervised Learning

Unsupervised learning involves a model analyzing a set of input data without the aid of known labels.

Contrary to supervised learning, the algorithm isn't equipped with predetermined answers here. It seeks to unearth hidden patterns and structures within the data and discern correlations.

Drawing from the fruit analogy, if the system receives a series of fruit images without any specific labeling, it can autonomously categorize them based on inherent similarities. However, it won't inherently discern that one cluster represents apples and another signifies strawberries.

The core objective of unsupervised learning is to delve into data, unveiling latent patterns or structures. While it doesn't necessitate "supervision" in the conventional understanding, the results might still warrant a human touch for interpretation.

Applications of unsupervised learning span areas like customer segment clustering, anomaly detection, and topic clustering for articles.

Machine Learning Techniques



Semi-supervised Learning

Semi-supervised learning strikes a balance between supervised and unsupervised techniques. It leverages a mix of a limited amount of labeled data and a substantial volume of unlabeled data. The primary objective is to harness the labeled samples to interpret better the structure inherent in the larger, unlabeled dataset.

A significant advantage of this method is its ability to produce accurate models while drastically reducing the expenses and effort of data labeling. This positions semi-supervised learning as a valuable tool in scenarios where manual annotation proves costly or time-intensive, yet unlabeled data are abundant.

Practical applications of semi-supervised learning span domains like speech analysis, natural language processing, and fraud detection.

Machine Learning Techniques



Reinforcement Learning

Reinforcement learning is a technique where a machine learns through interaction with its environment. Distinct from supervised learning, which provides the machine with explicit responses, reinforcement learning hinges on the principle of trial and error instead of relying on a specific dataset for training.

Within this paradigm, an agent (model or algorithm) carries out actions in a given environment, receiving rewards or penalties in return. The process centers on discerning which actions yield the highest reward under specific environmental conditions. Central to this approach is the “reward” – a quantitative measure of the system's success in achieving its objectives.

Consider the challenge of teaching a robot to walk. Instead of supplying step-by-step instructions for each limb's movement, the robot receives positive feedback or “rewards” for correct actions and “penalties” or negative feedback when it stumbles. . The robot refines its movements by consistently navigating this feedback to garner optimal outcomes.

Another illustrative example is mastering the game of chess. . Here, the model's actions equate to the moves it makes, the environment corresponds to the game board and the rival's plays, and the eventual reward hinges on whether the game is won or lost.

Machine Learning Techniques

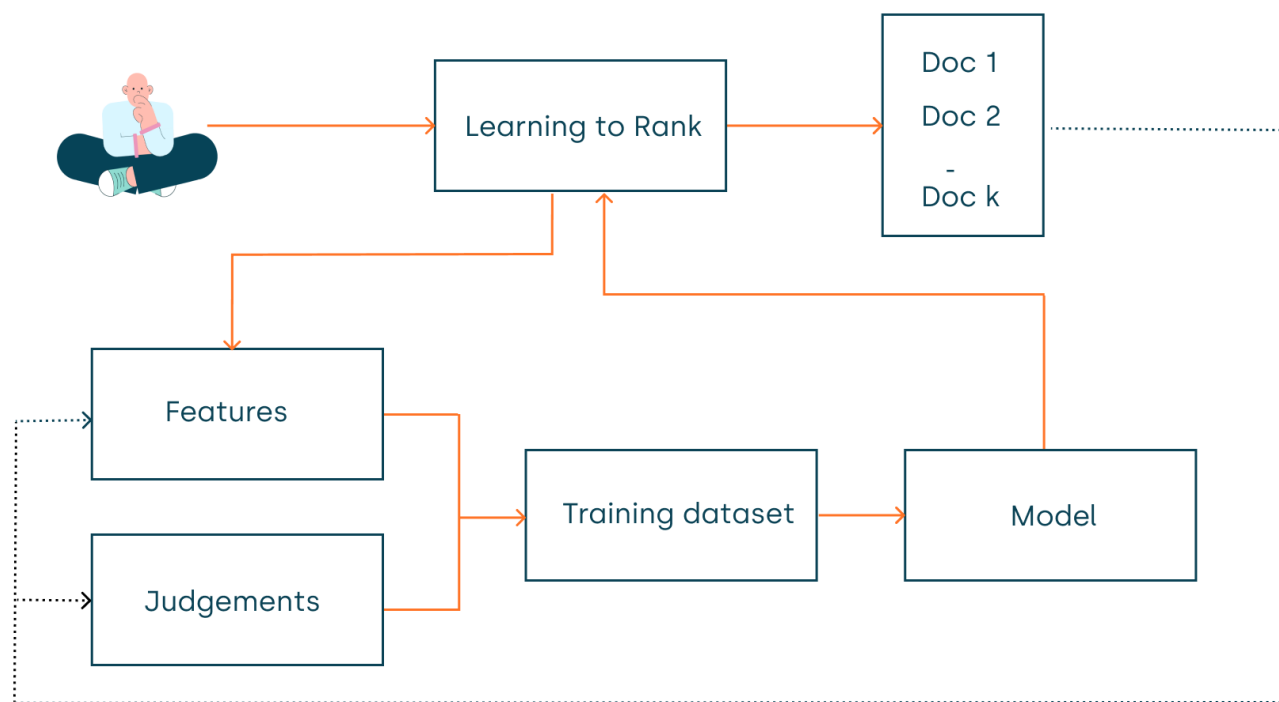


Reinforcement learning finds applications in diverse domains. Examples range from autonomous vehicles navigating impediments to trading algorithms making market moves to enhance gains or curtail losses. Additionally, games such as chess, Go, and Starcraft have been pivotal in advancing reinforcement learning. Interestingly, developing virtual assistants like ChatGPT also leans on reinforcement learning, among other methodologies.

What is Learning to Rank?


The field of “Learning to Rank” (LTR) is a type of machine learning that specializes in developing techniques and algorithms for sorting and ranking data sets in the appropriate order. It is typically applied in applications related to natural language processing (NLP), information retrieval, and recommendations.

The primary purpose of Learning to Rank is to teach a model how to arrange a set of documents or items in a manner that is consistent with the user's intentions. In a ranking recommender system, there is a catalog of items. When a user interacts with the system, it retrieves relevant items and arranges them in order of relevance.



Learning to Rank in Tourism

In the travel domain, LTR can be implemented to offer individualized recommendations to users by analyzing search queries. For instance, when a user searches for flights from New York to London, LTR can rank the results in order of relevance based on several factors, such as flight duration, price, and departure time. This ensures that the user receives the most suitable results and saves time by not having to sift through irrelevant options.



Example Application of Re-Ranking in the Tourism Industry



To illustrate how the re-ranking process works, let's consider the following scenario: Imagine a search engine system specializing in offering travel trip suggestions to users. The system's primary objective is to present travel itineraries that align with users' preferences and interests. However, it's equally important to balance offering popular, well-visited, and unique travel destinations.

Initially, the system arranges travel trip options based on their popularity. This ranking could consider metrics such as the number of travelers who have explored these trips or user reviews. Consequently, the most widely visited travel destinations rise to the top of the search results. However, some users may prefer discovering lesser-known or niche travel experiences.

To address this, the system integrates a recommender algorithm that factors in a user's personal preferences by interpreting their search queries. When a user initiates a search for travel trips, the system takes into consideration both the initial popularity ranking and the personalized recommendations. It might begin by showcasing a selection of popular travel itineraries at the beginning of the search results. The system can refine the search results and recommendations by considering a specific user's search.

In this way, the system combines the benefits of both popularity and personalization. It ensures that users are exposed to popular travel destinations while adapting their preferences, thereby avoiding the extremes of offering only mainstream travel options.

Types of Learning to Rank Algorithms



LTR algorithms belong to the category of [supervised learning](#), which means they require input and output data. In the context of LTR, the input is typically a user ID, while the output is a list ranked according to relevance to the user's search query. Some of the most commonly used LTR algorithms employ deep neural networks designed to learn and recognize complex patterns in data. Other LTR algorithms use boosting, combining multiple weaker models to create a stronger model, or decision trees, which are used to classify data based on a series of binary decisions. By utilizing these sophisticated algorithms, LTR systems can effectively learn to rank data based on a wide range of factors and improve their accuracy over time.

There are three main types of LTR algorithms: pointwise, pairwise, and listwise. The selection of the algorithm relies on the particular task at hand. In the following sections, we will discuss each type of algorithm in more detail.

What Is the Difference Between Predicting Grades and Ranking?



In machine learning, LTR is a crucial subfield that plays a significant role in search engine applications. LTR is a set of algorithms that are used to train recommenders and other data-driven applications to produce ordered lists.

Unlike rating recommendation systems, LTR doesn't focus on correctly predicting ratings but instead concentrates on creating lists of items ranked appropriately.

The primary goal of LTR is to provide the most relevant results for a particular user query.

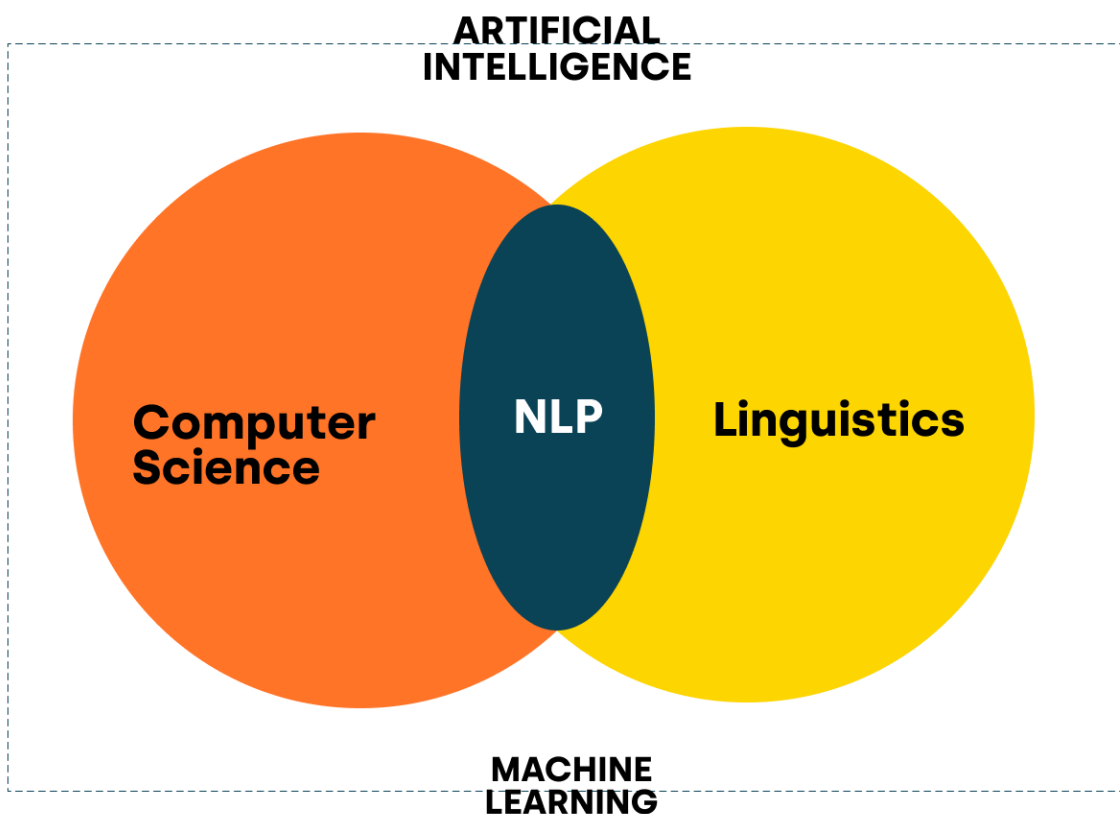
The difference between rating prediction and ranking is that ranking emphasizes the position of an item in the list and not the actual score. In other words, it doesn't matter if an item has a high score or is within a rating scale as long as the score represents the position of the item in the rank.

Jak działa przetwarzanie języka naturalnego?



Natural Language Processing is [a field of artificial intelligence](#) combining programming and linguistics elements, enabling computers to understand human language.

NLP utilizes advanced technologies, including computational linguistics, [machine learning](#), and deep learning models, to process human language in various forms, from text to voice recordings. The main goal of NLP-based solutions is not just to grasp the meaning of individual words but also to comprehend entire speech patterns, considering the speaker's context, intentions, and even emotions.



How does natural language processing work?



1. Data Gathering

The fundamental component of any solution based on machine learning models is data. Specifically for Natural Language Processing (NLP), these systems require a corpus – a vast collection of linguistic data. For instance, the GPT-3 version of a system like ChatGPT [has been trained on approximately 570 GB of data](#). This data includes a diverse range of sources such as books, web texts, Wikipedia, and various articles available on the Internet, amounting to 300 billion words.

2. Converting Human Language into a Format Understandable by Computers

After data collection, the system converts this data into a numerical format, enabling it to comprehend and process the information. The conversion technique varies based on the type of data, whether text or audio. For instance, with voice data, the system analyzes and extracts information about varying frequency components at different moments in the recording.

How does natural language processing work?



3. Understanding the Message

Remember, while words and sentences are understandable to humans, they are merely strings of characters to computers. To enable these systems to comprehend language and the meaning of statements, AI and NLP engineers must train them accordingly.

This is where the so-called NLU (Natural Language Understanding) comes in, involving the processing of input data to make sense of sentences in natural language.

It usually involves various natural language processing techniques, such as tokenization (breaking text into smaller parts), part-of-speech tagging, and stemming (reducing words to their basic forms).

These methods allow algorithms to grasp the rules of human language, leading to an understanding of sentence structure and meaning.

There are significant challenges at this stage notably with homonyms — words that share the exact spelling or pronunciation but have different meanings. The system must discern their context when such words are encountered by analyzing the entire utterance.

How does natural language processing work?



Once our system understands the questions or commands it receives, it generates an appropriate response. It is achieved through deep learning algorithms that enable it not only to read and comprehend statements but also to formulate its own responses. This process, known as Natural Language Generation (NLG), transforms structured data (data understandable to the computer) into natural language outputs that are understandable to humans.

Taking the example of ChatGPT, at this stage the system generates responses based on the user's input. It leverages billions of previously learned words and language rules and has an extensive knowledge base. Practically, this involves predicting subsequent text segments that should logically follow the provided input, guided by established language patterns.

How does natural language processing work?



5. Continuous Training and Fine-Tuning of the Model

You must constantly refine and update to ensure your system works at its best. The more data the system processes, the more accurate it becomes. It is particularly crucial for NLP solutions, where the model needs regular fine-tuning with new data alongside continuous refinement based on the feedback it receives. As a result, the system becomes increasingly precise and effective in understanding and responding to user queries.

For instance, to keep ChatGPT informed about current regulations in a specific location, it is essential to tune the model with new, up-to-date information. This ongoing training enables the AI to provide the most relevant and accurate responses.

Examples of NLP Applications



Voice Assistants

NLP is crucial for Siri, Alexa, and Google Assistant. It enables them to understand and naturally respond to user voice queries, performing tasks ranging from setting alarms to searching for information online.

Translators

Tools like Google Translate and DeepL employ NLP for real-time translation of texts into various languages.

Text Dictation Function

The text dictation feature, available on smartphones and other devices, allows users to dictate text that is subsequently transcribed into written form.

Sentiment Analysis

Companies utilize NLP to analyze customer reviews and opinions, helping them grasp the general sentiment toward their products or services. This insight aids in refining marketing strategies and enhancing customer service.

Examples of NLP Applications



Chatbots

Chatbots on websites and apps use NLP to respond quickly to customer inquiries.

Search Engines

Advanced or semantic search engines apply NLP to understand the intent and context behind search queries. This capability enables them to comprehend complex queries and deliver answers that grasp the whole meaning of the search content rather than merely matching keywords.

Industries Benefiting from Natural Language Processing (NLP)



Tourism

NLP helps identify ideal travel deals and more. For instance, [Qtravel Search](#), a [travel search engine](#), uses NLP to comprehend user queries effectively, matching them with suitable travel deals, hotels, or attractions based on a detailed query analysis.

Healthcare

NLP is crucial in analyzing medical records to expedite diagnoses and personalize patient care. NLP-powered systems can evaluate medical histories, doctor's notes, and medical publications, offering vital insights to healthcare professionals.

Finance

NLP aids in market analysis and fraud detection within the financial sector. It can process extensive financial data, including reports and market news, to provide investment recommendations or identify potential fraudulent trading patterns.

Industries Benefiting from Natural Language Processing (NLP)



Education

NLP is instrumental in creating personalized curricula and language learning tools. It analyzes students' responses to tailor teaching materials to their understanding level and learning style.

Accounting

NLP enhances accounting efficiency by automating data entry and extracting essential details from invoices, receipts, and other financial documents. This automation speeds up processes like invoice posting and reduces the likelihood of errors.

Benefits of Using Natural Language Processing in Business



Time Savings

NLP enables the automation of repetitive tasks like sorting customer service emails or analyzing market data, leading to substantial time savings for your employees.

Enhanced Customer Satisfaction

NLP-powered solutions like chatbots and semantic search engines deliver quick and precise responses, enhancing service quality and boosting customer satisfaction.

Deeper Customer Insights:

NLP's ability to analyze customer feedback and reviews offers more profound insights into customer needs and preferences, allowing for more customized service offerings.

Informed Business Decisions

NLP-driven data analysis yields accurate insights into market trends and customer behavior, aiding in more informed, data-driven decision-making.

Improved Conversions and Sales

In marketing and sales, NLP aids in precise ad targeting and offers personalization, increasing conversion rates and sales.

What is Computer Vision?



Computer Vision is a field of artificial intelligence that, as the name suggests, is tasked with detecting and interpreting visual data, such as digital images or video recordings.



Like humans use their eyes to perceive their surroundings, computer vision employs sensors like cameras or camcorders to identify certain patterns and objects. Then, akin to the human brain, a Computer Vision-based system utilizes various algorithms, such as neural networks — specifically Convolutional Neural Networks (CNN). These networks, inspired by the function of the visual cortex, can identify the content of the processed image.

Thus, software based on Computer Vision can, among other capabilities, distinguish objects, recognize faces, or even read and interpret images.



How does Computer Vision work?

Human Vision VS Computer Vision



Data Acquisition

The initial step involves obtaining data from a suitable device (sensor) to capture what the system is intended to see. For instance, a heritage recognition application could include taking a photo with a smartphone camera.

Pre-processing

The raw data extracted from devices often requires further modification to ensure uniformity or to enhance quality. In the heritage recognition example, users may take images with cameras of varying resolutions. During pre-processing, the dimensions of these images are standardized to a fixed value. Additional techniques include normalization, de-noising, or altering the color scheme.

How does Computer Vision work?



Feature Extraction

The acquired data is then processed by a selected algorithm to extract numerically expressed features from the image, which are essential for performing the target task. A popular approach here is convolutional neural networks. These networks extract features from spatial data, such as images, by capturing local relationships between neighboring elements (e.g. pixels). In successive layers, features of increasing levels of abstraction are extracted from the image. Initially, simple elements like straight/curved lines or textures are detected and assembled to identify more complex shapes and features. For a heritage recognition application, these might include the materials of the building, the shapes of elements (like windows and columns), and the overall structure of the building.

Decision-making

Based on the features extracted in the previous step, specific decisions are made depending on the task. It could involve assigning an appropriate name to the image for a heritage recognition application. Other potential choices might involve determining the type and position of objects in the image or assigning a label to each pixel (semantic segmentation).


Key Techniques of Computer Vision

Classification



fountain

Object Detection



historical buildings, fountain, landmark

Image classification

Refers to the process of categorizing images into different groups. This technology is utilized to identify and classify images based on their content. The algorithm evaluates each image and assigns a probability score for its belonging to various classes. For instance, if a picture contains a cat, the algorithm will likely assign higher probabilities to categories such as 'cat,' 'animal,' or 'pet.'

Object localization

This technique determines the locations of objects in a photo or video without classifying them. Objects are most often localized by specifying the rectangular areas they are situated.

Key Techniques of Computer Vision



Object detection

This technique addresses both 'where' and 'what' objects are in images or videos. It combines localization and classification to detect objects in images or videos and is often used to identify interesting segments for more detailed analysis. It can be useful, for example, in detecting animal species in their natural habitat for research or conservation purposes.

Object tracking

This refers to the process in video footage that involves tracking (associating occurrences of) the same objects across successive frames, such as tracking the movement of vehicles.

Content-based image retrieval

This method involves browsing, searching, and retrieving data from large collections based on the content of the input image. Commonly used for digital asset management and research purposes, A popular application of this method is [Google Lens](#).

Identification

This technique determines the specific instance of a particular object in an image. It is a more detailed classification form focused on identifying a single, unique occurrence of an object within a particular class. For example, identification might involve recognizing a specific person's face or fingerprint for biometric authentication.

Key Techniques of Computer Vision



Semantic segmentation

Involves dividing an image into segments, each representing different object classes. The technique assigns a specific class to each pixel. In autonomous vehicle technology, pixel classes might include roadways, sidewalks, lanes, buildings, etc.

Instance segmentation

This process identifies and separates each individual object in an image. Unlike semantic segmentation, instance segmentation recognizes each unique instance of an object within the same category. For example, it can distinguish individual people in a crowd, even though they all belong to the “people” category. It is useful in applications requiring precise recognition of the exact locations of individual objects, such as vehicles in traffic.

OCR – Optical Character Recognition

OCR enables machines to recognize and interpret text from images or documents. It converts text in a visual form into an editable digital format (such as a text file). OCR can recognize characters from various sources, such as document scans, images of handwritten text, or PDF files, eliminating the need for manual transcription.

Computer Vision Applications in Various Industries



Medicine

In the medical sector, Computer Vision analyzes medical images like MRIs and CT scans. It aids in faster diagnosis and more precise treatment.

Manufacturing

CV has numerous applications in manufacturing, including monitoring and optimizing processes and controlling product quality.

Automotive

The technology is utilized to analyze the road environment, encompassing tasks like obstacle detection and sign recognition and ensuring compliance with traffic regulations. Another common application is the recognition of vehicle license plate numbers at highway toll gates or in modern parking lots.

Everyday Uses of Computer Vision Systems



E-commerce

E-commerce platforms such as Alibaba and Amazon employ CV to analyze product images and recommend similar items to customers, enhancing the shopping experience.

Transportation and Logistics

Computer Vision enables the monitoring and optimizing of goods flow. It includes tasks such as automatically scanning and tracking shipments in logistics centers.

Agriculture

The technology has been successfully used to monitor the health of crops and livestock. It involves analyzing aerial photos to detect plant diseases or monitoring pasture conditions.

Tourism

Tourism companies can integrate CV into their applications to assist in navigating and exploring new places. It might include recognizing landmarks and providing real-time information about them.

AI glossary



Clustering: An algorithm that segments data into groups based on similarities, spotlighting patterns that might elude human observers.

Convolutional Neural Networks (CNN): A type of neural network that utilizes a convolution mechanism to detect and process patterns, such as in images. Instead of analyzing the entire image simultaneously, CNNs review small portions using filters, learning to recognize various features, such as shapes or colors. With each additional layer, the model extracts features from the image at an increasingly higher level of abstraction. Initially, these features identify simple shapes, which later become features that specify the presence of, for instance, a face or a specific hand gesture. Apart from image recognition, this is also useful in other tasks, such as Natural Language Processing (NLP) or Optical Character Recognition (OCR).

Data Analyst: An analyst who utilizes statistical and mathematical techniques to convert large data sets into useful information. Their role involves data analysis to aid organizations in making data-driven business decisions.

Data Engineer: An engineer tasked with building, testing, and maintaining data infrastructure, like databases and large data processing systems. They work on designing, building, and integrating techniques to gather, store, process, and analyze large volumes of data.

Data Mining: the process of uncovering patterns, correlations, trends, anomalies, or dependencies in large data sets that are often hidden or non-obvious. Data mining usually includes several steps, like data selection, data processing and cleaning, and the evaluation and interpretation of results. In the context of Big Data, data mining is particularly critical as it converts vast amounts of disordered data into valuable insights.

AI glossary



Data Science: A wide field concentrating on conducting research, analyzing, and interpreting information from data. It employs various techniques and tools like machine learning to gain valuable insights from large data sets.

Data Scientist: A specialist in data science who employs statistical, mathematical, and programming knowledge to analyze and interpret large data sets. They create predictive models, analyze trends, and solve complex business issues.

Decision Trees: usually represent decisions and potential outcomes in a tree-like format. Useful for both predicting values and classifying data. Their primary advantage lies in the straightforward interpretation and knowledge extraction from the model.

Deep Learning: A machine learning subdomain that employs deep neural networks for intricate tasks. Owing to its profound depth and complexity, it excels at analyzing intricate patterns in extensive datasets, proving instrumental in advanced image recognition, natural language processing, and audio analysis.

Deep Neural Networks: A specific subtype of neural networks with multiple layers between the input and output. Each layer consists of neurons that process data and relay it onward. More layers allow the model to discern intricate and abstract features in data.

Deconvolutional Neural Networks: Networks that function similarly to convolutional networks but in reverse. They use information that the convolutional network might have discarded during analysis to uncover hidden features or signals. They are especially useful in image analysis, helping the computer to “decode” their details.

Feed-forward Neural Networks: A type of neural network that processes data step by step, transmitting information from input to output. They help computers recognize images, understand language, or analyze complex patterns in data.

AI glossary



Generative Adversarial Networks (GAN): An advanced type of neural network that consists of two parts: one generates new data, and the other attempts to distinguish these new data from the real ones. They are used to create realistic simulations, such as generating images that look like actual photos.

Language Model: A system for predicting the sequence of words or sentences based on prior examples.

Linear Regression: An algorithm that predicts a continuous numerical value based on a linear combination of other variables. For example, it might estimate an apartment's rental price considering factors like square footage, room count, and location.

Logistic Regression: A model for categorical predictions, such as binary outcomes. It is commonly used in tasks like spam classification.

Machine Learning: A branch of artificial intelligence that involves using algorithms to analyze data, learn from it, and make predictions or decisions based on patterns within the data.

Midjourney – an interactive program-bot from the field of generative AI used to create images based on text. Its operation involves the user describing what they want to see, and then the program finds images that fit the description and combines them to create a unique composition.

Modular Neural Networks: Complex neural networks composed of many smaller, independent networks, each focusing on a specific part of the problem. It is akin to a team of specialists, each focusing on their expertise, enabling faster and more efficient problem-solving.

AI glossary



Morphological Analysis: Examining the structure of words, focusing on elements like roots, endings, and other morphological components.

Named Entity Recognition (NER): Identifying and classifying key information in text, including names of people, places, and organizations.

Natural Language Generation (NLG): The creation of human-understandable text by a computer.

Natural Language Understanding (NLU): The process by which a machine interprets, understands, and processes human speech or text to extract meaning.

Neural Networks: Models inspired by the human brain's structure, emphasizing pattern recognition. They are applied in language translation, image, and speech recognition.

Node / Neuron: The basic component of a neural network that receives one or more inputs, processes them and then returns an output.

Part-of-Speech Tagging: Assigning grammatical categories, such as nouns, verbs, and adjectives, to each word in a text.

Predictive Analytics: Process that forecasts future outcomes and trends based on historical and current data. Companies can utilize predictive analytics to anticipate sales trends, customer behaviors, financial performance, and even risk, among other factors. This data-driven approach enables better business decisions.

Random Forests: An ensemble method integrating multiple decision trees to bolster prediction accuracy and counteract overfitting.

AI glossary



Recurrent Neural Networks (RNN): Neural networks that can remember information from previous calculations. It is especially useful when past information impacts current results. They are typically used for understanding and generating speech and predicting trends.

Reinforcement Learning: A type of machine learning in which the model is trained through a system of rewards and penalties, enabling it to learn from its own decisions.

Semantic Analysis: A process of understanding and interpreting the meanings of sentences.

Stemming: Simplifying words to their base form (e.g., "running" to "run").

Structured Data : Data that is organized in a specific manner, thus making it straightforward to store, process, and analyze. It can, for example, be organized in tables, where each column represents a specific field (like first name, last name, address), and each row.

About Us

Qtravel.ai is a company revolutionizing the travel industry by providing cutting-edge, personalized travel search and discovery experiences. Our customizable AI-powered travel search engine not only aims to reduce costs but also seeks to improve customer conversion rates by offering efficient, tailor-made solutions for each individual traveler.

Contact Us



www.qtravel.ai



(+48) 607 999 040



contact@qtravel.ai



ul. Jabłoniowa 20/512, 80-175 Gdańsk, Polska



www.linkedin.com/company/qtravel-ai



www.facebook.com/qtravelai

Autors

Katarzyna Rojewska

Karolina Selwon

Paulina Kreft