

AI/ML - An overview of industry trends and Cisco CX use-cases

March 2021

Contents

Executive summary	3
Why AI/ML?	3
AI/ML Types	5
AI/ML Trends	7
Cisco CX AI/ML Use Cases	11
Customer segmentation use-case implementation journey	18
Conclusion	21
Authors	22
Contributors	22
Reviewers	22

Executive summary

Why should you care about leveraging AI in your organization? What is the current AI/ML landscape, business use-cases and evolution trends? How can you deploy AI/ML in your organization?

This white paper helps enterprise and service provider technical and business decision makers get answers to those questions. It is particularly aimed at helping organizations who are adopting AI/ML and are looking for IT Service Management use-cases / implementation references.

The whitepaper is split into 5 parts. In the first part it highlights why organizations should care about implementing AI/ML projects: overall AI/ML business size, growth potential and major use-cases currently being considered by the industry.

In the second part a taxonomy of the AI/ML algorithms and data usage types is presented with a special focus on the frameworks that are being actively developed and relevant for the industry.

The AI/ML domain future evolution is presented as part of the third section. Special focus is given to the obstacles, needs and frameworks relevant for evolving the AI/ML enterprise practice, technology accelerators and relevant AI/ML algorithms evolution trends.

As part of section 4, the paper illustrates relevant Cisco CX examples as implemented in our organization and in customer environments. These use-cases leverage our collection of intellectual capital, telemetry, customer interactions and can run on Cisco CX cloud or on our customer premises. They have helped improve customer satisfaction, reduce thousands of hours spent on solving issues and predict potential customer equipment problems before they occur. For each use-case, the alignment with the wider industry use-cases, Cisco CX service offer, as well as the expected business relevance, is highlighted.

In the last section, we present a customer AI/ML implementation example where Cisco CX has enabled and improved customer conversion rates.

AI/ML being one of the fastest developing industry segments, requires enterprises to adopt a multi-domain strategy for its proper adoption, considering technology and algorithmic evolutions as well as a rigorous company-wide practice approach. This paper presents some of the most critical areas to consider and shares the experience Cisco CX has in leveraging AI/ML models, as well as deploying them together with customers.

Why AI/ML?

Businesses across the world have not been short of challenges during 2020. The global pandemic is forcing a sudden readjustment of their supply chain, customer interactions and workforce to the new working from home setup. The sudden change of customer behavior and interaction patterns has placed tremendous pressure on finding more efficient ways to readjust their organizational processes and adapt to the “New Normal”.

This has brought new light into the emerging field of AI/ML and its inherent value in helping businesses become more agile by gaining actionable insights from the continuously increasing amount of data available.

The AI/ML field is not new. In fact, some of its underlying concepts predate 20th century (for example Bayes Theorem) and the first methodological approach started in the early 1950s (Alan Turing Learning

machine concept and Marvin Minsky first Neuronal Network Algorithm). AI/ML has had a non-linear development during the 20th century with advances during the 1960s / 1980s and late 1990s, but also periods of stagnation especially due to lack in available processing power, data, and systems.

After the 2000s AI/ML has constantly gained attention from businesses that are looking to gain insights on almost every area of their business, such as: operations, R&D improvements, sales and financial forecasting, customer interactions, marketing campaigns, business processes, and supply chain optimization.

Today AI/ML technologies are quickly becoming mainstream with 85% of enterprises evaluating or using artificial intelligence in production. ^[1]

IDC is indicating that worldwide revenues for the artificial intelligence (AI) market, including software, hardware, and services have totaled \$156.5 billion in 2020, an increase of 12.3% over 2019 and will likely surpass \$300 billion in 2024 with a five-year compound annual growth rate (CAGR) of 17.1%. ^[2]

From an industry use-case perspective, the main enterprise focus today is on the customer experience front. This involves improving customer interaction and generating customer insights fueled by increased focus on concepts like Natural Language Processing, Anomaly Prediction/Detection, Reinforcement Learning and Human-AI/ML collaboration, as shown in Figure 1 below.

Cisco Customer Experience (CX) is Cisco’s main interface towards customers and partners throughout our various product and solutions lifecycles. For this reason, CX is a natural place to leverage AI/ML to implement a significantly enhanced customer experience, aligned with the industry trends as we will see in the CX section below.




Figure 1. Algorithmia’s Third Annual Survey, 2021 Enterprise Trends In Machine Learning ^[3]

From the industry trends, it is clear that AI/ML will help you:

- Simplify and scale data operations.
- Increase accuracy of forecasts.
- Decrease time to market.
- Enable insights on otherwise unusable data.

AI/ML Types

Given its long history of continuous organic and inorganic growth in the industry, considerable number of methodologies, algorithms and use-case applicability, the AI/ML taxonomy is dynamic. Continuously adjusting with the latest evolutions in the field. The choice of the AI/ML algorithm may highly impact the business outcome you are targeting in many ways, such as training time, prediction accuracy or resources usage. This choice will be dictated, most of the times, by the type and amount of data you have available for your use-case.

Looking at the type of data, learning being employed and future evolution, the industry AI/ML landscape can be split in 4 major areas, 3 of them being well-established and the 4th one covering the areas of future development. Our view is represented in Figure 2.

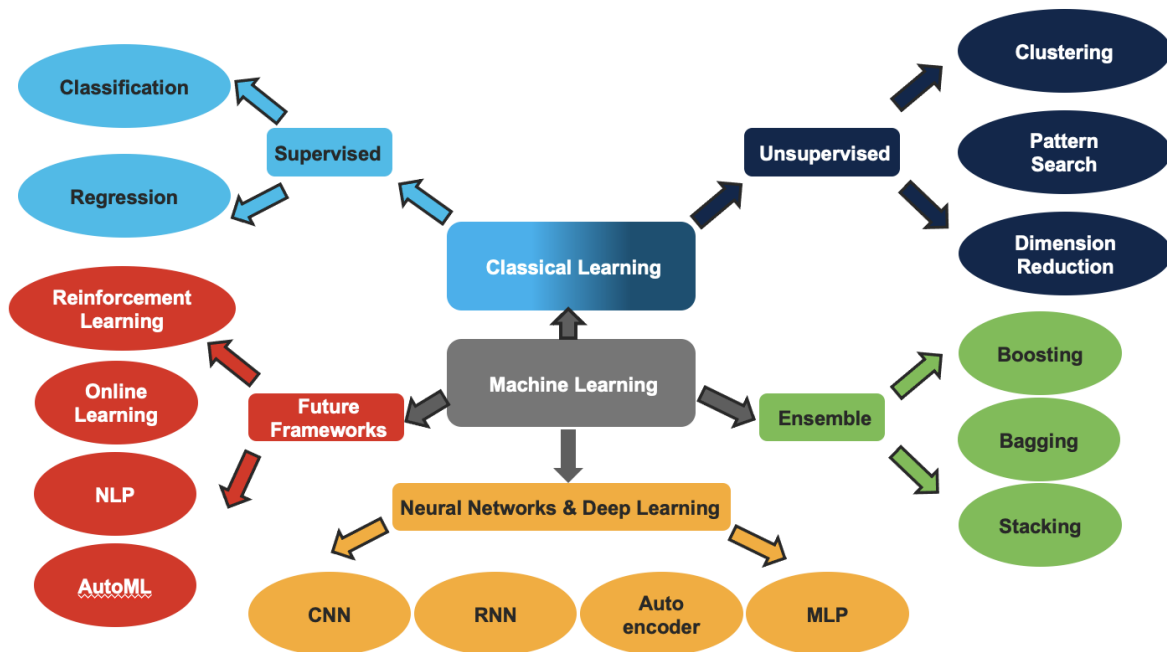


Figure 2. Machine learning landscape

Classical ML (when available data is simple or significant pre-processing has resulted in high data quality):

- **Supervised learning:** This assumes the data is pre-interpreted (labeled) and tries to predict an output given a specific set of data based on the past examples. This can be further split into: **Classification** (deciding in which category the output will belong) and **Regression** (predicting the value the output will have).

In your environment you might have hit interface saturation in the past which resulted in a traffic outage. In this situation, you might want to use a regression to predict interface load. You could store the interface counters in a database and use them to predict the future interface state, if the prediction indicates a probable future problem you could act on it before an outage occurs based on identifying data/usage patterns.

Another example are device crashes. Sometimes, a network device simply crashes/reboots. This typically causes a business impact; a ML classification algorithm can help in this situation. By collecting relevant properties (also known as features in AI/ML) from the devices (for example software and hardware version, statistics, configurations) you can train a model to predict if a device will crash or not in a near future based on the previous devices' features that have crashed. Here, model, represents the output of an algorithm run on data. In short, what was learned by the algorithm.

You will see both examples were implemented by Cisco CX and are reflected on Table 1.

- **Unsupervised learning:** This assumes the data is not pre-interpreted (unlabeled) and tries to find human interpretable patterns without any prior knowledge. This can be further split into **Clustering** (group based on similarity), **Association** (identify specific sequences of items in relation with each-other) and **Dimension Reduction** (reduce the number of specific features by assembling them into higher level groups).

In the previous example of classification, we had to know which device has crashed in order to label it. What if you want to know when a device crashes if there are others that may be at risk as well? You can achieve this by using clustering algorithms. After collecting devices characteristics, you can use a clustering algorithm to group them by similarity. If one of those devices' crashes, you then know which other devices are similar and may be at risk.

Ensemble methods (when data is of lower quality or sparse and produces high bias / variance errors or reduced predictions using Classical ML methods):

- **Stacking:** This assumes the decision is taken by multiple algorithms working together and their conclusions are forwarded to the final stage algorithm that has the final say. This helps improving the prediction power.
- **Bagging:** Use different subsets of the original data to train multiple instances of the same algorithm and aggregate their answers. This helps in reducing the variance errors.
- **Boosting:** Use sequentially trained algorithms to analyze the same data, each stage trying to address errors from the previous stage. This helps in reducing the bias and errors in the final output.

The algorithms may be supervised, unsupervised or a combination of both. Ensemble methods try to address the challenges of single models' methods, improving the robustness and accuracy of results.

Neural Networks and Deep Learning (when you have substantial amounts of unstructured/structured data or unclear features these may offer superior accuracy over Classical ML methods): Weights and functions replicating the natural neuron and connections, organized in layers and trained to solve specific problems (e.g., image, speech and text recognition). Training is based on feeding back how 'correct' a specific output is when an input is introduced in the network (through backward propagation).

- **CNN** (convolutional neural networks): used in picture and video detection and interpretation. Relies on identifying the most relevant features of the object and deciding on how many matches are present given a specific input.
- **RNN** (recurrent neural networks): interpreting sequential data like text, voice and music. Relies on adding additional memory to each neuron allowing it to remember its past decisions and thus recognize sequences in speech / text / music. Additionally, the memory space constraint is addressed by allowing neurons to decide on short term decisions that can be changed (flags) and long-term connections that support the model (the concept is called Long-Term Short-Term Memory cells - [LTSM](#)).
- **Autoencoder**: used with unsupervised data, typically for dimensionality reduction, the goal is to learn a representation (encoding) of data ignoring the noise in the data. They consist of two concepts: 1. The encoder, tries to represent the same information in a different (usually smaller format). 2. The decoder translates the encoded data into a representation as close as possible to the original.
- **MLP** (multilayer perceptron): used with classification or regression predictions. They are the 'classical' type of neural network, made of an input layer, one or more hidden layer and an output layer where predictions are made.

Future Frameworks: This is an area highlighting some of the most promising evolutions of the AI/ML domain. They are further detailed within the framework trends part of the next section.

At this moment, Supervised Learning and Neural Networks / Deep Learning are the most critical areas of investment for enterprises. In a recent study, [deep learning](#) displaced supervised learning as the most popular technique among organizations that are in the "evaluation" phase of AI adoption: (~55%) say they are using more deep learning than supervised learning (~54%) and close to 66% of respondents who work for "mature" AI adopters say they are using deep learning, making it the second most popular technique in the mature cohort—behind supervised learning. ^[4]

Most of the use-cases that are currently researched and invested upon by businesses (such as automated machine learning, reinforcement learning and conversational AI) employ a combination of the AI/ML categories listed above. For example, unsupervised learning can be used to simplify and label data that is fed into a supervised learning / neural network algorithm.

AI/ML Trends

While AI/ML is currently actively implemented by enterprises across the world, the entire domain continues to evolve at an incredible speed, fueled by advances in the **enterprise AI practice interest/maturity, technology, frameworks, and use-cases**.

Enterprise AI/ML practice evolution:

In a recent survey of the industry, just 16 percent of respondents say their companies have taken deep learning beyond the piloting stage. High-tech and telecom companies are leading the charge, with 30 percent of respondents from those sectors saying their companies have embedded deep-learning capabilities. ^[5]

At the same time looking at the barriers to a wider AI adoption, in another study, the shortage of ML modelers and data scientists topped the list, cited by close to 58% of respondents. The challenge of understanding and maintaining a set of business use cases came in at number two, cited by almost half of participants.

Slightly more than one-fifth of respondent organizations have implemented formal data governance processes and/or tools to support and complement their AI projects showing the relative low governance implementation.

Other obstacles listed are: unexpected outcomes/predictions, [interpretability and transparency](#) of ML models and fairness and ethics. ^[6]

Another evolution has been around transparency, interpretability, and bias. These are particularly important given the recent work from EU to address the topic within the wider context of GDPR. [\[7\]](#)

Most of the enterprises start their journey with creating ad-hoc proof-of-concept projects satisfying the local needs present in each department however the scattered nature of this approach prevents an optimal production roll-out.

AI/ML engineering

As AI systems and capabilities, models and algorithms are being developed at a rapid pace, it is a common challenge that many of these capabilities fail to go beyond their proof-of-concept stage as they work primarily in controlled environments and fail to deliver significant business values in the real-world scenarios.

The need for an engineering discipline to guide the development and deployment of AI capabilities is paramount. AI Engineering is an emerging engineering discipline that combines the principles of systems engineering, software engineering, computer science, and human-centered design. AI Engineering aims to provide a framework and tools that would equip practitioners to develop scalable, robust, secure and human-centered AI systems for the real-world use.

Considering that AI system is in fact, a software-intensive system, it is important to understand the core differences between AI engineering and software engineering, and why the tools and processes designed specifically for AI are critical for companies to achieve their AI transformation:

1. Data requirements - complex and difficult (discover, source, manage, version, legal and compliance).
2. Model customization and reuse requires vastly different skills.
3. No strict module boundaries.
4. Uncertain property of the changing world that can impact the models drastically.

A simplified workflow of AI Engineering should include DataOps, MLOps, and DevOps.

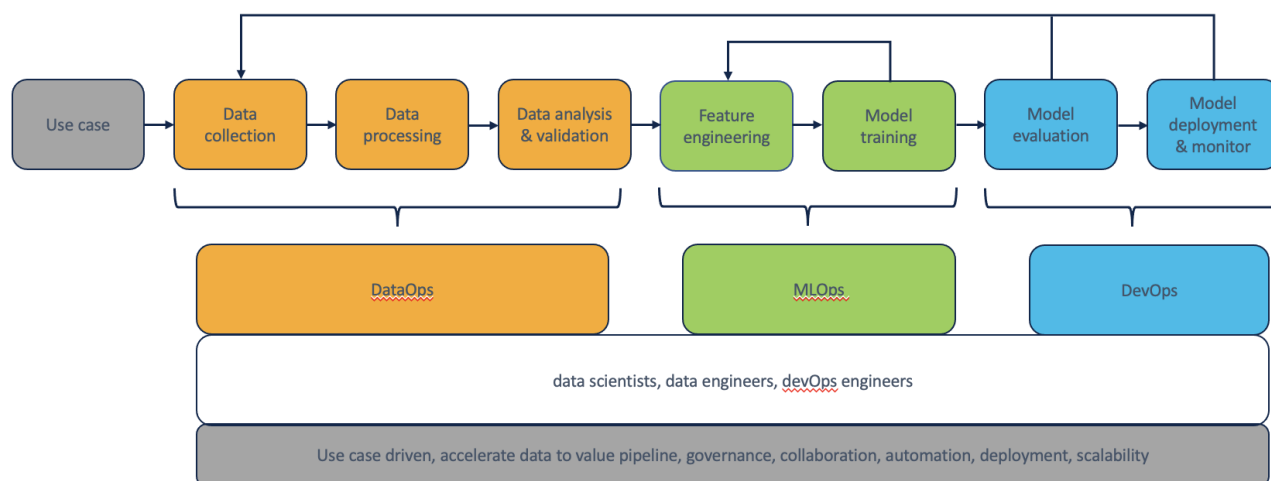


Figure 3. AI engineering workflow example

AI/ML enterprise practice evolution

Supporting the above-mentioned evolution, there is an increasing need for businesses to build a proper structured organization to formally manage their AI/ML approach.

Developing a mature AI/ML practice must cover multiple angles such as:

- **Ensuring** AI/ML strategic alignment with enterprise business goals.
- **Building** an enterprise-wide standardized data collection/storage, treatment and model training platform.
- **Defining** a common set of best practices supporting AI/ML deployments and operations into production.
- **Enforcing** a central governance process that allows review/raking and management of AI/ML projects in line with business needs.

A central data science office can facilitate the adoption of well-established standards throughout the organization, provide guidance and steer all the distinct AI/ML initiatives in alignment with the business goals set by the enterprise.

Using a common set of data standards and AI/ML model development framework is key to ensure consistency and broader insights across the different organization departments.

As enterprises move to operationalize at scale their AI/ML deployments, they will continue to improve their AI/ML practice with the new learnings discovered during their journey.

An overview of the AI/ML enterprise practice maturity evolution is presented below:

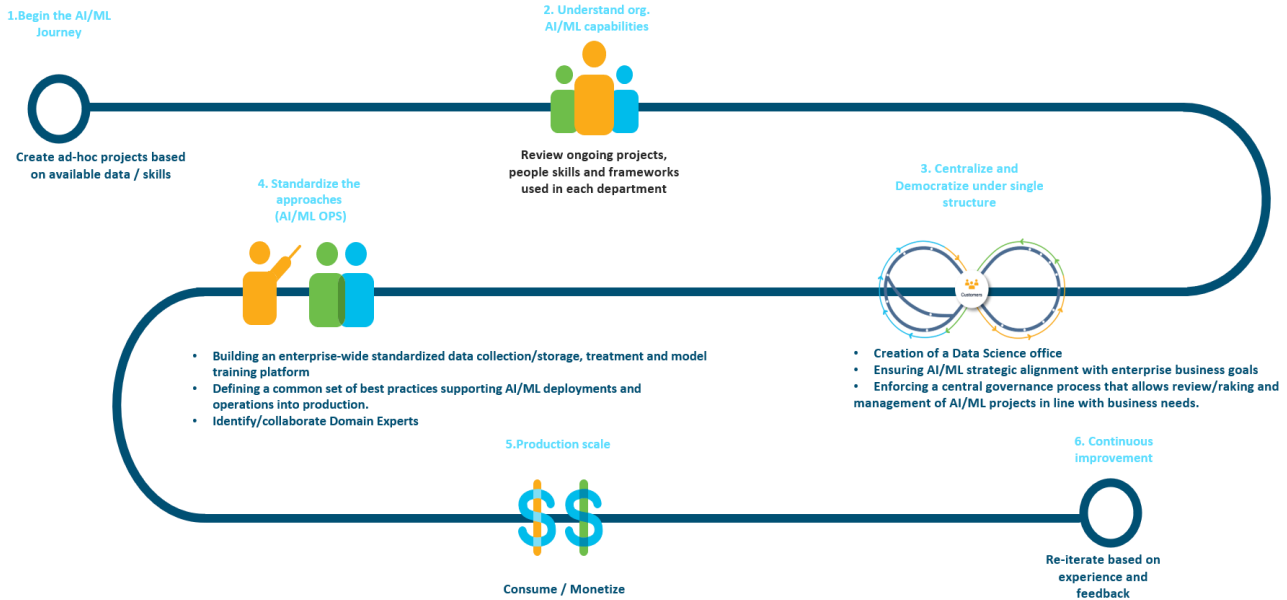


Figure 4. Enterprise AI/ML practice maturity journey

Technology accelerators:

AI/ML evolution is benefiting from the evolving technology trends in areas related to data, infrastructure, device adoption, skillset evolution and model reusability:

- **Increase in data/ compute and storage:** with the expected evolution of 5G the world continues to see increased data generation and number of new connections. According to Cisco VNI there will be close to 30 billion devices/connections by 2023 – 45% of those will be mobile. Globally, Internet traffic will grow 3.7-fold from 2017 to 2022, a compound annual growth rate of 30% [8]. At the same time compute power and storage size will continue to increase for a given price-point although the future alignment with Moore’s Law remains unclear. This will act as AI/ML enablers as businesses will struggle to make sense of the ever-increasing data amount.
- **Edge computing:** as AI/ML become more efficient in power and memory consumption, they will be able to be placed on the edge and become an essential part of data collection and data-driven decision distributed architecture.
- **Cloud Infrastructure:** the ubiquitous availability of the cloud platforms and cheap access to storage and compute power is easing the setup and adoption of AI/ML projects from proof of concept to production.
- **Cybersecurity:** with security breaches on the rise and the evolution of IP connected critical infrastructure and personal devices, AI-driven threat recognition and mitigation will continue to evolve and be an essential element in fighting cybercrime.
- **Increase in devices running AI:** As the cost of microprocessors decreases, connectivity increases (due to 5G adoption) more appliances will adopt AI-powered algorithms for providing basic tasks. From smart wearables and autonomous home appliances to autonomous vehicles, the AI/ML applicability in our daily life will only become more visible.

- **Pre-trained models:** Companies (for example, cloud providers) are making available pre-trained ML models, which other companies can use at a fraction of the cost and without the need of having a workforce trained in AI/ML. This eases the adoption of AI/ML throughout the industry.
- **Human and AI cooperation:** As companies continue to invest in AI/ML initiatives and tools to replace tedious / repetitive tasks, humans will need to shift their skills and learn to leverage AI/ML to achieve higher complexity tasks, such as design, strategy, innovation and art. IDC predicts that by 2025, 75% of organizations will be investing in employee retraining to fill skill gaps created by the need to adopt AI/ML. [\[9\]](#)

Framework evolution trends:

While the entire AI/ML algorithm stack will continue to improve, there are a few areas that are particularly well positioned to support the technology accelerators mentioned above:

- **NLP:** Natural Language Processing deals with the interaction between computers and humans. NLP can extract knowledge from machine to human or human to human interactions. Its ultimate purpose is to read and understand the communication in a way that is meaningful and usable in an algorithmic context. NLP has multiple exciting applications within the customer experience context, such as: sentiment analysis, language translation, topic extraction.
- **Conversational AI:** this is an algorithm often implemented by various business use-cases. It seeks to empower the human-computer contextual (including emotion recognition) interaction by leveraging multiple AI/ML frameworks such as Automatic Speech Recognition and NLP. This is a fundamental area enabling the human and AI cooperation of the future.
- **Reinforcement Learning:** (when data can be generated by interacting with the environment): Its main goal is to understand the rules of the environment by interacting with it and minimize the interaction errors. This is achieved by setting a system that rewards the system when following the rules and penalizes it when it does not. Although as of now the reinforcement learning algorithms do not top the business use-case implementation, this remains an exciting domain and one of the paths specialists see towards achieving the next major AI/ML breakthrough.
- **Online Learning:** this is not yet widely used outside of the biggest IT companies. Online learning is a methodology that uses data as it becomes available to improve and enhance existing ML models. Online learning aims to improve future predictions with every set of received new data, unlike the traditional models that are using batch learning models. With the increasing need of getting near real-time insights, it could become a mainstream method.
- **Auto-ML:** this is an area that is actively evolving and being improved. It aims to automate most of the time-consuming tasks within a given AI/ML deployment process with a special focus on feature engineering and ML model selection and building. Its associated practice is called ModelOps. AutoML significantly helps the AI/ML enterprise adoption by reducing the time and complexity required to define and train AI/ML models.

Cisco CX AI/ML Use Cases

Cisco Customer Experience (CX) focuses on helping customers realize the full value of their technology solutions. Its portfolio can be grouped in 5 major categories, aligned with various customer needs during the solution usage journey:

- **Support:** insights-driven technical assistance and problem resolution.
- **Guide:** analytics-driven proactive recommendations and assistance.
- **Operate:** management, monitoring and optimization of customer networks.
- **Assist:** design and implementation of emerging technologies.
- **Enable:** training, certifications, and developer community's facilitation.



Support



Guide



Operate



Assist



Enable

As part of its daily engagements, Cisco CX deals with an enormous amount of data, including but not limited to: customer issues and ways to solve them, telemetry data, software/hardware/functional specifications, log data and various operational details, best practices, and other types of intellectual capital. This data constitutes the foundation for delivering unique customer services and is a natural source for leveraging AI/ML techniques to derive additional insights supporting each of the 5 above mentioned categories.

The applications vary, whether it is resolving service requests faster, enhancing recommendations CX provides to its customers or providing faster services implementation. AI/ML is present across the entire CX portfolio, exploiting the power of data.

As an overview of the current AI/ML landscape within Cisco CX, the table below illustrates some use-cases tackled, the AI/ML framework used, alignment with CX services' offering and the expected business value.

Table 1. CX use cases

Use-case description	Framework / ML Type	CX Service Offering	Alignment with industry use-cases	Business relevance and impact
Predict hardware failures based on telemetry data	Clustering	Guide/Operate	Supply chain optimization	Allows for early detection of hardware failures, enabling early replacements and reducing business outages.
Increase DevOps system's effectiveness by test de-duplication and conversational AI for quicker insights and	Ensemble Bagging, Deep learning, NLP	Assist	Automating processes	DevOps testing quickly scales out of proportion especially when done by different vendors and teams.

Use-case description	Framework / ML Type	CX Service Offering	Alignment with industry use-cases	Business relevance and impact
issue resolution.				De-duplication of test cases, and recommendation of similar ones upon creation maintains test coverage while decreasing costs and required resources.
Assist in the quicker resolution of customer incidents and questions based on previously solved similar situations	Ensemble Boosting, Deep learning	Support	Improving customer experience	On-screen recommendations and conversational AI help customers find answers to their questions quicker, enabling in many cases faster MTTR.
Identify and quantify risk based on Cisco's network devices SW/HW and functional criticality within the network	Deep learning, NLP	Guide/Operate	Generating customer insights	Being able to automatically identify and quantify risk linked to network equipment business criticality, allows customers to act in a proactive manner focusing on the devices that are most relevant for their business.
Forecast bandwidth capacity shortage based on collected network telemetry	Deep learning, Clustering	Operate	Generating customer insights	Allows for early detection of capacity shortage allowing increased time to action, reducing the number of outages and accompanying business impact.
Encode CX intellectual capital into ML models for on-prem secure deployments in customer networks	Ensemble Bagging	Guide	Acquiring new customers	Allows secure access to Cisco IC stored locally without sending customer data outside the company. Facilitates onboarding data-sensitive customers, which cannot use the Cisco CX Cloud model: e.g., Financial Institutions and Public Sector.

Use-case description	Framework / ML Type	CX Service Offering	Alignment with industry use-cases	Business relevance and impact
<p>Identify groups of devices with common characteristics</p>	<p>Clustering, Dimensionality Reduction</p>	<p>Guide/Operate</p>	<p>Generating customer insights / Increase customer loyalty</p>	<p>Grouping devices by similarity helps in several ways, for example, when a device crashes, we know which other devices may experience the same failure. Another example is identifying configuration inconsistencies across an install base. This grouping saves engineers' and their business time.</p>
<p>Subscriber segmentation based on network activity data</p>	<p>Deep learning, Clustering</p>	<p>Operate</p>	<p>Recommender systems</p>	<p>Allows for customer segmentation based on their behavior in the network facilitating targeted marketing offers. In consequence, increasing revenue.</p>
<p>Ease the access into customer made ML models with an ML workbench.</p>	<p>Any</p>	<p>Assist</p>	<p>Improving customer experience</p>	<p>The workbench helps customers build and test their own models from network metrics using industry algorithms with no effort. It helps enabling network insights otherwise blocked behind a complete ML lifecycle. It is powered by a Spark back end.</p>
<p>Ease research within a vast intellectual capital repository by providing similar and relevant resources</p>	<p>NLP</p>	<p>All</p>	<p>Recommender systems</p>	<p>Searching for information is an activity that occupies most of the time spent by engineers. By facilitating this search with a ML recommendation system, we achieve a reduction in costs and an increase in engineers' time to focus on tasks only they can perform.</p>

Use-case description	Framework / ML Type	CX Service Offering	Alignment with industry use-cases	Business relevance and impact
Predict the role and importance a device plays in the network using its configuration	Ensemble Bagging	Guide	Recommender systems	Some devices play a more critical role than others. By predicting the role and importance of each equipment, we have a view of where the critical points are and address them.
Automation and improvement of Cisco's RMA process	Deep learning, NLP	Support	Improving customer experience / Automating process	By using ML to automate some of the RMA verifications steps, the time to provide replacements to customers is significantly reduced, thus increasing their satisfaction.
Forecast of counters evolution on a Cisco device	Deep learning	Guide/Operate	Generating customer insights	Using on-the-box telemetry to forecast a counter value (for example, interface errors) enables us to act before an outage occurs.
Predict network issues based on Syslog data	Pattern search	Guide/Operate	Generating customer insights / Increase customer loyalty	Typically, devices generate many log messages which can become cumbersome to utilize into actionable insights. By using ML, we predict future network issues minimizing outage risks and enhancing network visibility.
An AI chatbot, enabling self-service experience for common case inquiries and basic transactions.	NLP, Deep Learning	All	Improving customer experience	Managing cases can be a time-consuming activity. Having a ML powered self-service engine decreases time for issue resolution and enables customers to interact with Cisco's case management platform more easily.

Use-case description	Framework / ML Type	CX Service Offering	Alignment with industry use-cases	Business relevance and impact
Generate insights from CLI output	Deep learning	Operate/Assist	Automating processes	Troubleshooting is typically a time-consuming activity and requires advanced technical skills. With the aid of AI/ML, additional knowledge can be gained from the data gathered, decreasing average time for resolution.
Measure customer sentiment in the context of Service Requests communications	AutoML	Support	Improving customer experience	Allows for early detection of possible problematic situations engaging the required resources sooner. Enables faster conflicts resolution.
Anomaly detection based on network traffic data telemetry	Unsupervised learning	Operate	Automating processes	By using telemetry data from the network's data plane, anomalous entries can be identified, which enables quicker actions that may protect the business. An example could be identifying a malicious actor crafted packets.
Analysis and reporting of network equipment policies	Unsupervised learning	Guide	Generating customer insights	Policy configuration represents a large part of network equipment. Over time changes are made to these policies and are often undocumented. The result, in most cases, are different policies applied in places with the same business intent. This solution helps identifying and de-duplicating such policies. Business operations are improved when configurations are kept clean and

Use-case description	Framework / ML Type	CX Service Offering	Alignment with industry use-cases	Business relevance and impact
				consistent.

When looking at the business benefits, the above use-cases help especially in 3 areas:

- Reduce cost by:
 - **Improving operational efficiencies: assisting SR (Service Request) handling via sentiment analysis, allowing fast access to information via intelligent search, reducing effort by enabling test de-duplication**
 - **Reducing MTTR: by speeding up problem resolution via intelligent chat interactions, insightful CLI output and detecting traffic anomalies**
- Reduce risk by:
 - **Predicting network and equipment failures: predicting capacity shortage, HW and SW failures, determining common device risk profiles, forecasting failures based on counter or syslog evolution**
 - **Configuration standardization: ensuring similar devices with similar roles have a similar configuration**
- Increase revenue by:
 - **Enabling businesses to better serve their customers: facilitating more effective marketing campaigns by performing customer segmentation**
 - **Enabling Cisco to serve new categories of customers: by providing on-prem customer usable Cisco intellectual capital**

The list of AI/ML use-cases shows a balanced approach supporting Cisco customers across the CX service portfolio. While multiple AI/ML frameworks are being used, there is a clear focus on leveraging forward-looking frameworks such as NLP, AutoML and Deep Learning in over 50% of the use-cases.

From a deployment and customer interaction perspective there is a balance between use-cases that are delivered off-prem (especially centered around support services that are using data available on Cisco premises, such as Cisco intellectual capital or service requests data) and hybrid (centered around Guide/Assist and Operate services that are processing customer equipment data such as telemetry and logs either on-prem or off-prem).

The use-cases listed follow the industry distribution mentioned earlier. Most use-cases improve customer experience (in over 50% of the cases), followed by enhancing automation (in over 30% of the cases). There are also supply chain optimization and customer acquisition use-cases.

Developing AI/ML use-cases is a time and resource consuming task that typically starts by a business need. Followed by a selection and filter of relevant data, finding, and tuning the most effective algorithm and most importantly operationalizing the result for production use.

Cisco CX can also work directly with customers to develop specialized AI/ML uses cases. The section below shows an example where Cisco CX worked with a service provider customer to deploy a business-relevant use-case. It details the experience and journey developing and deploying the solution.

Customer segmentation use-case implementation journey

Customer segmentation is a term used to describe the process of dividing customers into homogeneous groups based on shared or common attributes (for example habits, tastes, or others). Segmentation enables telecommunication operators to communicate with their customers more effectively. Specifically, this allows service delivery and marketing to act on each segment, accordingly, enabling proper resource allocation and fulfilling specific business objectives.

Some widely used segmentation types are:

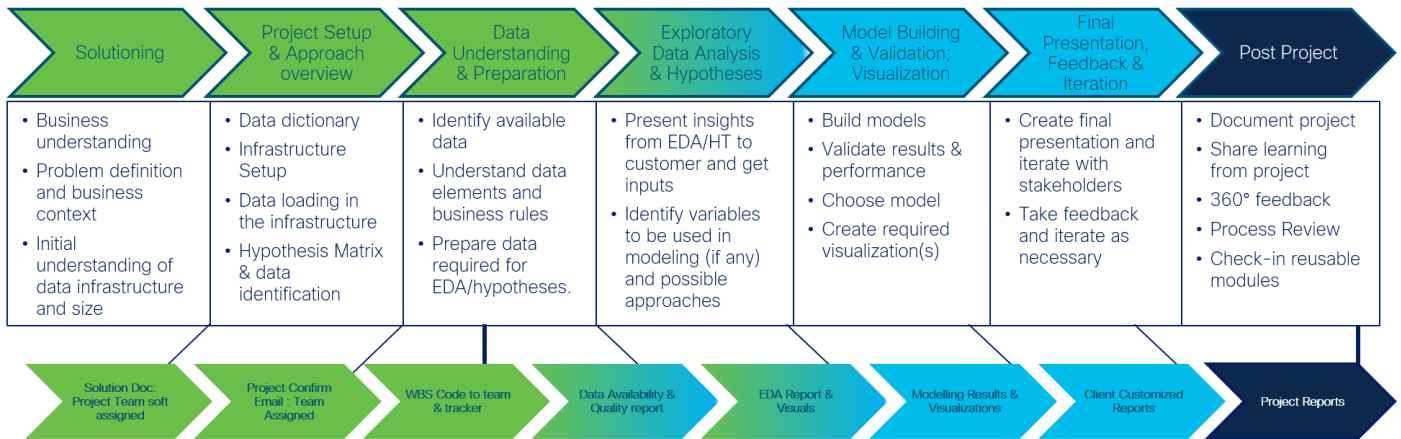
- **Value-based:** where customers are grouped based on actual value perceived and delivered to them.
- **Behavioral:** where customers are grouped based on their identified behavior and service usage patterns.
- **Propensity-based:** where customers are grouped based on their churn / cross-selling scores.
- **Needs/attitudinal:** where customers are grouped based on their needs, wants, attitudes, preferences, and perceptions pertaining to the company's services and products.

In this specific use-case, the customer was looking to leverage the significant amount of network data already collected by them, in order to define and track subscriber segmentation patterns fitting one of the above categories.

The nature of the anonymized data being collected (both mobile signaling data such as: SMS, voice and data call details and location information), as well as data traffic details (mainly coming from the various proxies present across the network) has helped the customer and the Cisco CX team to select the **Behavioral** segmentation as the use-case to be implemented.

The overall ideation, data selection, algorithm selection and operationalization process has followed a defined delivery process allowing for clear success metrics definition and quality gates at each stage of the AI/ML use-case implementation, as presented below:

Analytics Project Lifecycle



Adapted from CRISP-DM to

- Align with cisco-cx engagement & billing process
- Utilize cisco Analytics Infrastructure
- Recognize the role of information visualization in Analytics
- Continuously improve quality by leveraging experts and peer reviews
- Evaluate performance and identify learning opportunities for team

Figure 5. Cisco CX AI/ML implementation stages

The main stages that were performed during this project were:

- **Solutioning and Approach overview:** In this stage the focus was on getting the requirements from the lines of business and documenting them. In parallel, an audit of the infrastructure, data and people expertise needs was made supporting the approach. The various operations described above required different skillsets working in close cooperation through the entire project, such as:
 - Domain experts: They worked with the business user, understood the requirements and worked hand in hand with the data engineers and the data scientists.
 - Data Engineers were responsible for devising methods to connect into the data sources, clean, aggregate and organize the data from different data silos and transfer it into the data lake.
 - Data scientists were responsible for using the data sets to identify and reach meaningful conclusions in the form of strategic business decisions, they also selected and customized algorithms and predictive models for the analysis.
 - Production & Operations team were responsible for operationalizing the models built by the data scientists into production. Once operationalized, they were responsible for the day-to-day operations of the model, such as looking at the infrastructure utilization and health, as well as the overall performance of the model itself.
- **Data Understanding and Preparation:** Once the business requirement was documented, a domain expert along with a data science engineer identified how to solve the use case. In this phase, the team created a set of hypotheses on how the use-case could be addressed and identify the data sources / attributes that may need to be correlated to build it aligned with the hypothesis.

-
- **Exploratory Data Analysis:** In this phase data scientists looked for patterns, trends, outliers and unexpected results in existing data using visual and quantitative methods. Feature engineering is the process of using domain knowledge of the data to create features that help the ML model to be more accurate and efficient. Domain expertise was especially important while determining which features were meaningful to use for clustering within the given context of this specific engagement. For example, the features which did not define any meaningful customer behavior were eliminated at this stage.
 - **Model building and Validation:** Building the AI/ML model using the previously identified features and testing it. Visualization of the results was a key step in the model testing phase as it allowed the team to assess the validity of the chosen implementation and address potential deviations.

Once the results of the model validation were satisfactory, the model was handed off to the operations team for deployment

- **Iterative customer validation:** In this phase the results of the use-case implementation were presented to the customer stakeholders, feedback was collected on their interpretation, accuracy and business value. The collected feedback resulted in additional iterations of the data preparation and model building stages. In the final validation the results were accepted.
- **Operational implementation:** The operations team worked closely with the rest of the project team. When the model was accepted and moved into production, they are now responsible to ensure the model accuracy and the model resource utilization stays at an acceptable level on an ongoing basis.

As a result of all the process, the following categories of customer behavior were identified based on cluster analysis of the relevant features:

- **Business/Office Users:** These customers originate many voice calls. They call often during daytime and the call duration is high. During the weekend, their activity is very low.
- **Quick Talkers:** In this segment, the customers have a relatively small number of voice calls. Their call duration is also lower than average. However, their SMS and message application usage are relatively high.
- **Active users:** In this segment the customers have high usage of voice calls, data, and SMS. The usage is high throughout the week.
- **Night owls:** In this segment, the average number of calls in the evening and night is high.
- **Weekend warriors:** In this segment, the percentage of calls is higher on weekends than on weekdays.
- **Data Hungry:** In this segment customers have high data usage on applications and an average voice call usage.
- **Basic phone users:** These customers receive a small amount of voice calls. The average call duration is low. They also receive and originate a low number of SMS. They predominantly have calls during the day.

The associated data patterns allowing for their correct classification is shown below:

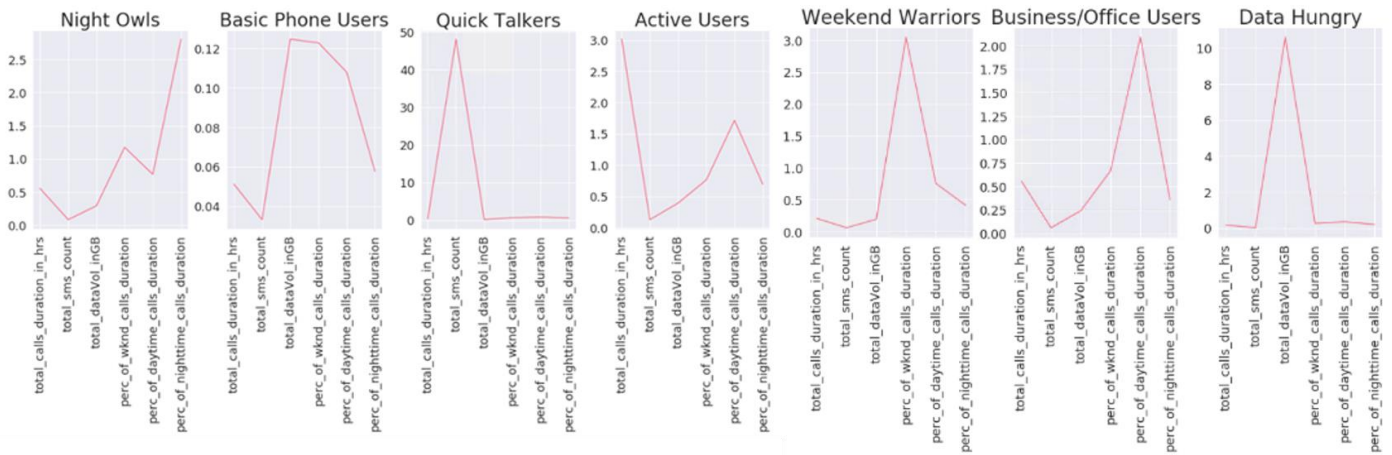


Figure 6. Example of customer behavior data pattern classification

The results of our clustering analysis were built into a user interface for an end user to consume. A heatmap of the geo-location of users of multiple clusters categories can be visualized to give an indication of the users' distribution in a geographic location that belongs to a specific category.

This service provider can now derive meaningful analysis of a population of users of a certain category in a specific location (for example, he is able to answer questions such as: “Which region has the highest concentration of business and office users?”).

This customer insight use-case allows for a proper identification of the service usage patterns and facilitates targeted marketing campaigns addressing users that are in a specific category / location (for example, an offering of increased data to an area identified with high density of Data Hungry users or an offering of easy switch to a weekend service where a substantial number of weekend warriors have been identified).

After the successful implementation of this use-case, the service provider has reported a 3-5 times increase in conversion rate of their targeted marketing campaigns when compared with their initial campaigns.

Conclusion

As part of this paper, we have presented the current AI/ML landscape, trends and industry use-cases which place customer experience at the top of AI/ML priorities.

Ensuring a long-lasting positive customer experience is at the heart of what Cisco CX does in all its daily interactions and this is well reflected by the AI/ML use-cases implemented within Cisco CX.

Each of the 19 use-cases presented is supportive of one or more of the CX Service offerings, highlighting the business value AI/ML can bring. It is also notable that Cisco CX is betting on some of the most promising algorithmic/framework trends in the AI/ML landscape, leveraging them in most of the presented use-cases.

Cisco CX has been leveraging AI/ML models, as well as deploying them alongside our customers. In the future, we will be adopting even more impactful models to provide a superior customer experience and create business value for us and our customers.

Authors

- Stefan-Alexandru Manza, Principal Architect CX EMEAR Service Provider
- Ivo Pinto, Sr Solutions Architect, CX Technical Transformation Group (TTG) EMEAR

Contributors

- John Garrett, Principal Architect CX Product Management, CX Business Critical Services
- Jane (Zizhen) Gao, Principal Engineer, CX CTO
- Anwin Kallumpurath, Director, CX Product Management, Cross Domain Architecture

Reviewers

- Vijay Raghavendran, Distinguished Engineer, CX CTO
- Ammar Rayes, Principal Architect, DevCX Product Management
- Ranjani Ram, Principal Architect, CX TAC Engine
- Ahmed Khattab, Principal Architect, CX Engineering

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)