ARTIFICIAL

INTELLIGANCE

Skill Enhance Course

(As per National Education Policy 2020)

Bangalore University

ABOUT AUTHORS:

Prof. RESHMA B
B.E, MTch.
Assistant Professor and Head,
BCA Department
Seshadripuram Academy of Business studies
Kengeri, Bangalore - 560060

Prof. AKSHATHA M.R
B.E, MTch.
Assistant Professor, BCA Department
Seshadripuram Academy of Business studies
Kengeri, Bangalore - 560060

Prof. AMABRISH R
M.Com, MBA, Dip. IFRS, UGC-NET
Principal, Dharmasagara First Grade College,
Dommasandra, Anekal Taluk,
Bangalore - 262125.

Prof. DADAVALI S.P
MCA, Mphil,
Assistant Professor,
Government First Grade College,
Kengeri, Bangalore - 560060

PREFACE

At least since the first century BCE, humans have been intrigued by the possibility of creating machines that mimic the human brain. In modern times, the term artificial intelligence was coined in 1955 by John McCarthy. In 1956, McCarthy and others organized a conference titled the "Dartmouth Summer Research Project on Artificial Intelligence." This beginning led to the creation of machine learning, deep learning, predictive analytics, and now to prescriptive analytics. It also gave rise to a whole new field of study, data science.

Today, the amount of data that is generated, by both humans and machines, far outpaces humans' ability to absorb, interpret, and make complex decisions based on that data. Artificial intelligence forms the basis for all computer learning and is the future of all complex decision making. As an example, most humans can figure out how to not lose at tic-tac-toe (noughts and crosses), even though there are 255,168 unique moves, of which 46,080 end in a draw. Far fewer folks would be considered grand champions of checkers, with more than 500 x 1018, or 500 quintillion, different potential moves. Computers are extremely efficient at calculating these combinations and permutations to arrive at the best decision. AI (and its logical evolution of machine learning) and deep learning are the foundational

Finally we profusely thank Branch Managers of Kalyani Publishers Mr. Mohammed Sameer for coming forward to publish our work. We will accept the comments and suggestion from our readers to improve the contents.

Prof. RESHMA B
Prof. AKSHATHA M.R
Prof. AMABRISH R
Prof. DADAVALI S.P

ACKNOWLEDGEMENT

We are thankful to a large number of intellectual individuals who have encouraged us in authoring this book. We would like to express our deep sense of gratitude to authors and scholars in the field of management who have shaped us in understanding through their rich and valuable contributions and also to our students fraternity for providing us the stimulus needed for shaping this book in a useful manner.

In particular we wish to thank profusely Dr. Ramachandra Gowda Vice Chancellor Rani Chennamma University, Dr.R.Chenraj Jain, Chairman Jain Group of Institutions, Dr. N. Sundararajan, Vice Chancellor, Jain University, Prof. V.H. Rajshekhar Principal GFGC Ramanagara Dr.B.T.Venkatesh, Principal, Jain University, Dr. B.Anuradha Principal, Al-Ameen Institute of Management Studies Bangalore; Mr.Ravikumar S of Eshcol Cybervilla, Kengeri, Bangalore for their invaluable encouragement, Guidance and Contributions in shaping this book.

Finally we express our deep sense of gratitude of our family members for their whole hearted support in authoring this book in a useful manner.

Prof. RESHMA B

Prof. AKSHATHA M.R

Prof. AMABRISH R

Prof. DADAVALI S.P

SYLLABUS

Course Content (Artificial Intelligence)

	Details of topic	Duration
Course – 1 - Azure AI Fundamentals (AI-900)	AI-900 pathway consists of 5 courses and 2 reading material: i. Introduction to AI on Azure ii. Use visual tools to create machine learning models with Azure Machine Learning iii. Explore computer vision in Microsoft Azure iv. Explore natural language processing v. Explore conversational AI vi. Tune Model Hyperparameters - Azure Machine Learning (Reading) vii. Neural Network Regression: Module Reference - Azure Machine Learning (Reading	05 hours
Practical	1. Prepare the data 2. Model the data 3. Visualize the data 4. Analyse the data 5. Deploy and maintain deliverables	13 hours
Course – 2 - Data Analyst Associate (DA-100)	DA-100 pathway consists of 5 courses and 2 reading material: 1. Get started with Microsoft data analytics 2. Prepare data for analysis 3. Model data in Power BI 4. Visualize data in Power BI 5. Data analysis in Power BI 6. Manage workspaces and datasets in Power BI 7. Key Influencers Visualizations Tutorial - Power BI 8. Smart Narratives Tutorial - Power BI Microsoft Docs	08 hours
Practical	1. Describe Artificial Intelligence workloads and considerations	13 hours

	2. Describe fundamental principles of machine learning on Azure 3. Describe features of computer vision workloads on Azure 4. Describe features of Natural Language Processing (NLP) workloads on Azure	

Course Outcomes (COs):

- Appraise the theory of Artificial intelligence and list the significance of AI.
- Discuss the various components that are involved in solving an AI problem.
- Illustrate the working of AI Algorithms in the given contrast.
- Analyze the various knowledge representation schemes, Reasoning and Learning techniques of AI.
- Apply the AI concepts to build an expert system to solve the real-world problems.

COURSE – 1

AZURE AI FUNDAMENTALS (AI-900)

1.1 INTRODUCTION:

AI is touching us in all aspects of our daily lives, most of the times unknowingly. Whenever we shop online, use our mobiles, drive to work daily, check our mail box or exercise, AI is coming into play and helping us, prodding us or controlling us. Since AI is already such an integral part of our lives, it makes sense to get more knowledge of this emerging technology. From chess-playing computers to self-driving vehicles, Artificial Intelligence (AI) is progressing rapidly and touching every aspect of our lives. In this module, you will learn how machines can be made to learn from data and carry out human tasks. AI enables us to build amazing software that can improve health care, enable people to overcome physical disadvantages, empower smart infrastructure, create incredible entertainment experiences, and even save the planet!

1.2 WHAT IS ARTIFICIAL INTELLIGENCE?

Artificial intelligence is a wide-ranging branch of computer science concerned with building smart machines capable of performing tasks that typically require human intelligence. Simply put, AI is the creation of software that imitates human behaviours and capabilities.

1.2.1 ARTIFICIAL INTELLIGENCE DEFINITION: BASICS OF AI

AI is intelligence demonstrated by machines, as opposed to the natural intelligence displayed by humans or animals.

Key workloads include:

- **Machine learning** - This is often the foundation for an AI system, and is the way we "teach" a computer model to make prediction and draw conclusions from data.
- **Anomaly detection** - The capability to automatically detect errors or unusual activity in a system.
- **Computer vision** - The capability of software to interpret the world visually through cameras, video, and images.
- **Natural language processing** - The capability for a computer to interpret written or spoken language, and respond in kind.
- **Knowledge mining** - The capability to extract information from large volumes of often unstructured data to create a searchable knowledge store.

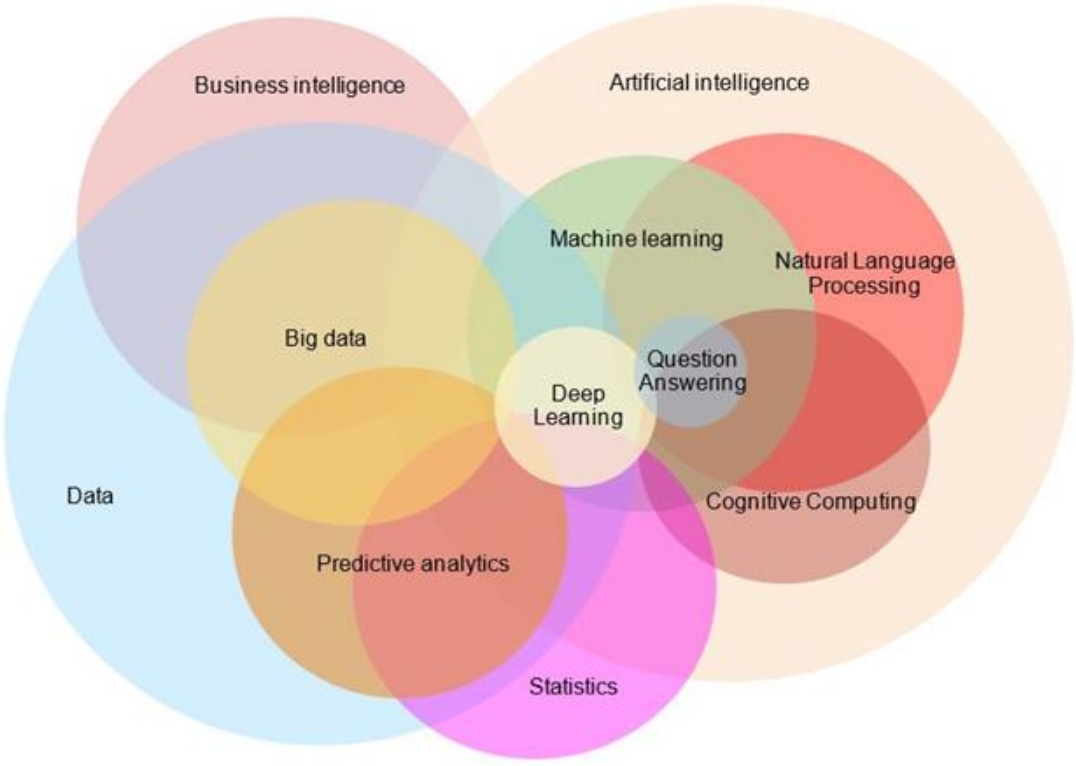

1.3 Use visual tools to create machine learning models with Azure Machine Learning:

The phrase “Every model is wrong, but some are useful” is especially true in Machine Learning. When developing machine learning models, you should always understand where it works as expected and where it fails miserably.

An important method of understanding a model is through model visualization. **Visualizing the model architecture** is important for:

- Model explain ability
- Results interpretation
- Model debugging
- Model comparison

Once you get some decent understanding of one model, you are good, right? No, it is wrong. !!!. Typically, you need to do some or a lot of experimenting with model improvement ideas, and visualizing differences between various ML experiments becomes crucial.

There are many methods that you can use to get that understanding:

- Look at evaluation metrics (also, you should know how to choose an evaluation metric for your problem)
- Look at performance charts like ROC, Lift Curve, Confusion Matrix, and others
- Look at learning curves to estimate overfitting
- Look at model predictions on best/worst cases
- Look how resource-intensive is model training and inference (they translate to serious costs and will be crucial to the business side of things)
- Use model interpretation tools and techniques to vet predictions

1.3.1 Tools for Machine Learning experiments visualization

1. **Neptune:** Neptune is a metadata store for MLOps, built for teams that run a lot of experiments. It gives you a single place to log, store, display, organize, compare, and query all your model-building metadata. Neptune is used for:

    - **Experiment tracking:** Log, display, organize, and compare ML experiments in a single place.
    - **Model registry:** Version, store, manage, and query trained models and model building metadata.
    - **Monitoring ML runs live:** Record and monitor model training, evaluation, or production runs live.

How can Neptune help you visualize experiments and models?

- Log any metadata type, including parameters and metrics, but also rich objects like images, video, audio, and interactive visualizations.
- Visualize logged metadata and analyze results in a preferred way – in the runs table, as charts, in dashboards, or in a folder structure.
- Compare hyper parameters and metrics across many runs with an intelligent compare table that highlights what was different.
- See how different parameters and con-figs affect the results, and debug and optimize models.
- Automatically monitor hardware utilization (GPU, CPU, memory).
- Register models and metadata associated with them.

2. **Weights & Biases:** Weights & Biases is a machine learning platform for developers to build better models faster. It lets you quickly track experiments, version and iterate on datasets, evaluate model performance, reproduce models, visualize results and spot regressions, and share findings with colleagues.

How can Weights & Biases help you visualize experiments and models?

- Monitor training runs information like loss, accuracy (learning curves)
- View histograms of weights and biases (no pun intended), or gradients
- Log rich objects like, charts, video, audio or interactive charts during training
- Use various comparison tools like tables showing auto-diffs, parallel coordinates plot and others
- Interactive prediction bounding box visualization for object detection models
- Interactive prediction masks visualization for semantic segmentation models
- Visualize live metrics like GPU and CPU utilization
- Build dataset dependency graphs
- Visualize parameter importance

3. **Comet:** Comet is an ML platform that helps data scientists track, compare, explain and optimize experiments and models across the model's entire lifecycle, i.e. from training to production. In terms of experiment tracking, data scientists can register datasets, code changes, experimentation history, and models. Comet is available for teams, individuals, academics, organizations, and anyone who wants to easily visualize experiments, facilitate work, and run experiments. It can be used as a hosted platform or deployed on-premise.

   Main advantages:

   - Fully-customizable experiment table within the web-based user interface;
   - Extensive comparison features—code, hyperparameters, metrics, predictions, dependencies, system metrics, and more
   - Dedicated modules for vision, audio, text, and tabular data that allow for easy identification of issues with the dataset.

4. **Sacred + Omniboard:** Sacred is open-source software that allows machine learning researchers to configure, organize, log, and reproduce experiments. Sacred does not come with its proper UI but there are a few dashboarding tools that you can connect to it, such as Omniboard (but you can also use others, such as Sacred board, or Neptune via integration). Sacred doesn't have the scalability of the previous tools and has not been adapted to team collaboration (unless integrated with another tool), however, it has great potential when it comes to individual research.

   Main advantages:

   - Possibility to connect it to the preferred UI;
   - Possibility to track any model training developed with any Python library;
   - Extensive experiment parameters customization options.

5. **MLflow:** MLflow is an open-source platform that helps manage the whole machine learning lifecycle. This includes experimentation, but also model storage, reproducibility, and deployment. Each of these four elements is represented by one MLflow component: Tracking, Model Registry, Projects, and Models. The MLflow Tracking component consists of an API and UI that support logging various metadata (including parameters, code versions, metrics, and output files) and later visualizing the results.

   Main advantages:

   - Focus on the whole lifecycle of the machine learning process;
   - Strong and big community of users that provide community support;
   - Open interface that can be integrated with any ML library or language.

1.4 Explore Computer vision

Computer vision is a field of artificial intelligence (AI) that enables computers and systems to derive meaningful information from digital images, videos and other visual inputs - and take actions or make recommendations based on that information. If AI enables computers to think, computer vision enables them to see, observe and understand. Computer vision works much the same as human vision, except humans have a head start. Human sight has the advantage of lifetimes of context to train how to tell objects apart, how far away they are, whether they are moving and whether there is something wrong in an image. Computer vision trains machines to perform these functions, but it must do it in much less time with cameras, data and algorithms rather than retinas, optic nerves and a visual cortex. Because a system trained to inspect products or watch a production asset can analyse thousands of products or processes a minute, noticing imperceptible defects or issues, it can quickly surpass human capabilities. Computer vision is used in industries ranging from energy and utilities to manufacturing and automotive – and the market is continuing to grow. It is expected to reach USD 48.6 billion by 2022.

1.4.1 How does computer vision work?

Computer vision is an area of artificial intelligence (AI) in which software systems are designed to perceive the world visually, though cameras, images, and video. There are multiple specific types of computer vision problem that AI engineers and data scientists can solve using a mix of custom machine learning models and platform-as-a-service (PaaS) solutions - including many cognitive services in Microsoft Azure.

Computer vision is one of the core areas of artificial intelligence (AI), and focuses on creating solutions that enable AI applications to "see" the world and make sense of it.

Of course, computers don't have biological eyes that work the way ours do, but they are capable of processing images; either from a live camera feed or from digital photographs or videos. This ability to process images is the key to creating software that can emulate human visual perception.

Some potential uses for computer vision include:

- **Content Organization**: Identify people or objects in photos and organize them based on that identification. Photo recognition applications like this are commonly used in photo storage and social media applications.
- **Text Extraction**: Analyze images and PDF documents that contain text and extract the text into a structured format.
- **Spatial Analysis**: Identify people or objects, such as cars, in a space and map their movement within that space.

To an AI application, an image is just an array of pixel values. These numeric values can be used as *features* to train machine learning models that make predictions about the image and its contents.

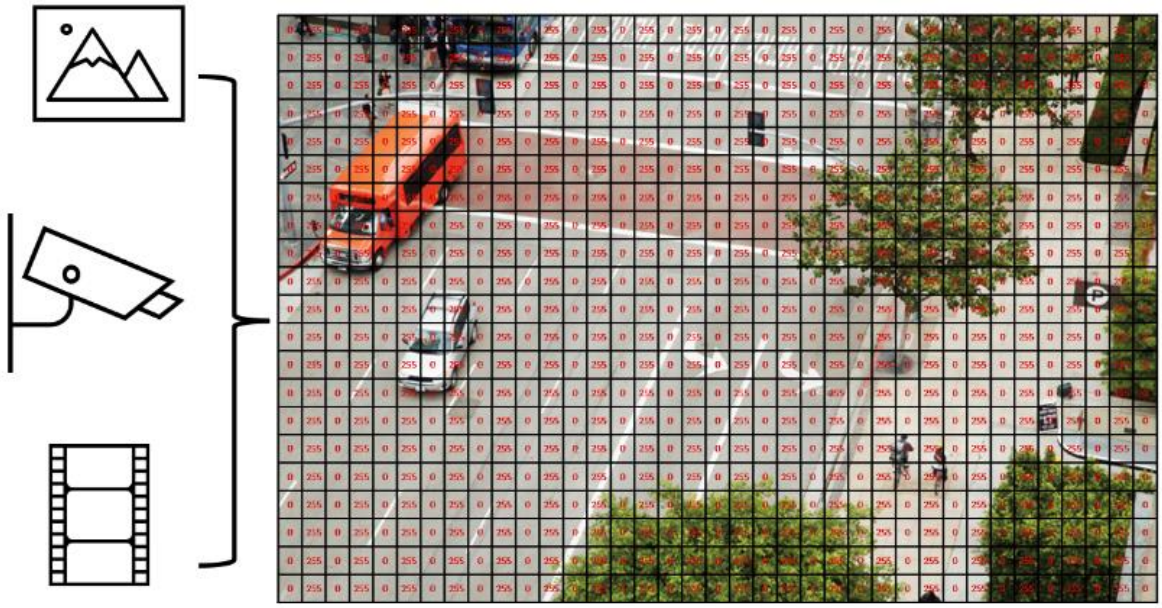

Training machine learning models from scratch can be very time intensive and require a large amount of data. Microsoft's Computer Vision service gives

you access to pre-trained computer vision capabilities.

Computer Vision models and capabilities

Most computer vision solutions are based on machine learning models that can be applied to visual input from cameras, videos, or images. The following table describes common computer vision tasks.

- **Image classification**: Image classification involves training a machine learning model to classify images based on their contents. For example, in traffic monitoring solution you might use an image classification model to classify images based on the type of vehicle they contain, such as taxis, buses, cyclists, and so on.
- **Object detection:** Object detection machine learning models are trained to classify individual objects within an image, and identify their location with a bounding box. For example, a traffic monitoring solution might use object detection to identify the location of different classes of vehicle.
- **Semantic segmentation**: Semantic segmentation is an advanced machine learning technique in which individual pixels in the image are classified according to the object to which they belong. For example, a traffic monitoring solution might overlay traffic images with "mask" layers to highlight different vehicles using specific colors.

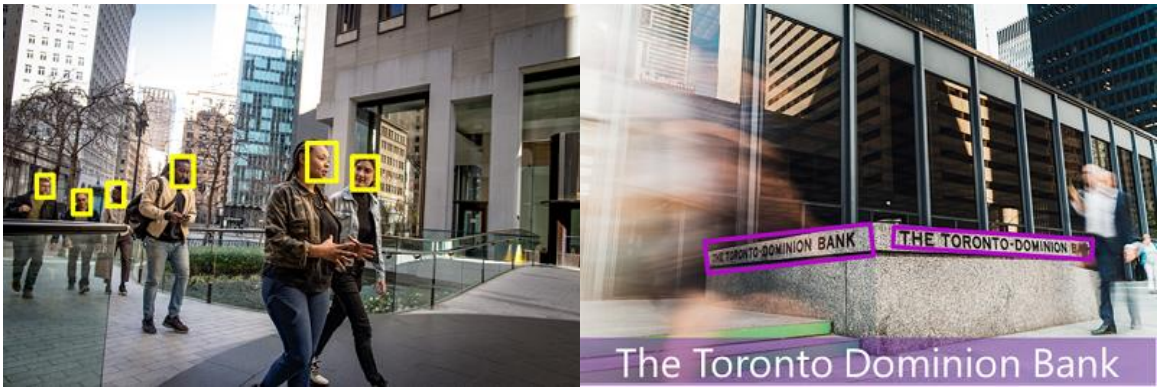

- **Image analysis**: You can create solutions that combine machine learning models with advanced image analysis techniques to extract information from images, including "tags" that could help catalogue the image or even descriptive captions that summarize the scene shown in the image.

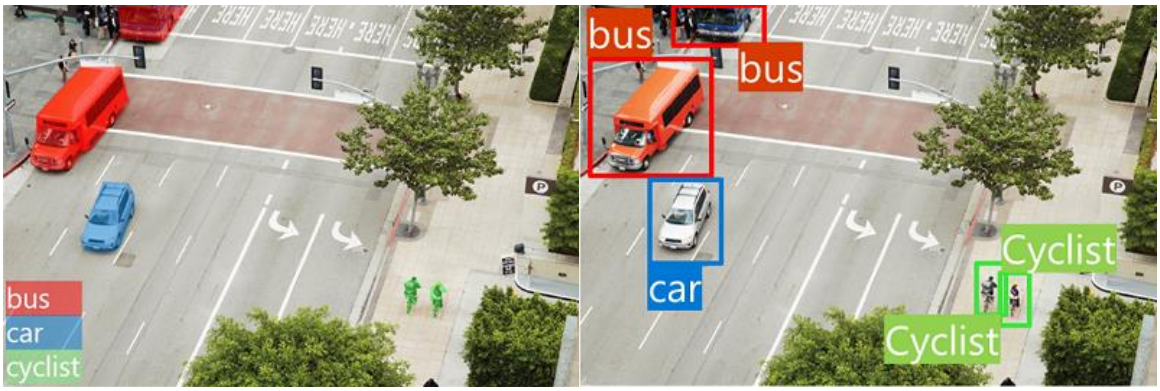

- **Face detection, analysis, and recognition**: Face detection is a specialized form of object detection that locates human faces in an image. This can be combined with classification and facial geometry analysis techniques to recognize individuals based on their facial features.
- **Optical character recognition (OCR):** Optical character recognition is a technique used to detect and read text in images. You can use OCR to read

text in photographs (for example, road signs or store fronts) or to extract information from scanned documents such as letters, invoices, or forms.

1.4.2 Explore Natural Language Processing

Natural language processing supports applications that can see, hear, speak with, and understand users. Using text analytics, translation, and language understanding services, Microsoft Azure makes it easy to build applications that support natural language.

Analyzing text is a process where you evaluate different aspects of a document or phrase, in order to gain insights into the content of that text. For the most part, humans can read some text and understand the meaning behind it. Even without considering grammar rules for the language the text is written in, specific insights can be identified in the text.

As an example, you might read some text and identify some key phrases that indicate the main talking points of the text. You might also recognize names of people or well-known landmarks such as the Eiffel Tower. Although difficult at times, you might also be able to get a sense for how the person was feeling when they wrote the text, also commonly known as sentiment.

Text Analytics Techniques

Text analytics is a process where an artificial intelligence (AI) algorithm, running on a computer, evaluates these same attributes in text, to determine specific insights. A person will typically rely on their own experiences and knowledge to achieve the insights. A computer must be provided with similar knowledge to be able to perform the task. There are some commonly used techniques that can be used to build software to analyze text, including:

- Statistical analysis of terms used in the text. For example, removing common "stop words" (words like "the" or "a", which reveal little semantic information about the text), and performing frequency analysis of the remaining words (counting how often each word appears) can provide clues about the main subject of the text.
- Extending frequency analysis to multi-term phrases, commonly known as N-grams (a two-word phrase is a bi-gram, a three-word phrase is a tri-gram, and so on).
- Applying stemming or lemmatization algorithms to normalize words before counting them - for example, so that words like "power", "powered", and "powerful" are interpreted as being the same word.
- Applying linguistic structure rules to analyze sentences - for example, breaking down sentences into tree-like structures such as a noun phrase, which itself contains nouns, verbs, adjectives, and so on.
- Encoding words or terms as numeric features that can be used to train a machine learning model. For example, to classify a text document based on the terms it contains. This technique is often used to perform sentiment analysis, in which a document is classified as positive or negative.
- Creating vectorized models that capture semantic relationships between words by assigning them to locations in n-dimensional space. This modeling technique might, for example, assign values to the words "flower" and "plant" that locate them close to one another, while "skateboard" might be given a value that positions it much further away.

While these techniques can be used to great effect, programming them can be complex. In Microsoft Azure, the Language cognitive service can help simplify application development by using pre-trained models that can:

- Determine the language of a document or text (for example, French or English).
- Perform sentiment analysis on text to determine a positive or negative sentiment.

- Extract key phrases from text that might indicate its main talking points.
- Identify and categorize entities in the text. Entities can be people, places, organizations, or even everyday items such as dates, times, quantities, and so on.

1.4.3 Explore conversational AI

Conversational AI is an umbrella term used to describe various methods of enabling computers to carry on a conversation with a human. This technology ranges from fairly simple natural language processing (NLP) to more sophisticated machine learning (ML) models that can interpret a much wider range of inputs and carry on more complex conversations.

One of the most common applications of conversational AI is in chat bots, which use NLP to interpret user inputs and carry on a conversation. Other applications include virtual assistants, customer service chatbots, and voice assistants.

Savvy consumers expect to communicate via mobile app, web, interactive voice response (IVR), chat, or messaging channels. They look for a consistent and enjoyable experience that's fast, easy, and personalized.

For businesses, the key to meeting and exceeding these expectations across channels and at scale is intelligent automation. Conversational artificial intelligence (AI) powers interactions that are near human, improving CX, boosting satisfaction, driving loyalty, and increasing customer lifetime value (LTV).

1.4.4.1 Components of Conversational AI

Conversational AI can be broken down into five core components. These five core components work together to enable a computer to understand and respond to human conversation:

1. **Natural language processing:** NLP is the ability of a computer to

understand human language and respond in a way that is natural for humans. This involves understanding the meaning of words and the structure of sentences, as well as being able to handle idiomatic expressions and slang.

NLP is made possible by machine learning, which is used to train computers to understand language. NLP algorithms use large data sets to learn how words are related to each other, and how they are used in different contexts.

2. **Machine learning:** Machine learning is a field of artificial intelligence that enables computers to learn from data without being explicitly programmed. Machine learning algorithms can automatically improve their performance as they are exposed to more data.

   Machine learning is used to train computers to understand language, as well as to recognize patterns in data. It is also used to create models of how different things work, including the human brain.

3. **Text analysis:** Text analysis is the process of extracting information from text data. This involves identifying the different parts of a sentence, such as the subject, verb, and object. It also includes identifying the different types of words in a sentence, such as nouns, verbs, and adjectives.

   Text analysis is used to understand the meaning of a sentence, as well as the relationships between different words. It is also used to identify the topic of a text, as well as the sentiment (positive or negative) of the text.

4. **Computer vision:** Computer vision is the ability of a computer to interpret and understand digital images. This involves identifying the different objects in an image, as well as the location and orientation of those objects. Computer vision is used to identify the contents of an image, as well as the relationships between different objects in the image. It is also used to interpret the emotions of people in photos, and to understand the context of a photo.
5. **Speech recognition:** Speech recognition is the ability of a computer to

understand human speech. This involves recognizing the different sounds in a spoken sentence, as well as the grammar and syntax of the sentence.

Speech recognition is used to convert spoken words into text, and to understand the meaning of the words. It is also used to interpret the emotions of people speaking in a video, and to understand the context of a conversation.

1.4.4.2 How Does Conversational AI Work?

Driven by underlying machine learning and deep neural networks (DNN), a typical conversational AI flow includes:

- An interface that allows the user to input text into the system or Automatic Speech Recognition (ASR), a user interface that converts speech into text.
- Natural language processing (NLP) to extract the user's intent from the text or audio input, and translate the text into structured data.
- Natural Language Understanding (NLU) to process the data based on grammar, meaning, and context; to comprehend intent and entity; and to act as a dialogue management unit for building appropriate responses.
- An AI model that predicts the best response for the user based on the user's intent and the AI model's training data. Natural Language Generation (NLG) infers from the above processes, and forms an appropriate response to interact with humans.

In many cases, the user interface, NLP, and AI model are all provided by the same provider, often a conversational AI platform provider. However, it's is also possible to use different providers for each of these components.

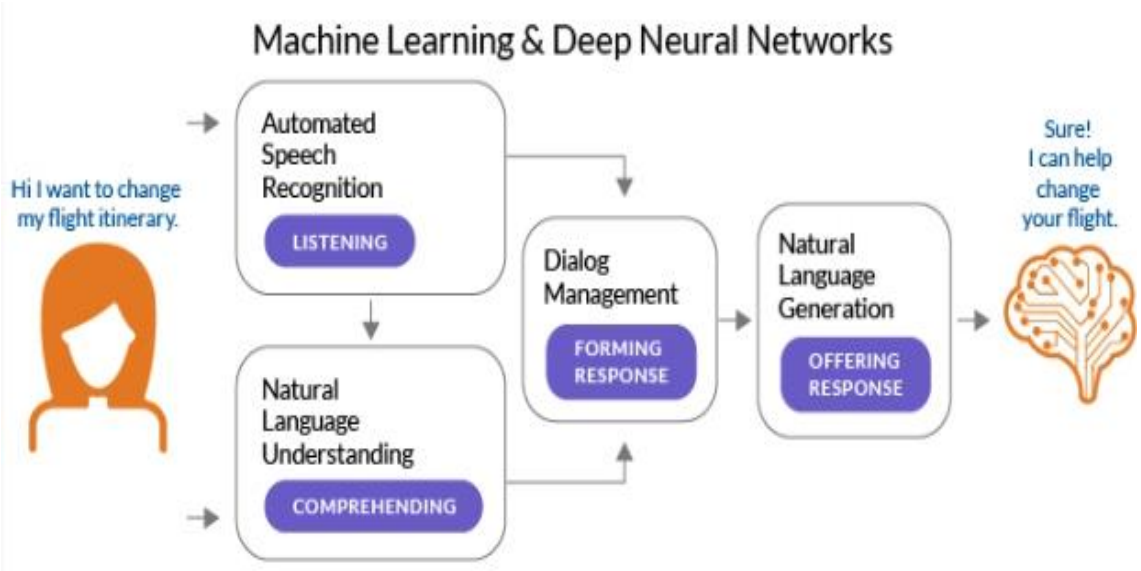

1.4.4.3 How to create Conversational AI?

There is no one-size-fits-all answer to this question, as the best way to create conversational AI depends on the specific needs and use cases of your organization. However, some tips on how to create conversational AI include:

1. **Start by understanding your use cases and requirements:** The first step in creating conversational AI understands your organization's specific needs and use cases. What are you trying to achieve with your chatbot? What type of conversations do you want it to be able to have? What data do you need to collect and track? Defining these requirements will help you determine the best approach to creating your chatbot.
2. **Choose the right platform and toolkit:** There are several different platforms and toolkits that you can use to create conversational AI. Each platform has its own strengths and weaknesses, so you need to choose the platform that best suits your needs. Some popular platforms include [24]7.ai Conversations, Microsoft Bot Framework, Amazon Lex, Google Dialogflow, and IBM Watson
3. **Build a prototype:** Once you have defined your requirements and chosen a platform, it's time to start building your prototype. Building a prototype will help you test your chatbot and iron out any kinks before deploying it to your users.
4. **Deploy and test your chatbot:** Once your prototype is finished, it's time

to deploy and test your chatbot. Make sure to test it with a small group of users first to get feedback and make any necessary adjustments.

5. **Optimize and improve your chatbot.:** The final step is to continually optimize and improve your chatbot. You can do this by tweaking the algorithms, adding new features, and collecting user feedback.

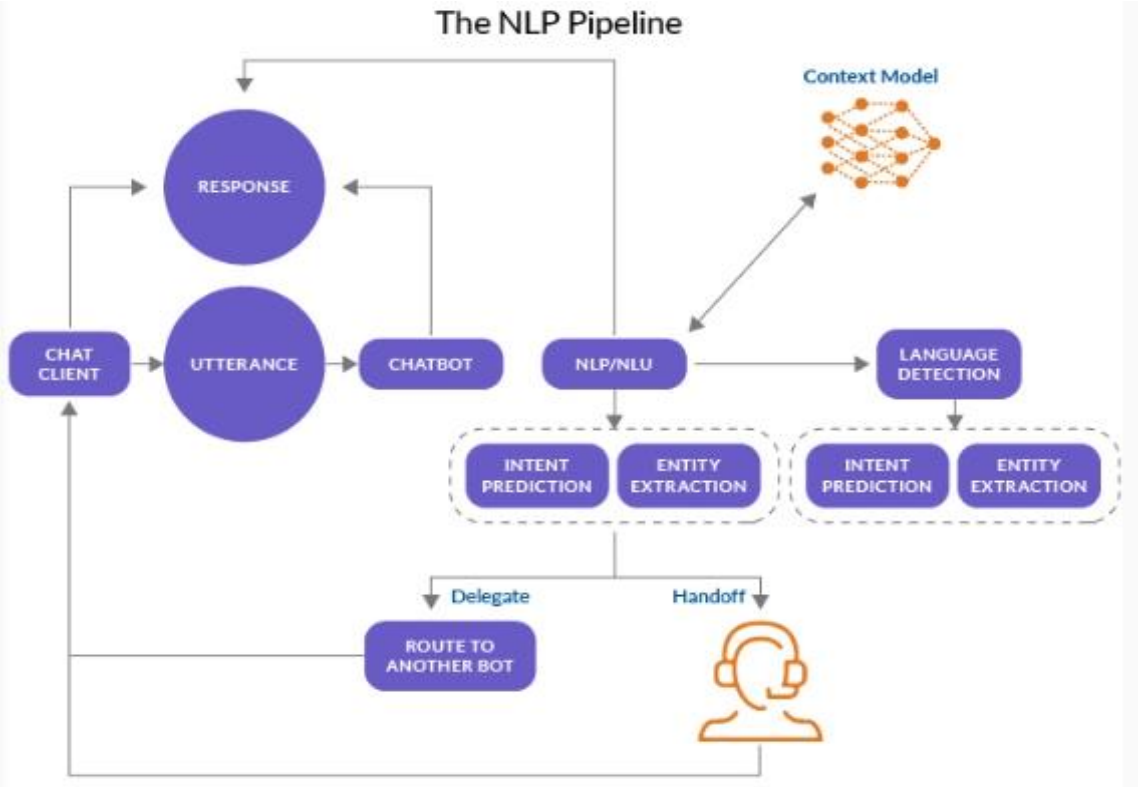

Difference between Traditional chatbots and AI – Powered Chatbots:

Traditional Chatbots	**AI-powered chatbots**	
Low complexity ▪ Basic answer and response machines ▪ Allow for simple integration ▪ Based on limited	**Focused, transactional** ▪ Can manage complex dialogues ▪ Integrate with multiple legacy/back-end systems	**Complex, contextual** ▪ Goes beyond conversations ▪ Contextually aware and intelligent ▪ Can self-learn and improve over time

scope ▪ Need explicit training for every scenario ▪ Require low back-end effort	▪ Based on large scope ▪ Sepecilize in completing tasks interacting with multiple systems ▪ Require high back-end effort	▪ Can anticipate user needs ▪ Require massive back-end effort

Conversational AI Challenges

Conversational AI's maturity has steadily increased over the past few years to the point where it can now deliver excellent business value and outcomes for companies. Nevertheless, challenges abound since this is also a fast-evolving conversational commerce category where very few vendors are constantly innovating and bringing new technologies to the market. Challenges include:

- Developing natural language processing (NLP) capabilities that can understand and interpret human interactions. This is a complex task that requires significant effort and investment in research and development.
- Understanding the context of a conversation in order to provide accurate responses. This can be particularly challenging in conversations that involve multiple people or multiple topics.
- The need for sophisticated design and development efforts to create a customer experience that engages users and keeps them engaged in the conversation.
- Deploying and integrating a Conversational AI solution into an existing business or application can be a significant challenge. Proper planning and execution are essential to ensure a successful deployment.
- Ensuring the security and privacy of data exchanged via multiple conversational AI-powered channels—this applies to any CX-related information exchange. Compliance with standards such as GDPR, CCCP, and other country-specific regulations is also critical.

- As conversational AI permeates global CX platforms, local language support becomes a high priority. Leading brands operating worldwide can't rely on availability in just one language to meet local needs at scale. Building a robust conversational AI platform to operate in regional languages, dialects, slang, noisy environments, with crosstalk, etc., is a huge challenge.
- Dialogue management and conversation design are non-trivial parts of conversational AI. Annotating the intelligence gathered from real agent conversations and building the right model-training data requires ongoing human-in-the-loop expertise.
- Building a conversational AI-based application that takes into consideration intent, entity extraction, sentiment analysis, and empathy is challenging and very few vendors offer solutions with these features.
- Explainable AI—not all conversational AI platforms use this data science tool, which eliminates algorithmic black boxes and helps answer the "why" within the model's functionality. Explainable AI also improves trust in the platform's ability to produce accurate, fair, and transparent results.
- Keeping automated conversations relevant can also be a real challenge, with customer needs and preferences changing faster than ever before. As a result, you may need people with coding skills, multiple-persona models, or IT input, making the solution more expensive. Conversational AI platforms that have no-code/low-code self-serve capabilities can enable business users to build and deploy voice and digital bots and context-aware conversational flows in just a few days.

State-of-the-Art Conversational AI

Technology behind conversational bot experiences is based on the latest advances in artificial intelligence, NLP, sentiment analysis, deep learning, and intent prediction. Together, these features encourage engagement, improve customer experience and agent satisfaction, accelerate time to resolution, and

grow business value.

Natural Language Processing (NLP)

Most conversational AI uses NLU to intelligently process user inputs against multiple models, enabling a bot to respond in a more human-like way to non-transactional journeys. The core technology understands slang, local nuances, colloquial speech, and can be trained to emulate different tones by using AI-powered speech synthesis.

Sentiment Analysis

This leading conversational AI technology layer abstracts pre-built sentiment and social models to prioritize and seamlessly escalate to an agent when it detects that a customer needs expert advice. Sentiment detection will recognize, for example, an upset customer and immediately route them to an agent. You can also prioritize unhappy customers in the system, placing them in special queues or offering exceptional services.

Deep Learning

This machine learning technique is inspired by the human brain or 'neural network' and allows AI to learn by association, just like a child. The more data AI is exposed to, the better it gets—and the more accurately it can respond over time. AI models trained with many years of contact center data from various voice and digital channels result in smarter and more accurate responses to human inquiries. Response accuracy can be further improved over time by learning from interactions between customers, chatbots, and human agents, and optimizing intent models using AI-powered speech synthesis.

Intent Prediction

Using behavioral analysis and tagging activities, conversational AI technologies can understand the true meaning behind each consumer's

request. Knowing intent allows companies to deliver the right response at the right moment through an automated bot or human agent.

The future roadmap for conversational AI platforms includes support for multiple use cases, multi-domain, and multiple vertical needs, along with explainable AI.

The Conversational AI Vendor Market

According to Gartner, over 1500 conversational AI providers now offer various levels of capability, language support, use-cases, and business models. Sophistication swings widely depending on what's supported, such as:

- Number of integrations with back-end systems such as CRM
- Number and type of channels (voice, text-based chatbots, messaging, etc.).
- Customization of chatbots and virtual assistants for vertical specific use cases and applications for faster adoption into production
- Number of languages, slang, dialects, local lingo, shorthand, phonetic spelling, grammatical structures, intents, entity, etc.

Horizontal solutions are the most flexible and controllable but take longer to implement, while vertical specific ones come with pre-built capabilities that are a better fit for a specialized use cases in a target domain. Vendors that offer vertical solutions built on an established horizontal platform give companies full flexibility in customizing to meet their precise needs.

How to pick the right Conversational AI Solution?

When it comes to selecting a conversational AI solution, there are a few key factors to consider.

1. First, consider the needs of your business. What questions or tasks do your customers commonly ask or need help with? What areas of your business could benefit from automation?
2. Next, evaluate the capabilities of different conversational AI solutions. Some platforms are better suited for specific tasks or industries, while others are more versatile.
3. Finally, consider the cost and complexity of implementing different solutions. Some platforms are more expensive or require more technical expertise to set up and use.

Once you have a better understanding of your business needs and the capabilities of different conversational AI solutions, you can begin to narrow down your options and select the right platform for your business.

1.5 Tune Model Hyperparameters – Azure Machine Learning

Here we will see how to use the Tune Model Hyper parameters component in Azure Machine Learning designer. The goal is to determine the optimum hyper parameters for a machine learning model. The component builds and tests multiple models by using different combinations of settings. It compares metrics over all models to get the combinations of settings.

The terms parameter and hyper parameter can be confusing. The model's parameters are what you set in the right pane of the component. Basically, this component performs a parameter sweep over the specified parameter settings. It learns an optimal set of hyper parameters, which might be different for each specific decision tree, dataset, or regression method. The process of finding the optimal configuration is sometimes called tuning.

The component supports the following method for finding the optimum settings for a model: integrated train and tune. In this method, you configure a set of parameters to use. You then let the component iterate over multiple combinations. The component measures accuracy until it finds a "best" model. With most learner components, you can choose which parameters

should be changed during the training process, and which should remain fixed.

Depending on how long you want the tuning process to run, you might decide to exhaustively test all combinations. Or you might shorten the process by establishing a grid of parameter combinations and testing a randomized subset of the parameter grid.

How to configure Tune Model Hyper parameters

Learning the optimal hyper parameters for a machine learning model requires considerable use of pipelines

Train a model by using a parameter sweep

This section describes how to perform a basic parameter sweep, which trains a model by using the Tune Model Hyper parameters component.

1. Add the Tune Model Hyper parameters component to your pipeline in the designer.
2. Connect an untrained model to the leftmost input.
3. Add the dataset that you want to use for training, and connect it to the middle input of Tune Model Hyper parameters.

   Optionally, if you have a tagged dataset, you can connect it to the rightmost input port (Optional validation dataset). This lets you measure accuracy while training and tuning.

4. In the right panel of Tune Model Hyper parameters, choose a value for Parameter sweeping mode. This option controls how the parameters are selected.

   - **Entire grid**: When you select this option, the component loops over a grid predefined by the system, to try different combinations and identify the best learner. This option is useful when you don't know what the

best parameter settings might be and want to try all possible combinations of values.

- **Random sweep**: When you select this option, the component will randomly select parameter values over a system-defined range. You must specify the maximum number of runs that you want the component to execute. This option is useful when you want to increase model performance by using the metrics of your choice but still conserve computing resources.

5. For Label column, open the column selector to choose a single label column.
6. Choose the number of runs:

   - **Maximum number of runs on random sweep**: If you choose a random sweep, you can specify how many times the model should be trained, by using a random combination of parameter values.

7. For Ranking, choose a single metric to use for ranking the models.

   When you run a parameter sweep, the component calculates all applicable metrics for the model type and returns them in the **Sweep results** report. The component uses separate metrics for regression and classification models.

   However, the metric that you choose determines how the models are ranked. Only the top model, as ranked by the chosen metric, is output as a trained model to use for scoring.

8. For **Random seed**, enter an integer number as a pseudo random number generator state used for randomly selecting parameter values over a pre-defined range. This parameter is only effective if **Parameter sweeping mode** is **Random sweep**.
9. Submit the pipeline.

Results of hyper parameter tuning

When training is complete:

- To view the sweep results, you could either right-click the component, and then select Visualize, or right-click left output port of the component to visualize.

  The Sweep results includes all parameter sweep and accuracy metrics that apply to the model type, and the metric that you selected for ranking determines which model is considered "best."

- To save a snapshot of the trained model, select the Outputs+logs tab in the right panel of the Train model component. Select the Register dataset icon to save the model as a reusable component.

Train a model by using a parameter sweep

This section describes how to perform a basic parameter sweep, which trains a model by using the Tune Model Hyper parameters component.

1. Add the Tune Model Hyper parameters component to your pipeline in the designer.
2. Connect an untrained model to the leftmost input.

Neural Network Regression:

Let us discuss about a component in Azure Machine Learning designer. This component is used to create regression model using a customizable neural network algorithm.

Although neural networks are widely known for use in deep learning and modelling complex problems such as image recognition, they are easily adapted to regression problems. Any class of statistical models can be termed a neural network if they use adaptive weights and can approximate non-linear

functions of their inputs. Thus, neural network regression is suited to problems where a more traditional regression model cannot fit a solution.

Neural network regression is a supervised learning method, and therefore requires a tagged dataset, which includes a label column. Because a regression model predicts a numerical value, the label column must be a numerical data type. The model can be trained by providing a model and the tagged dataset as an input to train the model. The trained model can then be used to predict values for the new input examples.

Configure Neural Network Regression:

- **Create a neural network model using the default architecture:** If you accept the default neural network architecture, use the Properties pane to set parameters that control the behaviour of the neural network, such as the number of nodes in the hidden layer, learning rate, and normalization. Start here if you are new to neural networks. The component supports many customizations, as well as model tuning, without deep knowledge of neural networks.
- **Define a custom architecture for a neural network:** Use this option if you want to add extra hidden layers, or fully customize the network architecture, its connections, and activation functions. This option is best if you are already somewhat familiar with neural networks. You use the Net# language to define the network architecture.

1. **Create a neural network model using the default architecture:**

- Add the Neural Network Regression component to your pipeline in the designer. You can find this component under Machine Learning, Initialize, in the Regression category.
- Indicate how you want the model to be trained, by setting the Create trainer mode option.
  - **Single Parameter**: Choose this option if you already know how you want to configure the model.
  - **Parameter Range:** Select this option if you are not sure of the best

parameters, and want to run a parameter sweep. Select a range of values to iterate over, and the Tune Model Hyper parameters iterates over all possible combinations of the settings you provided to determine the hyper parameters that produce the optimal results.

- In Hidden layer specification, select fully connected case. This option creates a model using the default neural network architecture, which for a neural network regression model, has these attributes:
  - The network has exactly one hidden layer.
  - The output layer is fully connected to the hidden layer and the hidden layer is fully connected to the input layer.
  - The number of nodes in the hidden layer can beset by the user (default value is 100).

    Because the number of nodes in the input layer is determined by the number of features in the training data, in a regression model there can be only one node in the output layer.

- For Number of hidden nodes, type the number of hidden nodes. The default is one hidden layer with 100 nodes. (This option is not available if you define a custom architecture using Net#.)
- For learning rate, type a value that defines the step taken a teach iteration, before correction. A larger value for learning rate can cause the model to converge faster, but it can overshoot local minima.
- For Number of learning iterations, specify the maximum number of times the algorithm processes the training cases.
- For The momentum, type a value to apply during learning as a weight on nodes from previous iterations.
- Select the option, Shuffle examples, to change the order of cases between iterations. If you deselect this option, cases are processed in exactly the same order each time you run the pipeline.
- For Random number seed, you can optionally type a value to use as the seed. Specifying a seed value is useful when you want to ensure repeatability across runs of the same pipeline.
- Connect training data set and train the model:

- If you set Create trainer mode to Single Parameter, connect a tagged data set and the Train Model component.
- If you set Create trainer mode to Parameter Range, connect a tagged data set and train the model by using Tune Model Hyper parameters.

- Submit the Pipeline

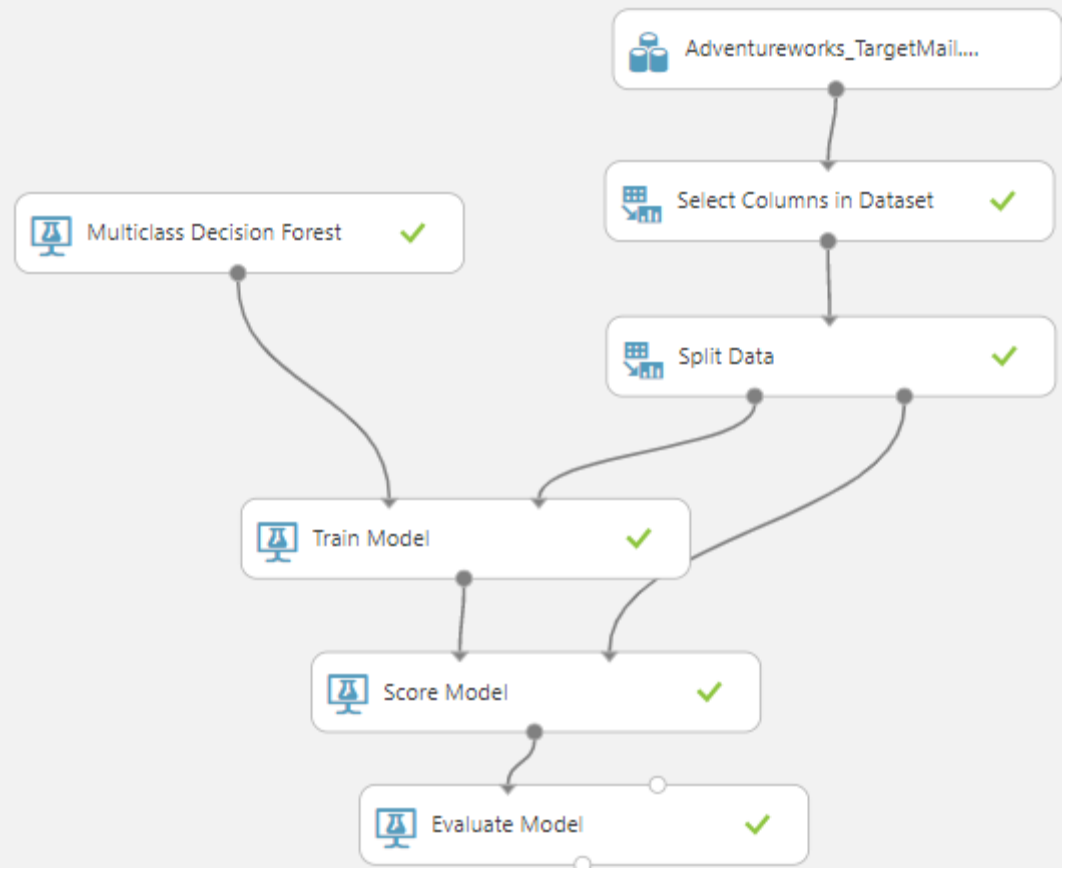

Course 2

Data Analyst Associate (DA -100)

Introduction to Microsoft data analytics

Before data can be used to tell a story, it must be run through a process that makes it usable in the story. Data analysis is the process of identifying, cleaning, transforming, and modelling data to discover meaningful and useful

information. The data is then crafted into a story through reports for analysis to support the critical decision-making process.

As the world becomes more data-driven, storytelling through data analysis is becoming a vital component and aspect of large and small businesses. It is the reason that organizations continue to hire data analysts.

Data-driven businesses make decisions based on the story that their data tells, and in today's data-driven world, data is not being used to its full potential, a challenge that most businesses face. Data analysis is, and should be, a critical aspect of all organizations to help determine the impact to their business, including evaluating customer sentiment, performing market and product research, and identifying trends or other data insights.

While the process of data analysis focuses on the tasks of cleaning, modelling, and visualizing data, the concept of data analysis and its importance to business should not be understated. To analyse data, core components of analytics are divided into the following categories:

- Descriptive
- Diagnostic
- Predictive
- Prescriptive
- Cognitive

- **Descriptive analytics:** Descriptive analytics help answer questions about what has happened based on historical data. Descriptive analytics techniques summarize large datasets to describe outcomes to stakeholders. By developing key performance indicators (KPIs), these strategies can help track the success or failure of key objectives. Metrics such as return on investment (ROI) are used in many industries, and specialized metrics are developed to track performance in specific industries.An example of descriptive analytics is generating reports to provide a view of an organization's sales and financial data.
- **Diagnostic analytics:** Diagnostic analytics help answer questions about why events happened. Diagnostic analytics techniques supplement basic

descriptive analytics, and they use the findings from descriptive analytics to discover the cause of these events. Then, performance indicators are further investigated to discover why these events improved or became worse. Generally, this process occurs in three steps:

1. Identify anomalies in the data. These anomalies might be unexpected changes in a metric or a particular market.
2. Collect data that's related to these anomalies.
3. Use statistical techniques to discover relationships and trends that explain these anomalies.

- **Predictive analytics:** Predictive analytics help answer questions about what will happen in the future. Predictive analytics techniques use historical data to identify trends and determine if they're likely to recur. Predictive analytical tools provide valuable insight into what might happen in the future. Techniques include a variety of statistical and machine learning techniques such as neural networks, decision trees, and regression.
- **Prescriptive analytics:** Prescriptive analytics help answer questions about which actions should be taken to achieve a goal or target. By using insights from prescriptive analytics, organizations can make data-driven decisions. This technique allows businesses to make informed decisions in the face of uncertainty. Prescriptive analytics techniques rely on machine learning as one of the strategies to find patterns in large datasets. By analysing past decisions and events, organizations can estimate the likelihood of different outcomes.
- **Cognitive analytics:** Cognitive analytics attempt to draw inferences from existing data and patterns, derive conclusions based on existing knowledge bases, and then add these findings back into the knowledge base for future inferences, a self-learning feedback loop. Cognitive analytics help you learn what might happen if circumstances change and determine how you might handle these situations. Inferences aren't structured queries based on a rules database; rather, they're unstructured hypotheses that are gathered from several sources and expressed with varying degrees of confidence. Effective cognitive analytics depend on machine learning algorithms, and will use several natural language processing concepts to make sense of

previously untapped data sources, such as call centre conversation logs and product reviews.

2. Prepare data for analysis

Your manager wants to see a report on your latest sales and profit figures by the end of the day. But the latest data is in files on your laptop. In the past, it's taken hours to create a report, and you're beginning to feel anxious. With Power BI, you can create a stunning report and share it in Microsoft Teams in no time!

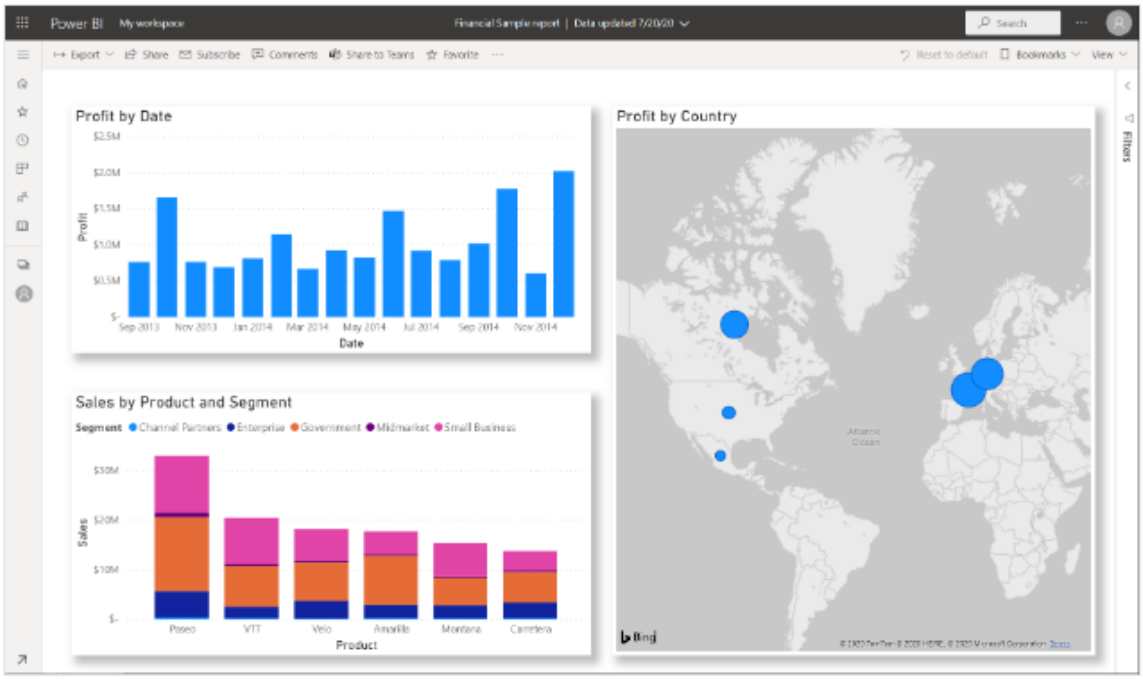

Here, we upload an Excel file, create a new report, and share it with colleagues in Microsoft Teams, all from within Power BI. You'll learn how to:

- Prepare your data in Excel.
- Download sample data.
- Build a report in the Power BI service.
- Pin the report visuals to a dashboard.
- Share a link to the dashboard.
- Share the dashboard in Microsoft Teams

Prerequisites

- Sign up for the Power BI service.
- Download the Financial Sample workbook and save it your computer or to OneDrive for Business.

PRACTICAL

Prepare data in Excel

Let's take a simple Excel file as an example.

1. Before you can load your Excel file into Power BI, you must organize your data in a flat table. In a flat table, each column contains the same data type; for example, text, date, number, or currency. Your table should have a header row, but not any columns or rows that display totals.

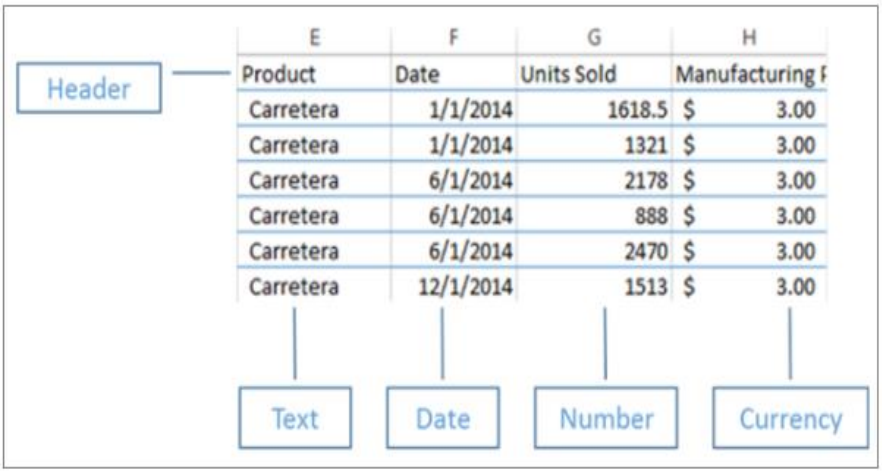

1. Next, format your data as a table. In Excel, on the **Home** tab, in the **Styles** group, select **Format as Table**.
2. Select a table style to apply to your worksheet. Your Excel worksheet is now ready to load into Power BI.

Segment	Country	Product
Government	Canada	Aliqui UR
Channel Partners	Canada	Aliqui UR
Small Business	Canada	Barba UM
Midmarket	Canada	Barba UM
Enterprise	Canada	Barba UM
Government	Canada	Barba UM
Midmarket	France	Fama UE
Enterprise	France	Fama UE
Enterprise	Germany	Fama UE
Small Business	Germany	Pirum RP

Upload your Excel file to the Power BI service

The Power BI service connects to many data sources, including Excel files that live on your computer.

3. To get started, sign in to the Power BI service. If you haven't signed up, you can do so for free.

2. In **My workspace**, select **New** > **Upload a file**.

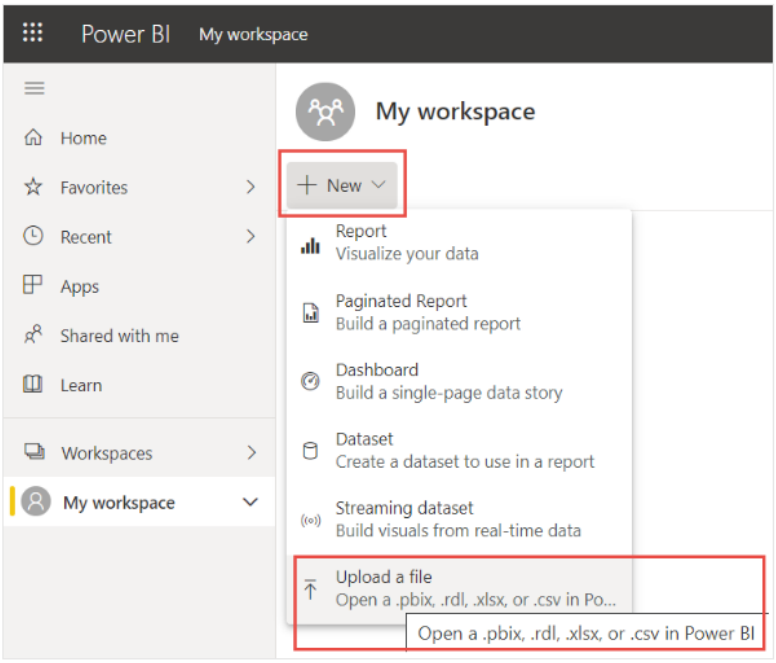

1. Select **Local File**, browse to where you saved the Financial Sample Excel file, and select **Open**.
2. On the **Local File** page, select **Import**: Now you have a Financial

Sample dataset. Power BI also automatically created a blank dashboard. If you don't see the dashboard, refresh your browser.

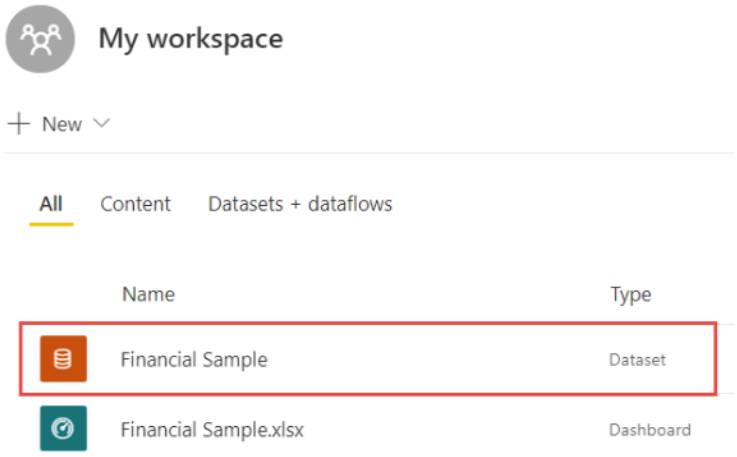

3. You want to create a report. Still in **My workspace**, select **New** > **Report**.

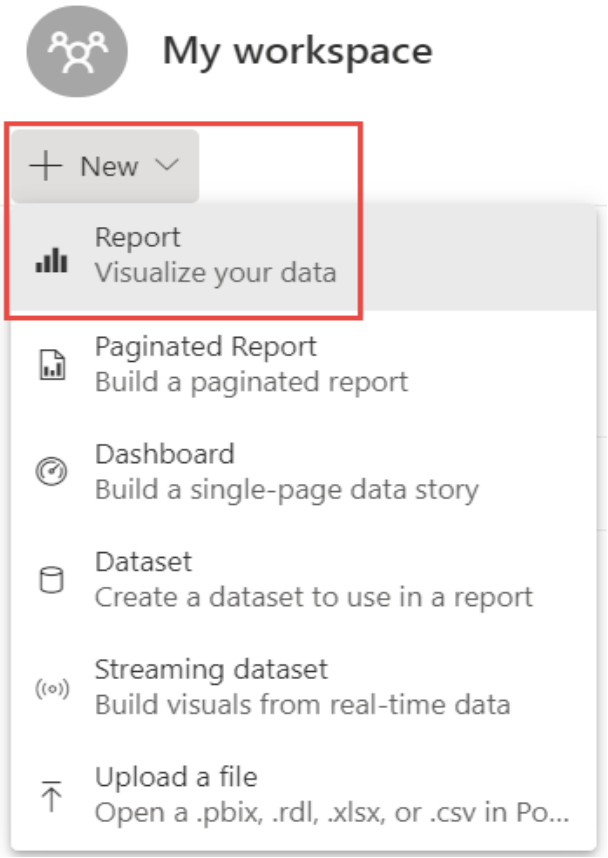

4. In the **Select a dataset to create a report** dialog box, select your **Financial Sample** dataset > **Create**.

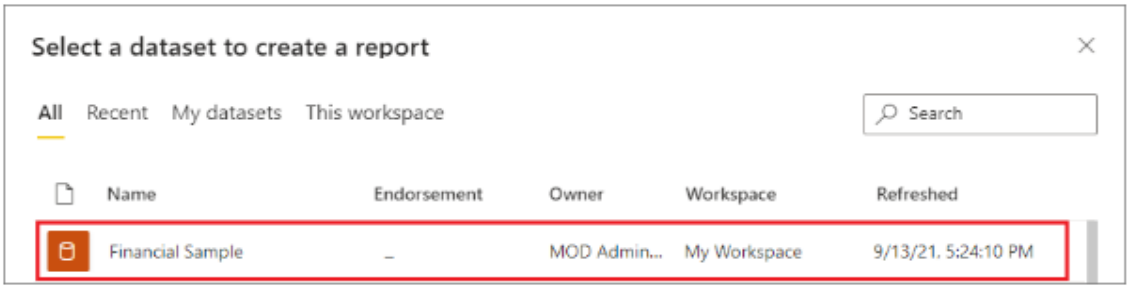

Build your report: The report opens in editing view and displays the blank report canvas. On the right are the **Visualizations**, **Filters**, and **Fields** panes. Your Excel workbook table data appears in the **Fields** pane. At the top is the name of the table, **financials**. Under that, Power BI lists the column headings as individual fields. You see the Sigma symbols in the Fields list? Power BI has detected that those fields are numeric. Power BI also indicates a geographic field with a globe symbol.

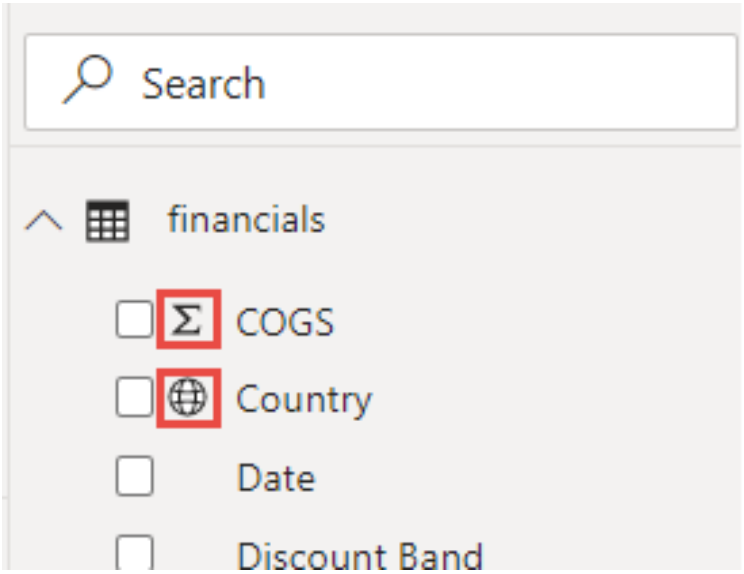

To have more room for the report canvas, select **Hide the navigation pane**, and minimize the **Filters** pane.

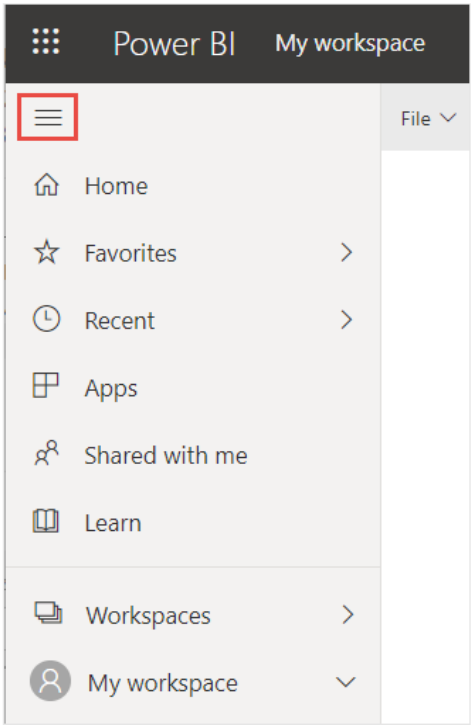

Now you can begin to create visualizations. Let's say your manager wants to see profit over time. In the **Fields** pane, drag **Profit** to the report canvas. By default, Power BI displays a column chart with one column.

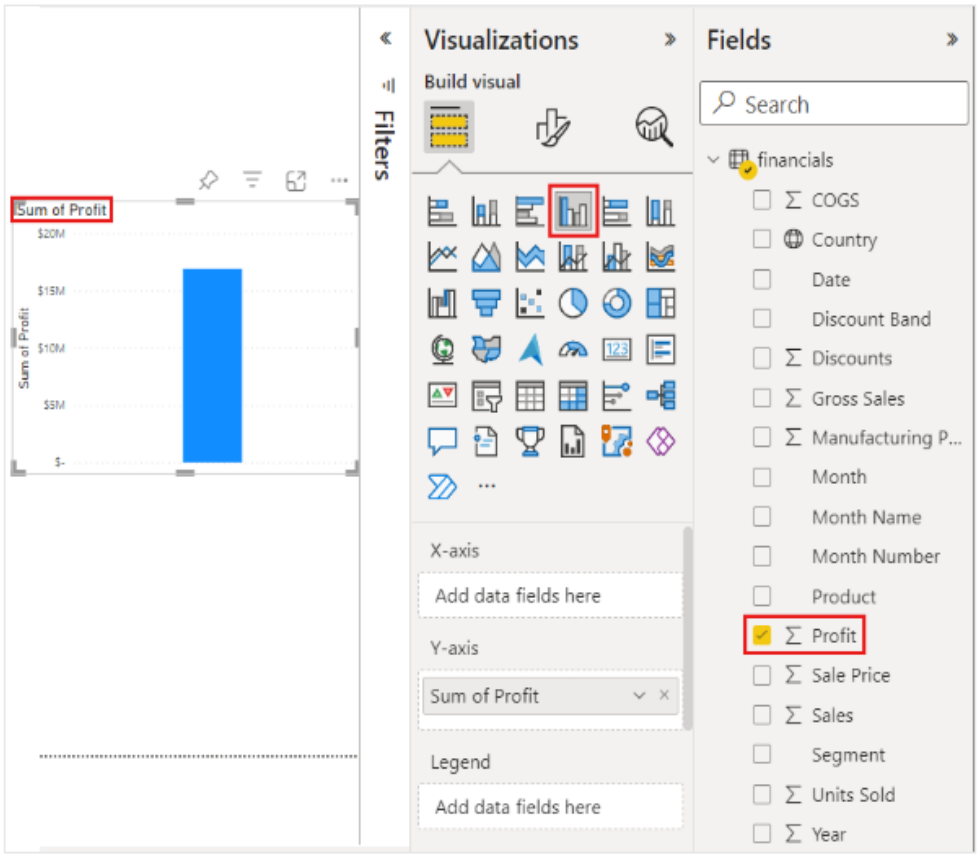

- Drag **Date** to the report canvas.
- Power BI updates the column chart to show profit by date.

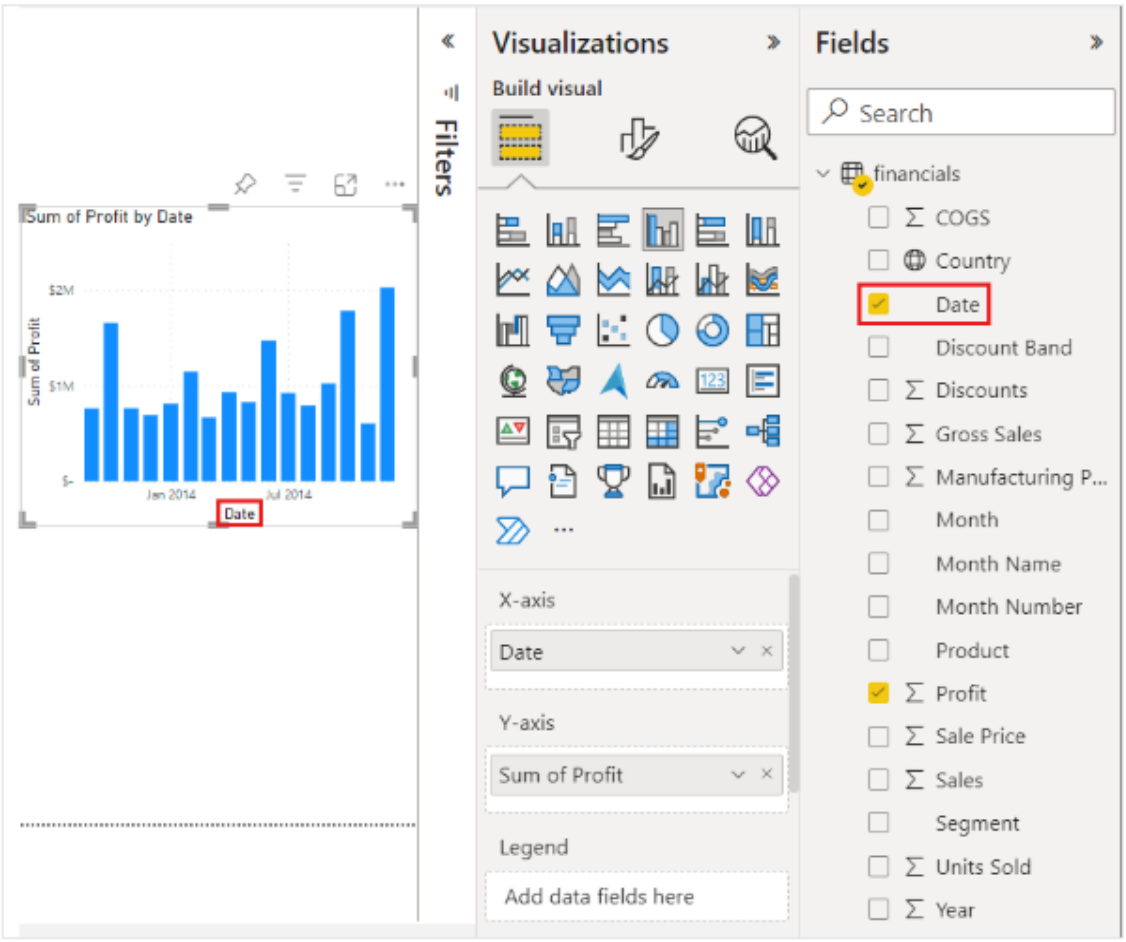

December 2014 was the most profitable month.

4. **Model data in Power BI:** Often, you'll connect to multiple data sources to create your reports. All that data needs to work together to create a cohesive report. *Modelling* is how to get your connected data ready for use.

Tasks in this module:

- Create relationships between your data sources
- Create a new field with calculated columns

- Optimize data by hiding fields and sorting visualization data
- Create a measure to perform calculations on your data
- Use a calculated table to create a relationship between two tables
- Format time-based data so that you can drill down for more details

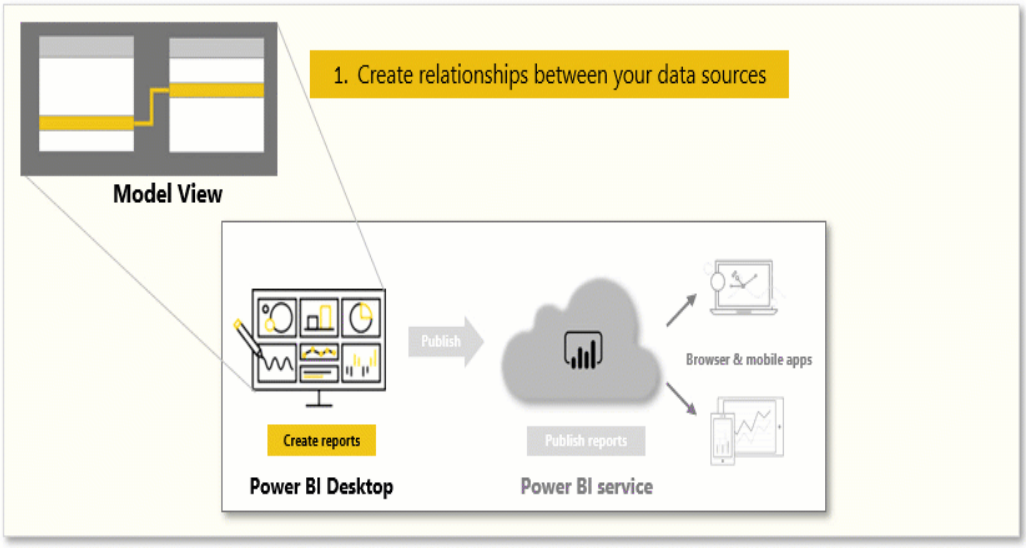

Fig: Create relationships between your data sources

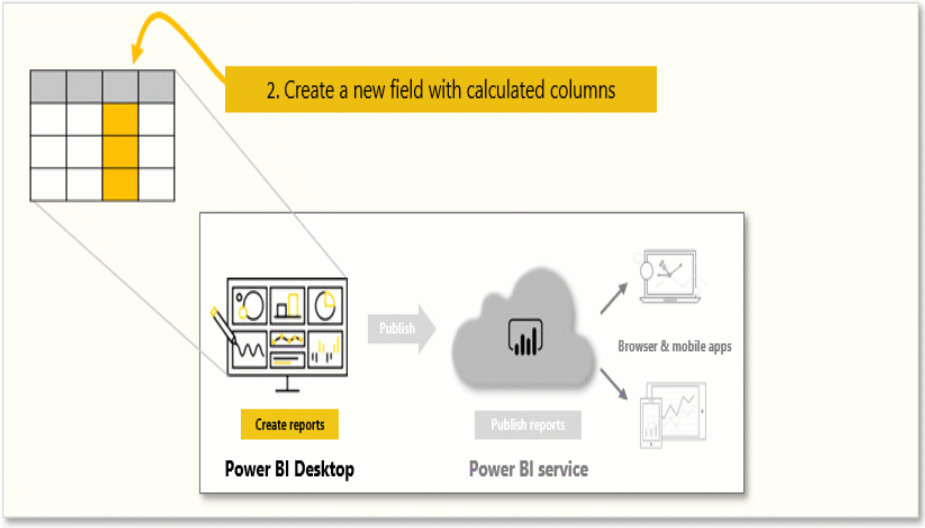

Fig: Create a new field with calculated columns

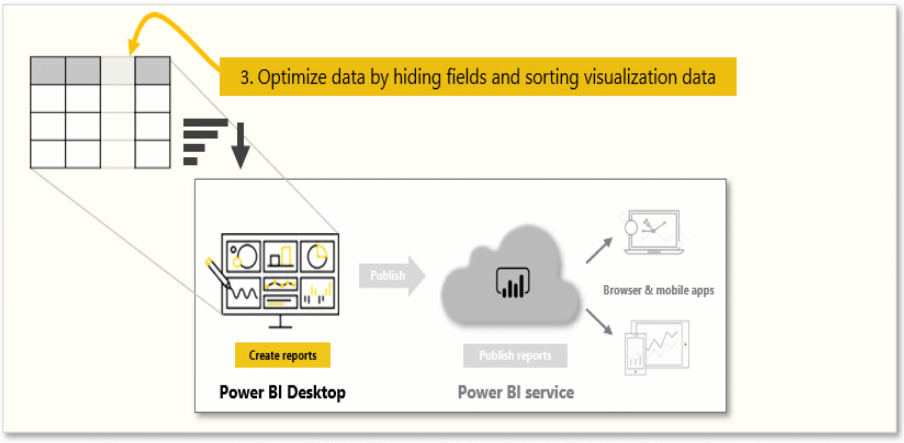

Fig: Optimize data by hiding fields and sorting visualization data

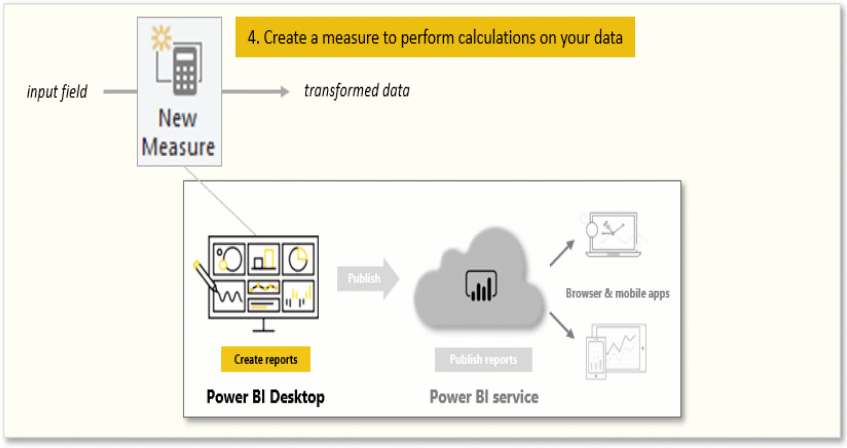

Fig: Create a measure to perform calculations on your data

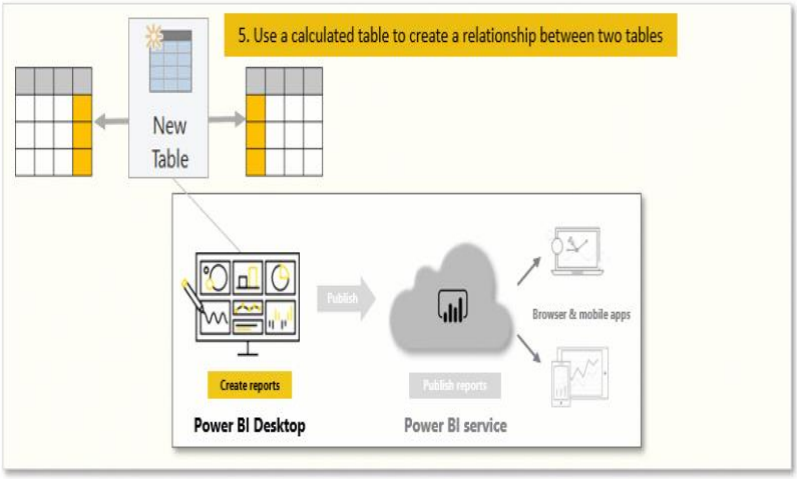

Fig: Use a calculated table to create a relationship between two tables

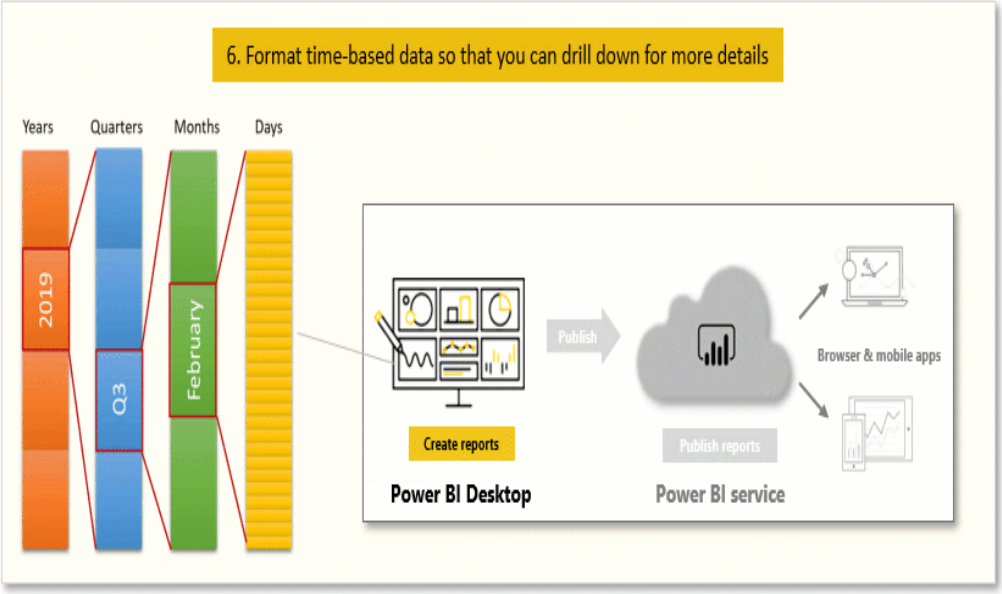

Fig: Format time-based data so that you can drill down for more details

In Power BI, you can create a *relationship* to create a logical connection between different data sources. A relationship enables Power BI to connect tables to one another so that you can create visuals and reports. This module describes data-centric relationships and how to create relationships when none exists.

5. Visualize data in Power BI

Visualizations are used to effectively present your data and are the basic building blocks of any Business Intelligence tool. Power BI contains various default data visualization components that include simple bar charts to pie charts to maps, and also complex models such as waterfalls, funnels, gauges, and many other components.

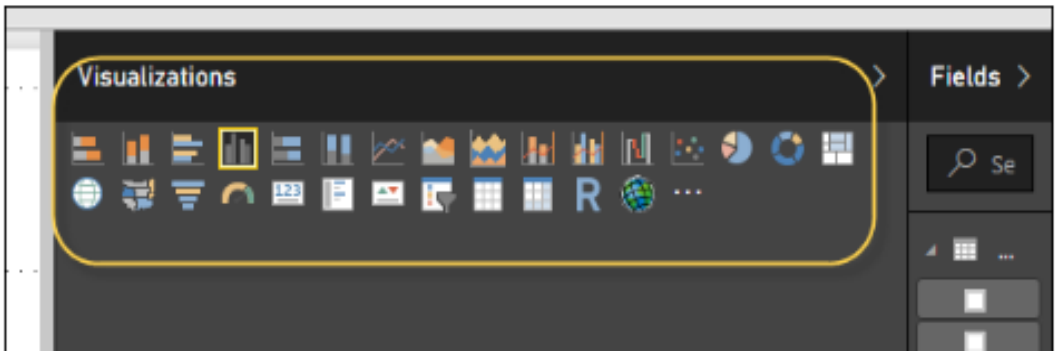

In Power BI, you can create visualization in two ways. First is by adding from the right-side pane to Report Canvas. By default, it is the table type visualization, which is selected in Power BI. Another way is to drag the fields from right side bar to the axis and value axis under Visualization. You can add multiple fields to each axis as per the requirement.

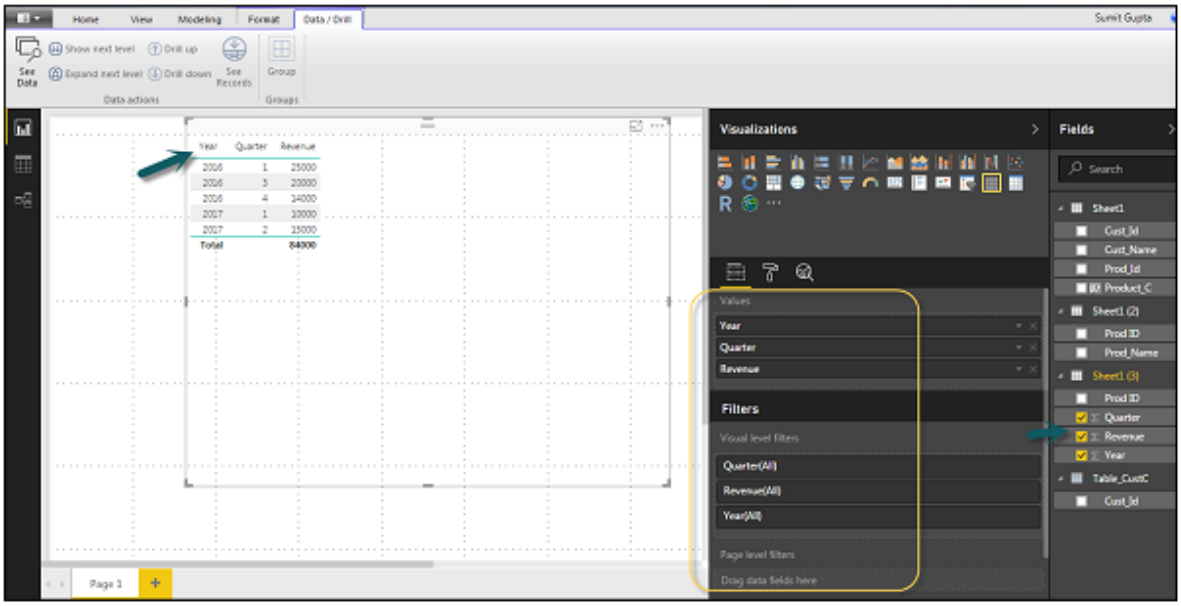

In Power BI, it is also possible to move your visualization on the reporting canvas by clicking and then dragging it. You can also switch between different type of charts and visualizations from the Visualization pane. Power BI attempts to convert your selected fields to the new visual type as closely as possible.

Creating Map Visualizations

In Power BI, we have two types of map visualization - bubble maps and shape maps. If you want to create a bubble map, select the map option from the visualization pane.

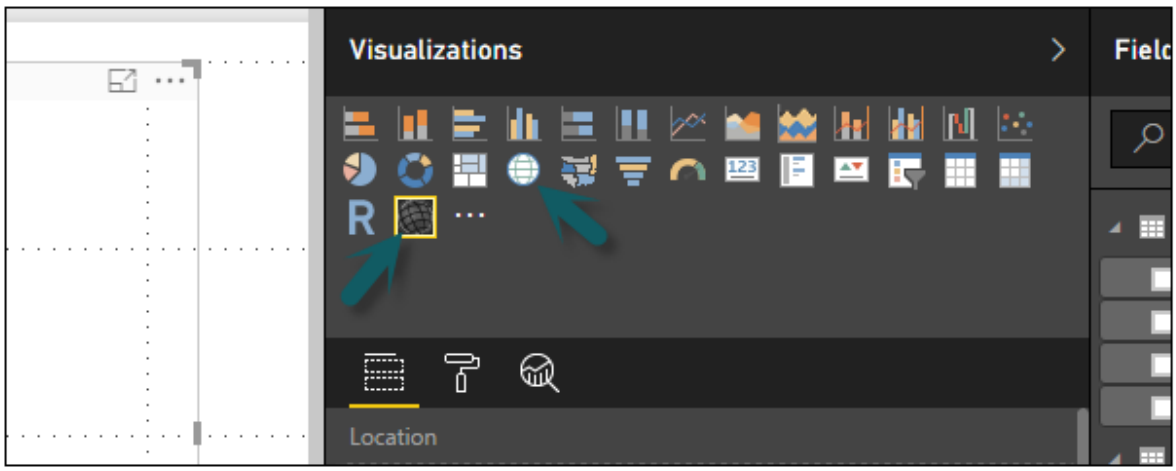

To use a bubble map, drag the map from Visualizations to the Report Canvas. To display values, you have to add any location object to the axis.

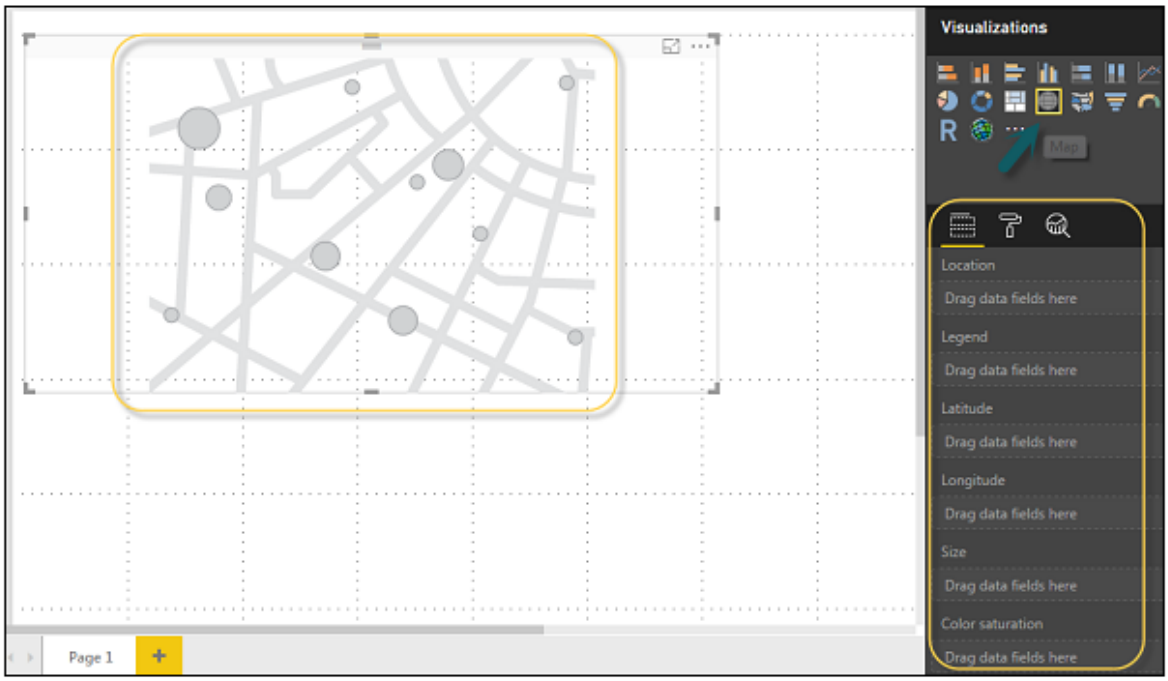

In the value fields, you can see that it accepts values axis such as City and State and or you can also add longitude and latitude values. To change the bubble size, you need to add a field to the value axis.

You can also use a filled map in data visualization, just by dragging the filled map to the Report Canvas.

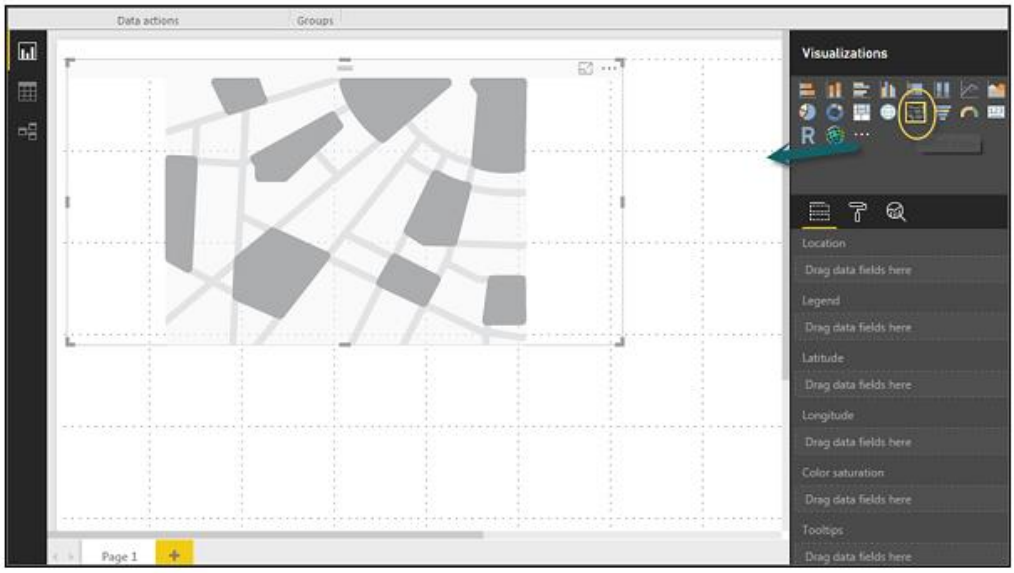

Note – If you see a warning symbol on top of your map visualization, it means that you need to add more locations to your map chart.

Using Combination Charts

In data visualization, it is also required to plot multiple measures in a single chart. Power BI supports various combination chart types to plot measure values. Let us say you want to plot revenue and unit sold in one chart. Combination charts are the most suitable option for these kinds of requirement.

One of the most common Combination charts in Power BI is Line and Stacked column charts. Let us say we have a revenue field and we have added a new data source that contains customer-wise unit quantity and we want to plot this in our visualization.

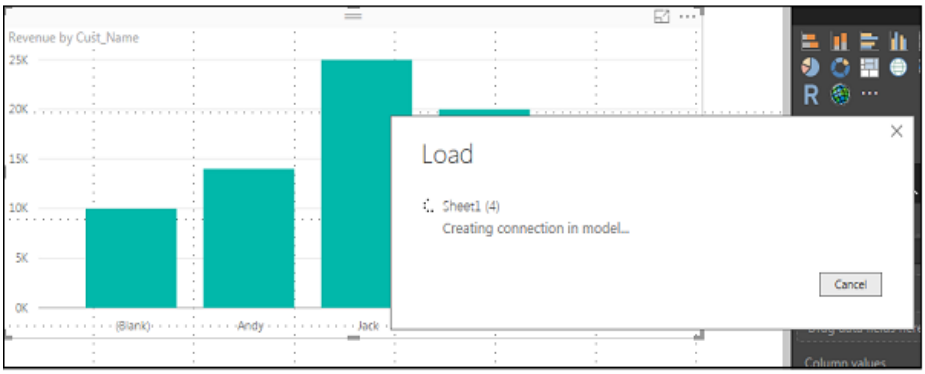

Once you add a data source, it will be added to the list of fields on the right side. You can add units to the column axis as shown in the following screenshot.

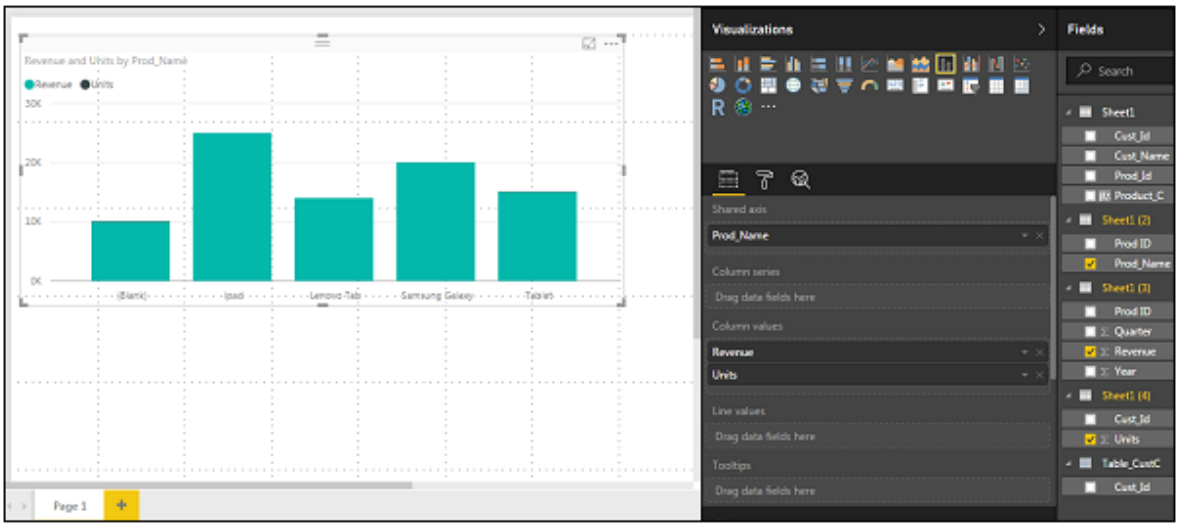

You have other type of combine chart that you can use in Power BI - Line and Clustered Column.

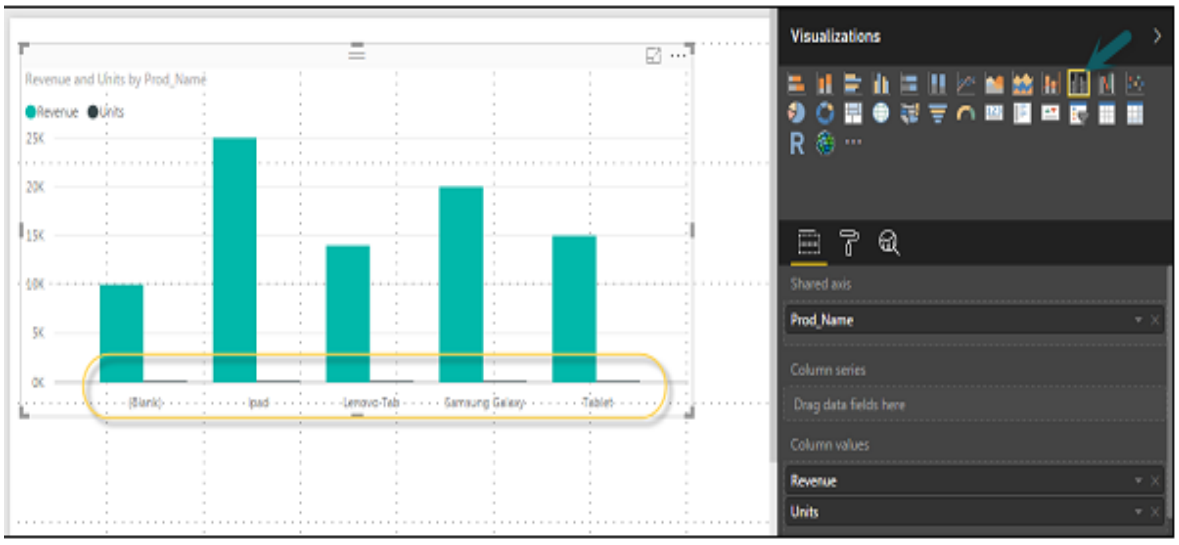

Using Tables

In Power BI, when you add a dataset to your visualization, it adds a table chart to the Report canvas. You can drag the fields that you want to add to the report. You can also select the checkbox in front of each field to add those to the Report area.

With the numerical values in a table, you can see a sum of values at the bottom.

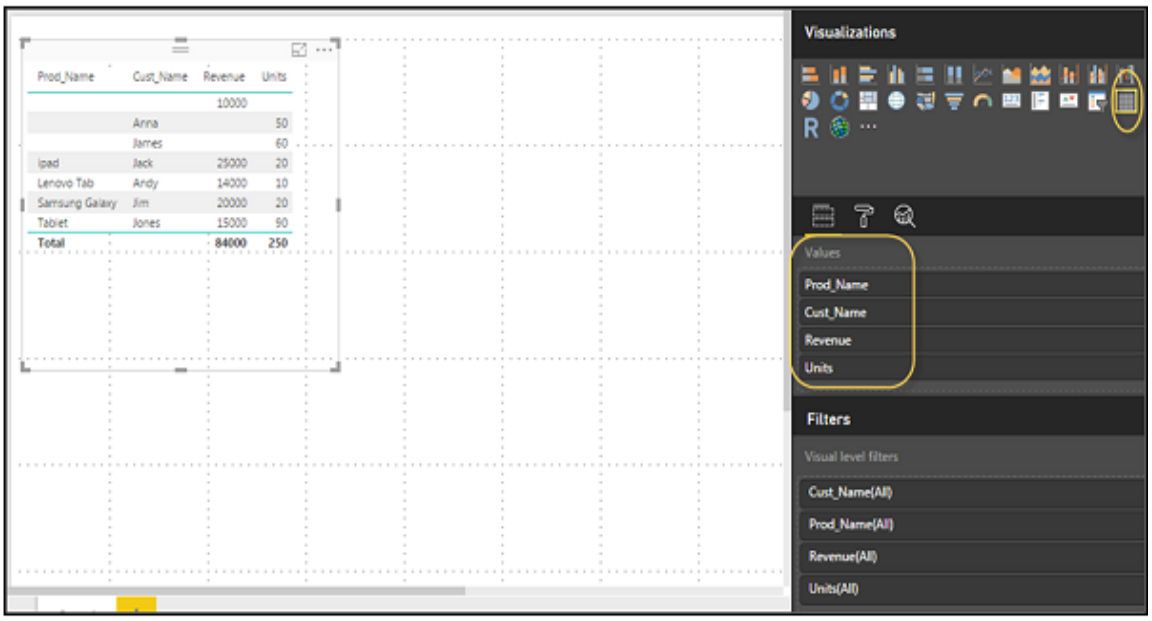

You can also perform a sort in the table using an arrow key at the top of the column. To perform ascending/descending sort, just click the arrow mark, and the values in the column will be sorted.

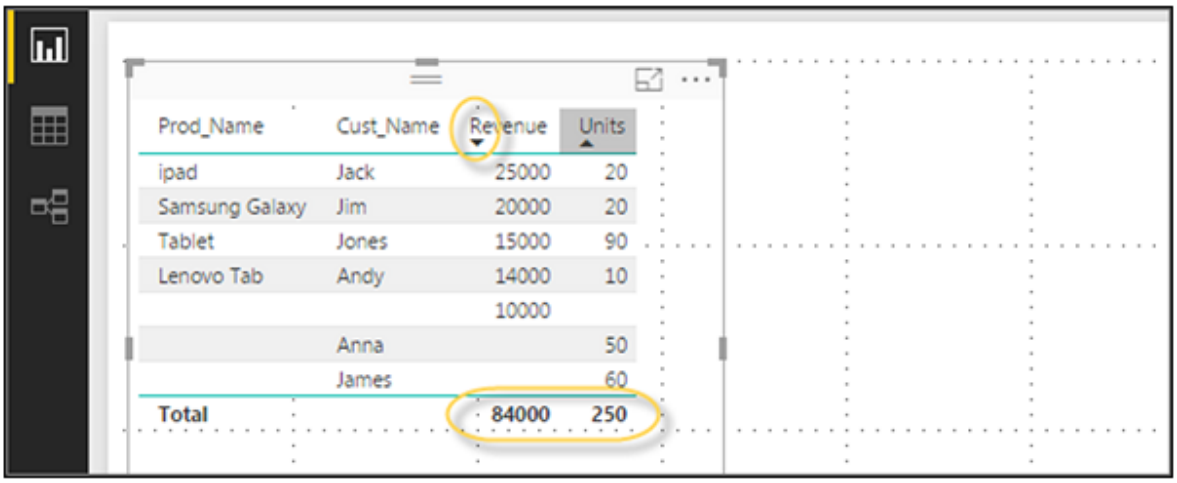

The order of the columns in a table is determined by the order in the value bucket on the right side. If you want to change the order, you can delete any column and add the other one.

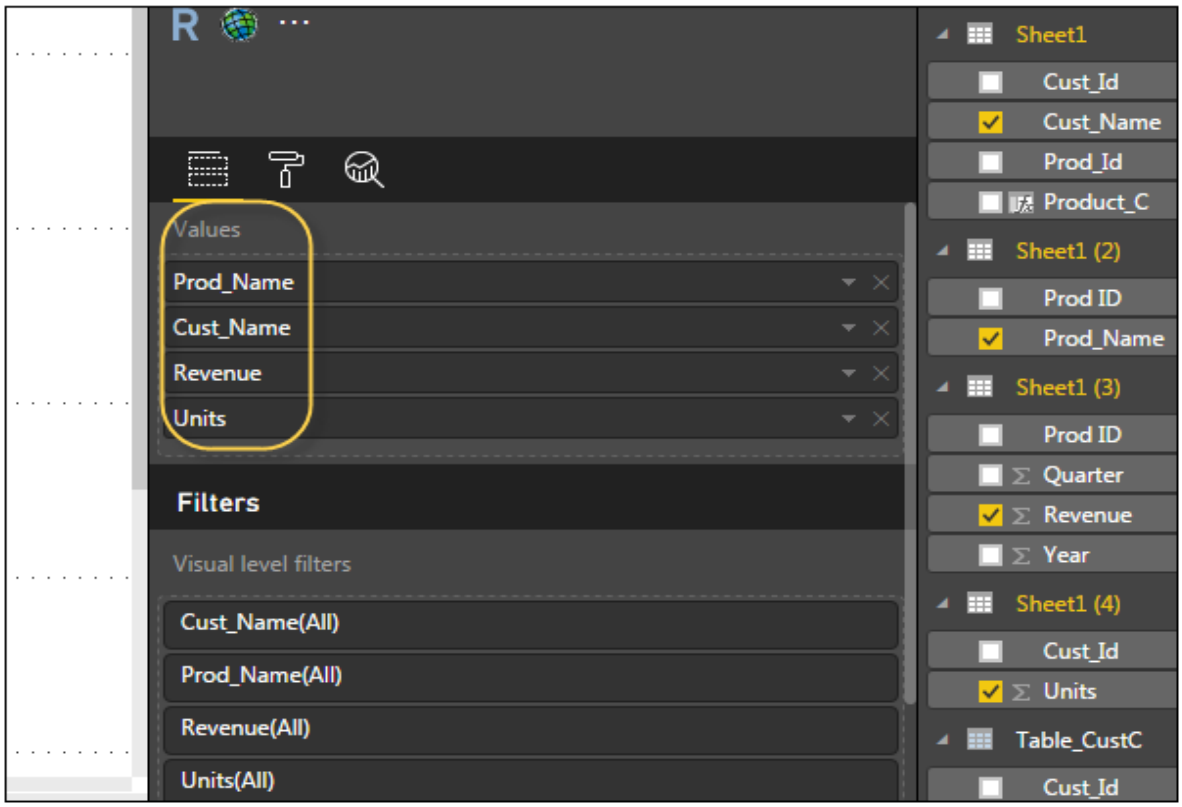

You can also undo summarize or apply different aggregate function on numerical values in the table. To change the aggregation type, click the arrow in the value bucket in front of the measure and you will see a list of formulas that can be used

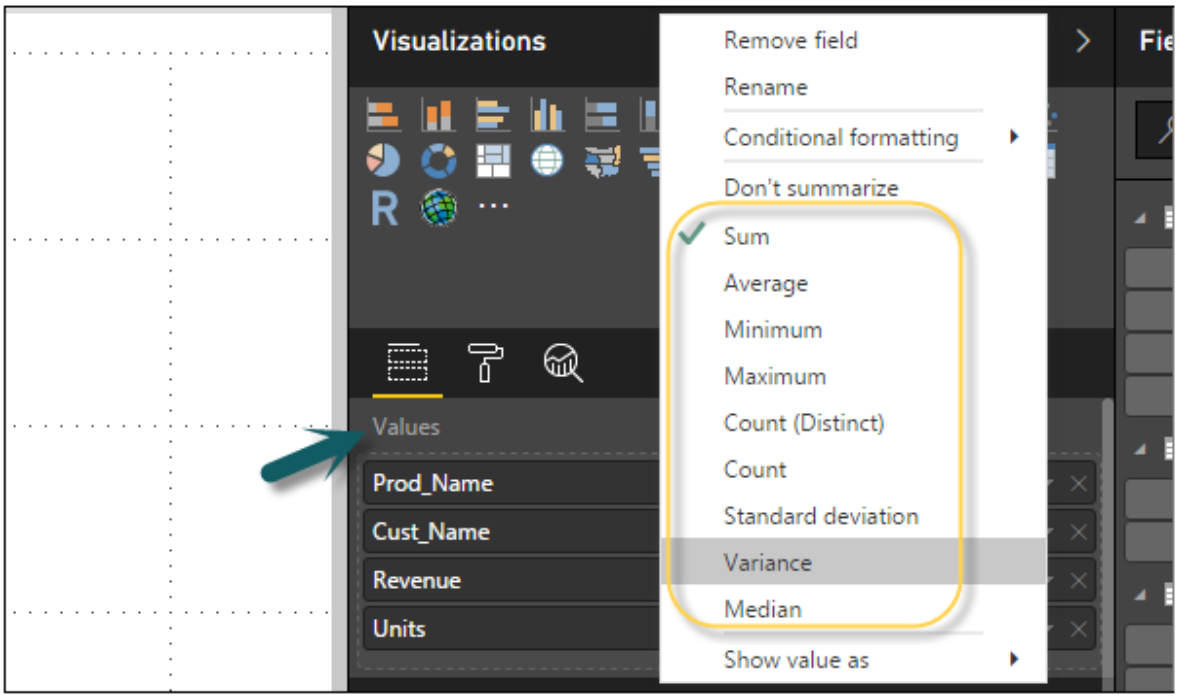

Another table type in Power BI is the matrix table that provides a lot of features such as auto sizing, column tables, and setting colors, etc.

Prod_Name	Units	Revenue
Tablet	90	15000
Samsung Gal...	20	20000
Lenovo Tab	10	14000
ipad	20	25000
	110	10000
Total	**250**	**84000**

Modify Colours in Charts

In Power BI, you can also modify the colours in the chart. When you select any visualization, it has an option to change the colour. Following options are available under the Format tab −

- Legend
- Data colours
- Detail Label
- Title
- Background
- Lock Aspect
- Border
- General

To open these options, go to the Format tab as shown in the following screenshot. Once you click, you can see all the options available.

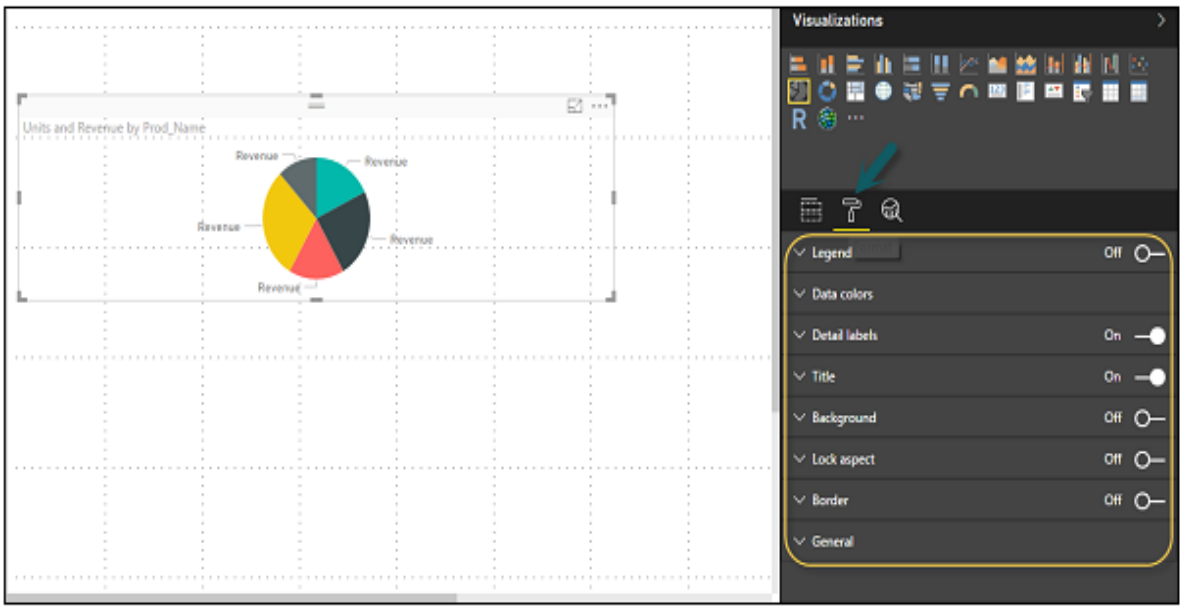

When you expand the Legend field, you have an option where you want to display the legend. You can select –

- Position
- Title
- Legend Name
- Colour
- Text Size
- Font Family

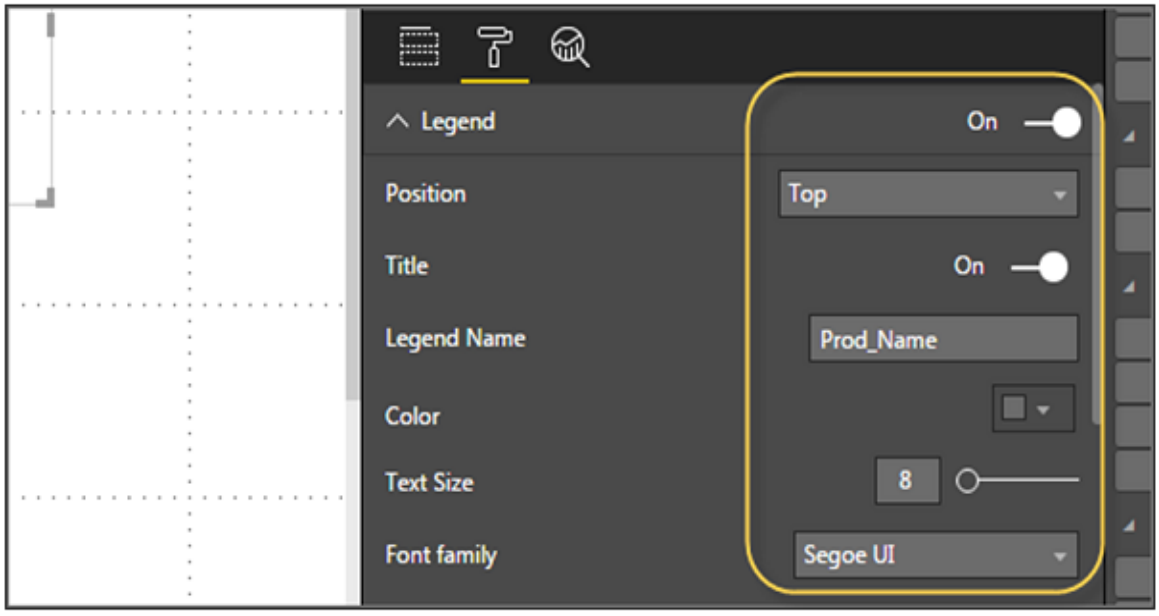

Similarly, you have data colours. In case, you want to change the colour of any data field, you can use this option. It shows all objects and their corresponding colours in the chart.

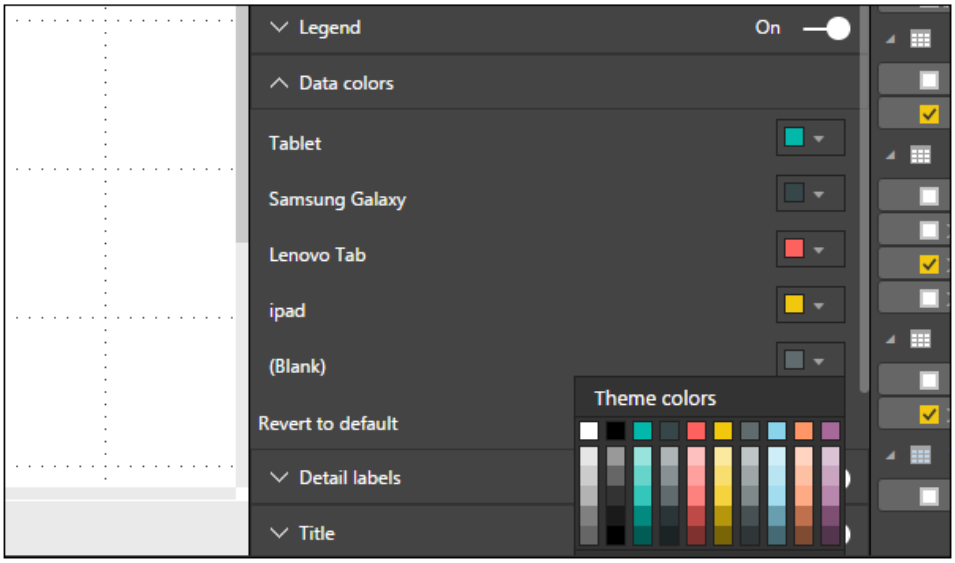

You also have Analytics feature in the tool, where you can draw lines as per requirement in data visualization. You have the following line types in data visualization −

- Constant Line
- Min Line
- Max Line
- Average Line
- Median Line
- Percentile Line

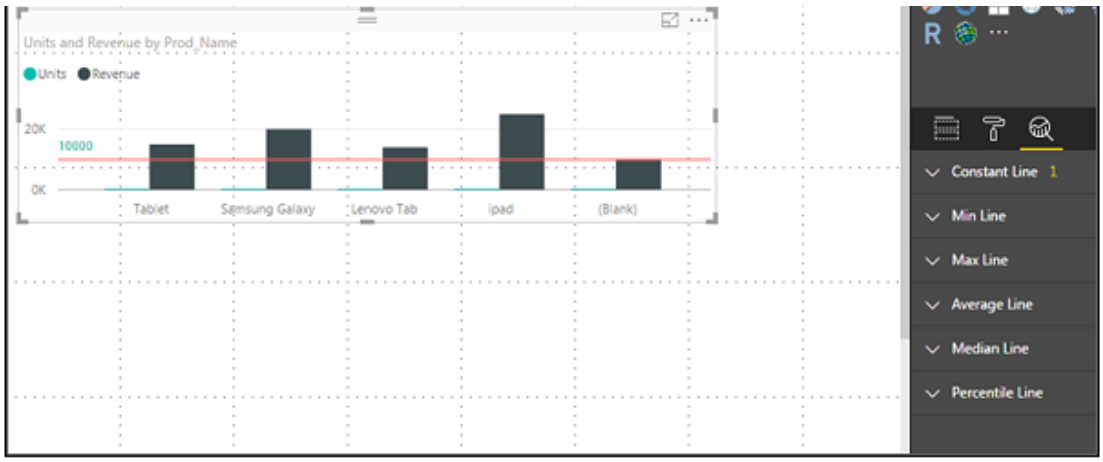

You can opt for a dashed, dotted, or a solid line. You can select Transparency level, color, and position of the line. You can also switch on/off data label for this line.

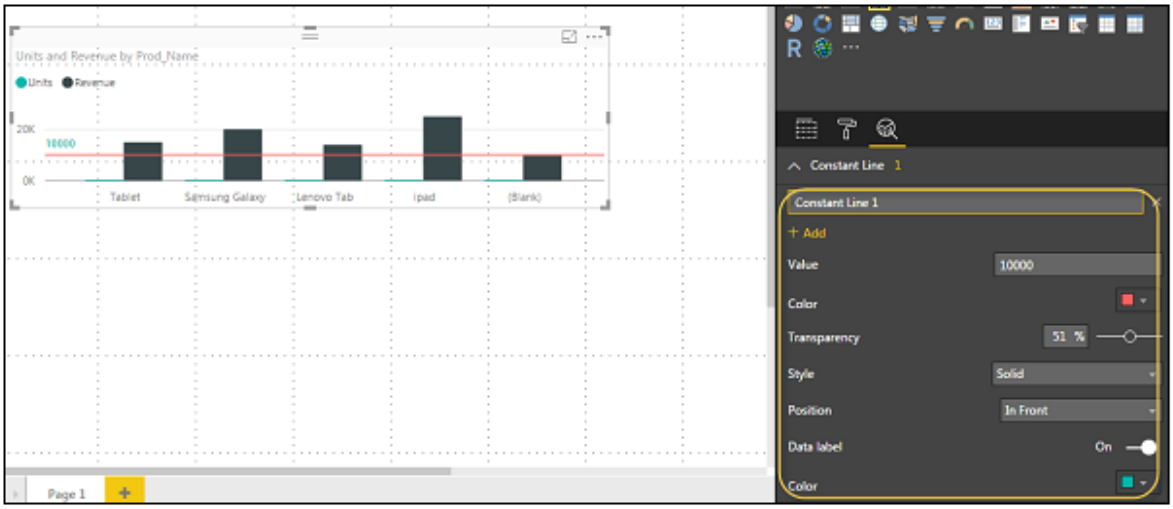

Adding Shapes, Images and Text box

Sometimes it is required that you need to add static text, images, or shapes to your visualization. In case you want to add header/footer or any static signatures, messages to data visualization this option can be used.

You can also add URLs in the text box and Power BI uses those link to make it live.

To add shapes, images and text box, navigate to the Home tab and at the top you will find an option to add images.

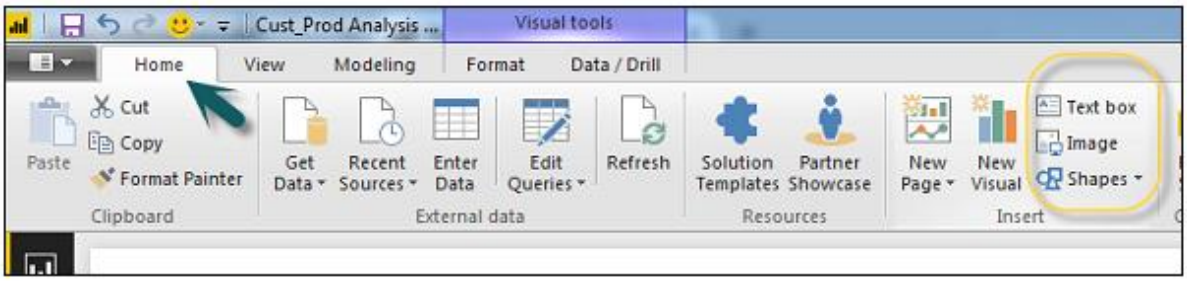

You can insert different shapes in data visualization. To see the available shapes, click the arrow next to the Shapes button.

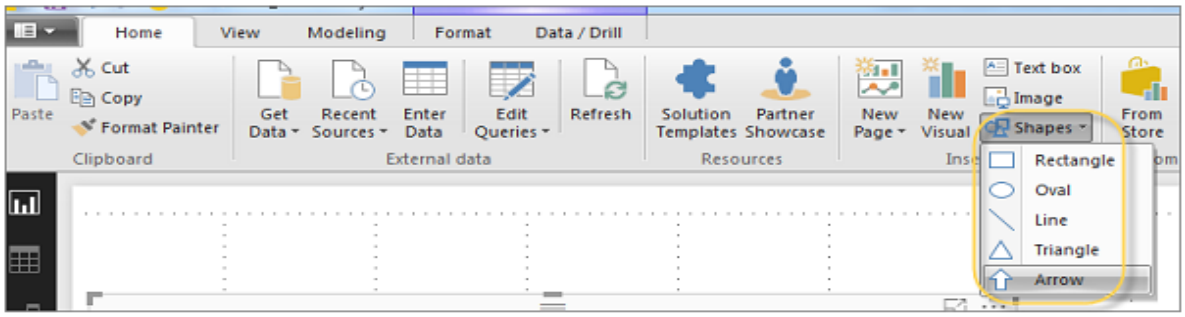

When you click on the text box, it adds a text box in your Report canvas. You can enter any text in the text box and use the rich text editor to make formatting changes.

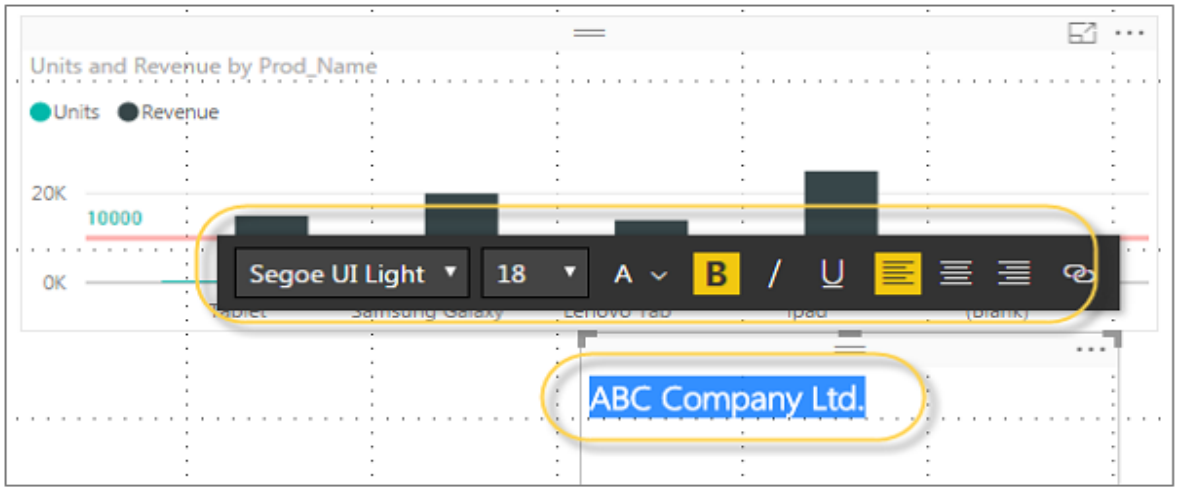

Similarly, images can be added to data visualization to add logos or other images to data visualization. When you click the Image option, it asks for a path to pass the image file.

You can add shapes by selecting any shape from the dropdown list. You can also resize it using different options.

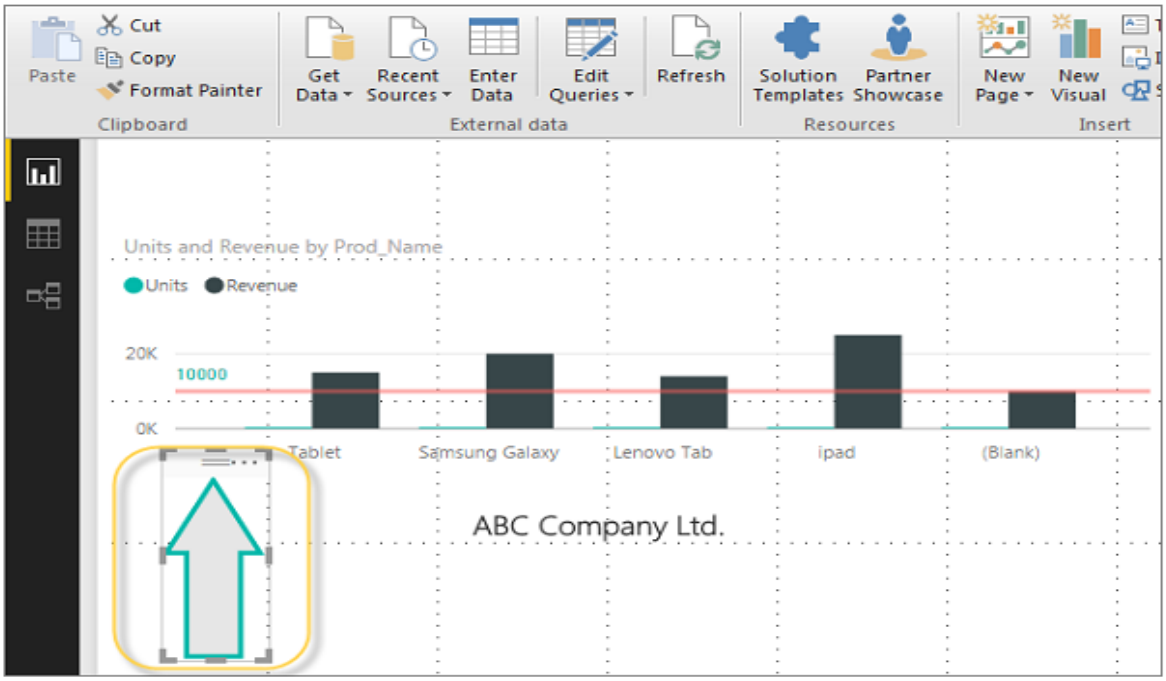

Styling Reports

In Power BI, you have flexible options to adjust the page layout and formatting such as orientation and page size of your report. Navigate to Page View menu from the Home tab and the following options are provided.

- Fit to Page
- Fit to Width
- Actual Size

By default, the page size in a report is 16:9; however, it is also possible to change the page size of the report. To change the page size, navigate to the Visualization pane and select Paint brush.

Note – To change page size, no visualization should be added to the Report canvas. You have the following options available under Page layout –

- Page Information
- Page Size
- Page Background

Under Page Information, you have Name and Q&A.

Under Page Size, you can select from the following options –

- Type
- Width
- Height

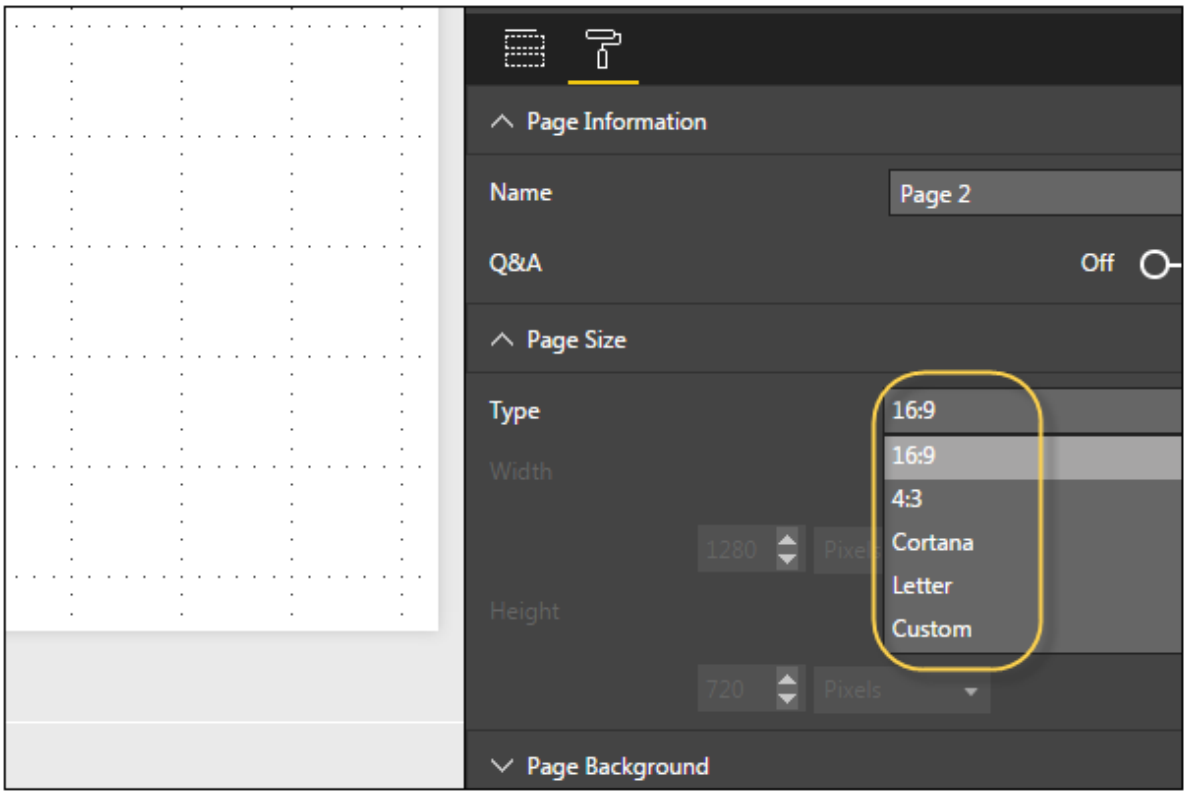

Under Page Background, you can select from the following options:

- Colour
- Transparency
- Add Image

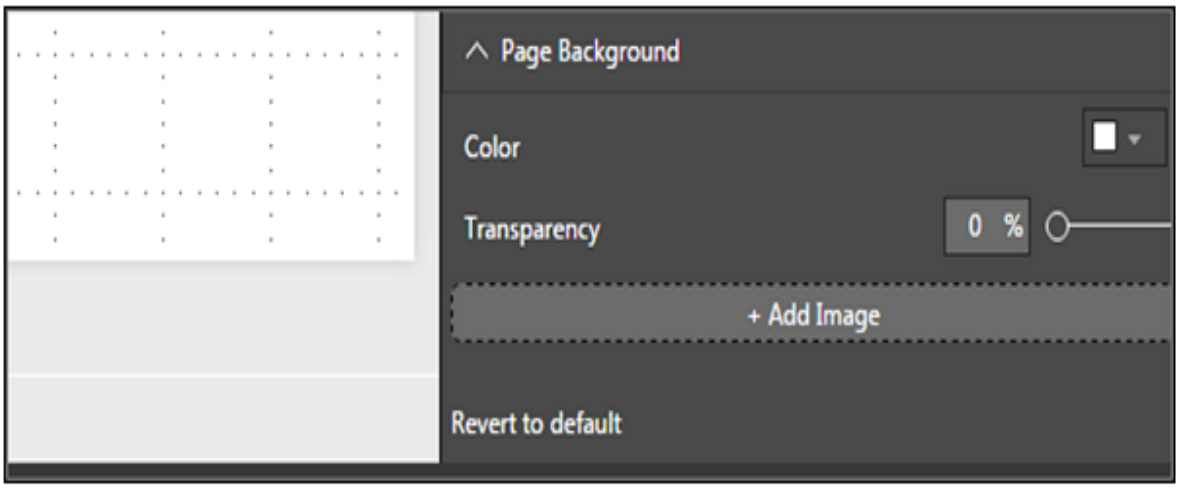

Duplicating Reports

In some scenarios, you may want to use the same layout and visuals for different pages. Power BI provides an option to create a copy of the page. When you use Duplicate Page option, a new page is added with similar layout and visuals.

To duplicate a page, right-click the Page and select Duplicate Page option. This will create a copy of the same page with the name - Duplicate of Page1.

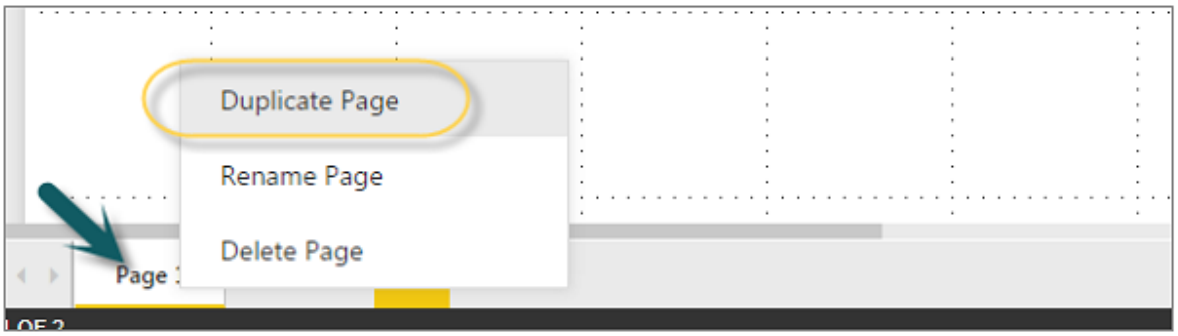

5.Data Analysis in Power BI

Microsoft Power BI is the leading data analytics, business intelligence and reporting tool in Gartner's latest Magic Quadrant. Beyond the opinion of a leading global technology consulting company, Power BI is also one of the most widely used data analytics tools by professionals because of its business approach.

The tool enables data scientists and analysts to turn data sets into engaging, interactive and insight-enhancing reports, dashboards and visuals used to monitor the course of business strategies. Power BI allows us to develop customized solutions for each of our customers.

Data analytics with Power BI promotes the discovery of information and the democratisation of data through its unique data visualisation capabilities, that improve the quality of companies' informative systems and help executives make data-driven decisions.

Power BI streamlines data analysts' work and makes it easy to connect, transform and visualise data.

Nevertheless, there are many other similar tools on the market. For example, Excel also by Microsoft is another essential tool for data experts. In fact, using Excel and Power BI together is very common.

Important reasons why you should do data analytics with Power BI

1. **Data connection**

   One of Power BI's greatest advantages is its extensive data connectivity. The tool connects to multiple tabular databases and integrates with a host of corporate tools and systems to make importing and exporting data, dashboards and reports as simple and fast as possible.

2. **Data visualization**

   Power BI is **one of the most complete platforms for data visualisation**.

In the tool's appsource you will find a lot of Power BI visuals validated by Microsoft, but you can also create your own custom visuals.

In addition, Power BI adds new visuals from time to time and you can even extend its visualisation capabilities with Zebra BI, which works both for Excel and Power BI.

3. **Data visualization**

Power BI is **one of the most complete platforms for data visualisation**. In the tool's appsource you will find a lot of Power BI visuals validated by Microsoft, but you can also create your own custom visuals.

In addition, Power BI adds new visuals from time to time and you can even extend its visualisation capabilities with Zebra BI, which works both for Excel and Power BI.

4. **Advanced analytics**

Power BI is the optimal platform for increasing the value of your regular Excel data analysis with **advanced analytics**. You can enrich business data by ingesting, transforming and integrating data services to other Microsoft suite's tools.

5. **Data governance**

For anyone who works with data, **data governance is a must** to ensure the smooth running of any type of process, especially in the business environment, since organisations often have a large amount of data that, when not well organised, can lose all its value.

Power BI includes features that support data control, authority and management. However, the tool has limited data governance capabilities and some organisations require specialised data governance solutions to work with Power BI.

Find out how you can apply data governance policies to Power BI

With **Power BI Viewer** you can organise your reports by category, assign permissions according to users' role, access all your reports without having a Power BI licence and much more.

With **Power BI Analytics** you will be able to analyse the activity of all your users with a historical repository with no space or time limits.

6. Data exploration

Power BI contains extensive data exploration options as well as automated queries. With this tool, discovering insights from data will be much easier. It is also the ideal platform for working with a top-down methodology.

7. UX & UI

Finally, it should be noted that Power BI was conceived as an enterprise tool designed to be used by business users. Clearly, its main users are data analysts and BI consultants. However, the business conception of the platform makes it one of the BI tools with the best usability and user interface. In addition, with Power BI you can adapt your reports to your brand image and automate the process by designing themes that can be applied to all your reports.

8. Manage workspaces and datasets in Power BI

Workspaces are places to collaborate with colleagues to create collections of dashboards, reports, datasets, and paginated reports. Here we describe how to manage, access and how to use them to create and distribute apps.

Creating a workspace

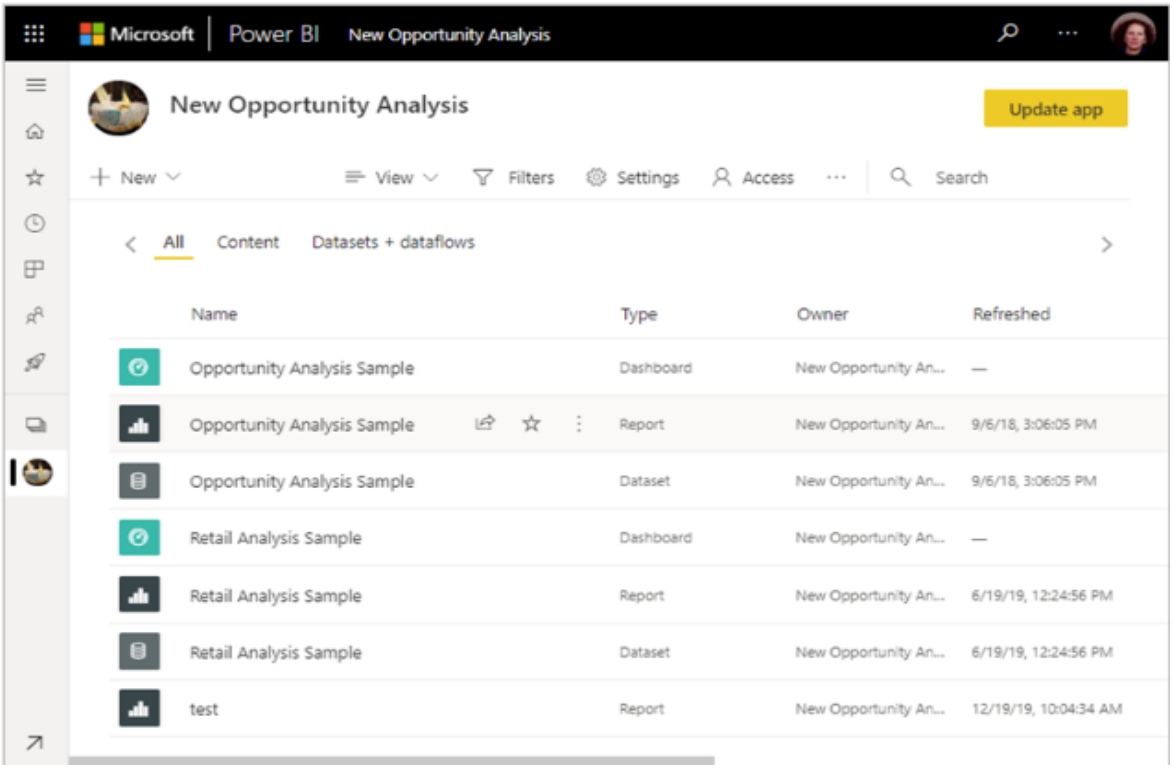

Working with workspaces

Here are some useful tips about working with workspaces.

- **Use granular workspace roles** for flexible permissions management in the workspaces: Admin, Member, Contributor, and Viewer.
- **Contact list**: Specify who receives notification about workspace activity.
- **Create template apps**: You can create *template apps* in workspaces. Template apps are apps that you can distribute to customers outside of your organization. Those customers can then connect to their own data with your template app.
- **Share datasets**: You can share datasets between workspaces.

Workspace contact list

The **Contact list** feature allows you to specify which users receive notification about issues occurring in the workspace. By default, any user or group specified as a workspace admin in the workspace is notified. You can add to that list. Users or groups in the contact list are also listed in the user interface

(UI) of the workspaces, so workspace end-users know whom to contact.

Create a workspace in Power BI

Here, we discuss how to create *workspaces*, spaces to collaborate with colleagues. In them, you create collections of dashboards, reports, and paginated reports. If you want, you can also bundle that collection into an *app* and distribute it to a broader audience.

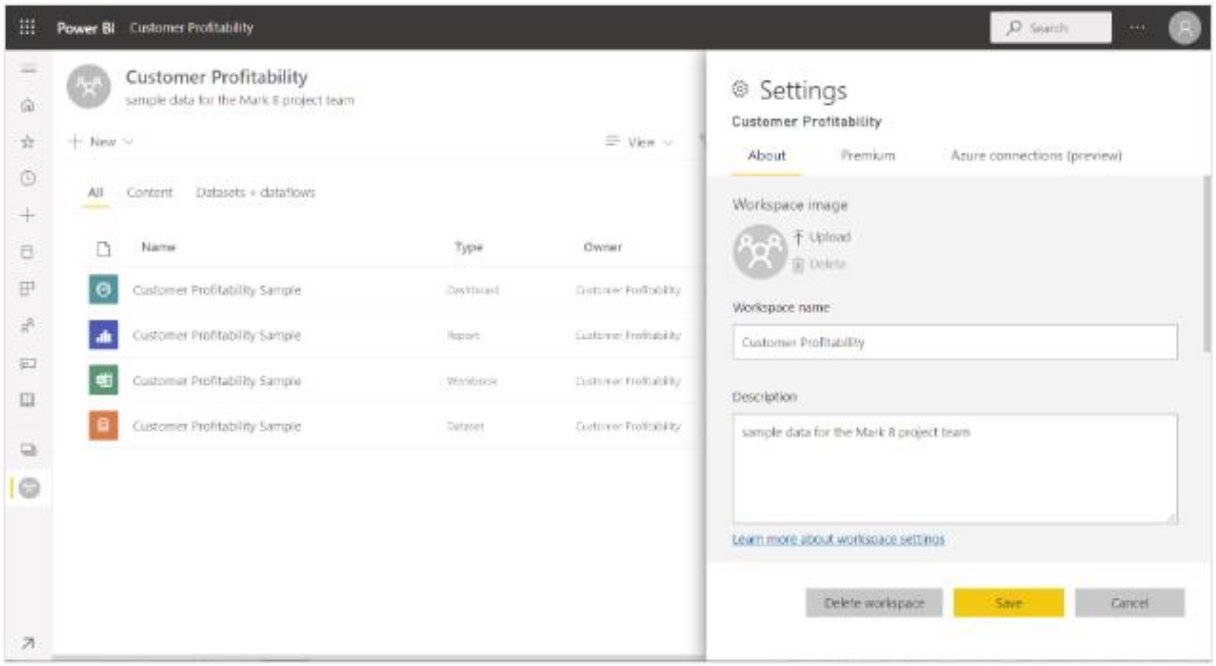

Create a workspace

1. Select **Workspaces** > **Create workspace**.

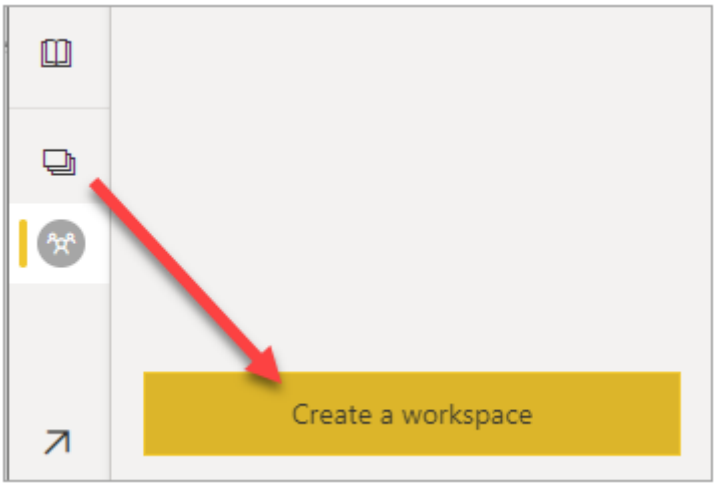

2. Give the workspace a unique name. If the name isn't available, edit it to come up with a name that's unique.

   When you create an app from the workspace, by default it will have the same name and icon as the workspace. You can change both when you create the app.

3. Here are some optional settings for your workspace.

   - Upload a Workspace image. Files can be .png or .jpg format. File size has to be less than 45 KB.
   - Specify a Workspace One Drive to use a Microsoft 365 Group file storage location (provided by SharePoint).
   - Add a Contact list, the names of people to contact for information about the workspace. By default, the workspace admins are the contacts.
   - Allow contributors to update the app for the workspace
   - Assign the workspace to a Premium capacity.
   - Connect the workspace to an Azure Data Lake Gen2 storage account (in preview).

4. Select **Save**.

   Power BI creates the workspace and opens it. You see it in the list of workspaces you're a member of. Microsoft 365 and OneDrive: Power BI doesn't create a Microsoft 365 group behind the scenes when you create a

workspace. All workspace administration is in Power BI. Still, you might find it useful to have a OneDrive associated with the workspace.

- You can **manage user access** to content through Microsoft 365 groups, if you want. You add a Microsoft 365 group in the workspace access list.

  Power BI doesn't synchronize between Microsoft 365 group membership and permissions for users or groups with access to the workspace. You can synchronize them: Manage workspace access through the same Microsoft 365 group whose file storage you configure in this setting.

- You can also **store Power BI content in OneDrive for Business**. With the Workspace One Drive feature in workspaces, you can configure a Microsoft 365 group whose SharePoint Document Library file storage is available to workspace users. You create the group outside of Power BI.

Note :

Power BI lists all Microsoft 365 groups that you're a member of in the workspaces list.

Roles and licenses

Roles let you manage who can do what in workspaces, so team members can collaborate. To grant access to a workspace, assign those user groups or individuals to one of the workspace roles: Admin, Member, Contributor, or Viewer.

- **Licensing enforcement**: Publishing reports to a workspace enforces existing licensing rules. Users collaborating in workspaces or sharing content to others in the Power BI service need a Power BI Pro or Premium Per User (PPU) license. Users without a Pro or PPU license see the error "Only users with Power BI Pro licenses can publish to this workspace."
- **Read-only workspaces**: The Viewer role in workspaces gives users read-only access to the content in a workspace.

- **Users without a Pro or Premium Per User (PPU) license** can access a workspace if the workspace is in a Power BI Premium capacity, but only if they have the Viewer role.
- **Allow users to export data**: Even users with the Viewer role in the workspace can export data if they have build permission on the datasets in that workspace.
- **Assign user groups to workspace roles**: You can add Active Directory security groups, distribution lists, or Microsoft 365 groups to these roles, for easier user management.

Administering and auditing workspaces

Administration for workspaces is in the Power BI admin portal. Power BI admins decide who in an organization can create workspaces and distribute apps.

Admins can also see the state of all the workspaces in their organization. They can manage, recover, and even delete workspaces.

Auditing

Power BI audits the following activities for workspaces.

Friendly name	Operation name
Created Power BI folder	CreateFolder
Deleted Power BI folder	DeleteFolder
Updated Power BI folder	UpdateFolder
Updated Power BI folder access	UpdateFolderAccess

Considerations and limitations

Limitations to be aware of:

- Workspaces can contain a maximum of 1,000 datasets, or 1,000 reports per dataset.

- Power BI publisher for Excel isn't supported.
- Certain special characters aren't supported in workspace names when using an XMLA endpoint. As a workaround, use URL encoding of special characters, for example, for a forward slash **/**, use **%2F**.

5. Key Influencers Visualizations Tutorial - Power BI

The key influencers visual helps you understand the factors that drive a metric you're interested in. It analyses your data, ranks the factors that matter, and displays them as key influencers. For example, suppose you want to figure out what influence's employee turnover, which is also known as churn. One factor might be employment contract length, and another factor might be commute time.

When to use key influencers

The key influencers visual is a great choice if you want to:

- See which factors affect the metric being analysed.
- Contrast the relative importance of these factors. For example, do short-term contracts affect churn more than long-term contracts?

Features of the key influencers visual

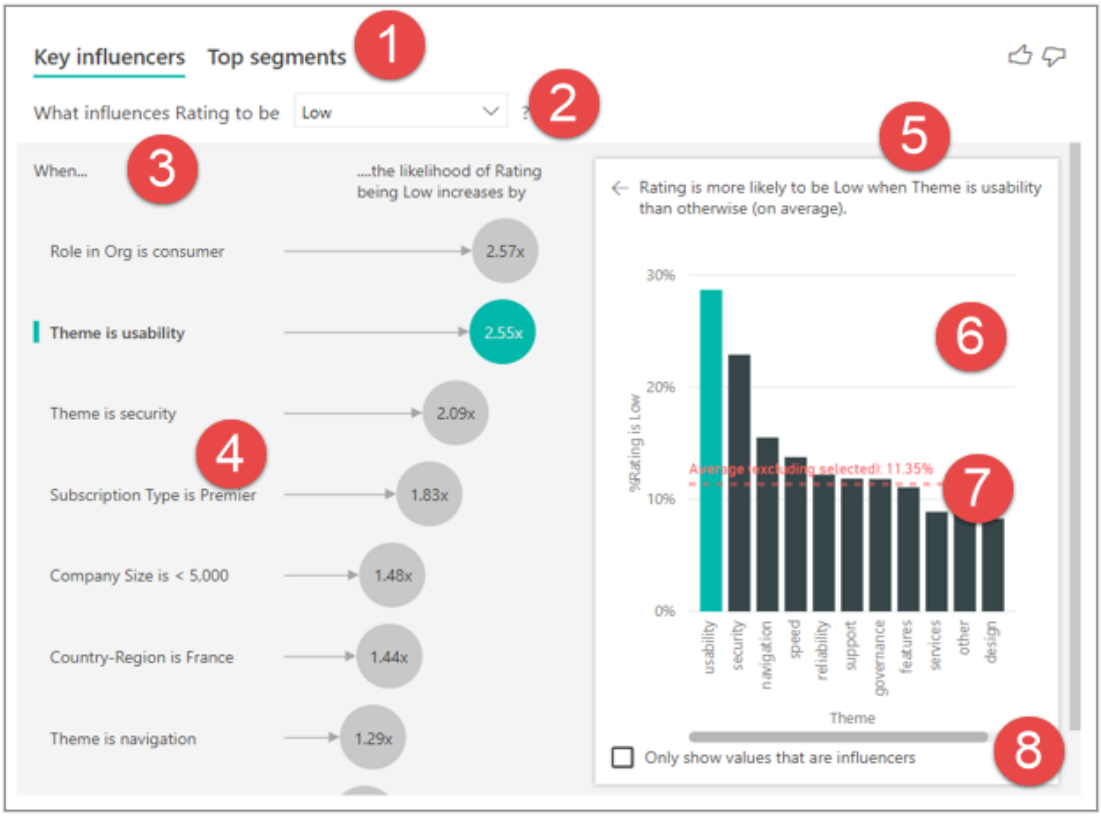

1. **Tabs**: Select a tab to switch between views. **Key influencers** show you the top contributors to the selected metric value. **Top segments** show you the top segments that contribute to the selected metric value. A *segment* is made up of a combination of values. For example, one segment might be consumers who have been customers for at least 20 years and live in the west region.
2. **Drop-down box**: The value of the metric under investigation. In this example, look at the metric **Rating**. The selected value is **Low**.
3. **Restatement**: It helps you interpret the visual in the left pane.
4. **Left pane**: The left pane contains one visual. In this case, the left pane shows a list of the top key influencers.
5. **Restatement**: It helps you interpret the visual in the right pane.
6. **Right pane**: The right pane contains one visual. In this case, the column chart displays all the values for the key influencer **Theme** that was selected in the left pane. The specific value of **usability** from the left pane is shown in green. All the other values for **Theme** are shown in black.
7. **Average line**: The average is calculated for all possible values for **Theme** except **usability** (which is the selected influencer). So, the calculation

applies to all the values in black. It tells you what percentage of the other **Themes** had a low rating. In this case 11.35% had a low rating (shown by the dotted line).

8. **Check box**: Filters out the visual in the right pane to only show values that are influencers for that field. In this example, the visual is filtered to display usability, security, and navigation.

Analyse a metric that is categorical

Here, we learn how to create a key influencer visual with a categorical metric. Then follow the steps to create one.

1. Your Product Manager wants you to figure out which factors lead customers to leave negative reviews about your cloud service. To follow along in Power BI Desktop, open the Customer Feedback PBIX file
2. Under **Build visual** on the **Visualizations** pane, select the **Key influencers** icon.
3. Move the metric you want to investigate into the **Analyse** field. To see what drives a customer rating of the service to be low, select **Customer Table** > **Rating**.
4. Move fields that you think might influence **Rating** into the **Explain by** field. You can move as many fields as you want. In this case, start with:

    - Country-Region
    - Role in Org
    - Subscription Type
    - Company Size
    - Theme

5. Leave the **Expand by** field empty. This field is only used when analysing a measure or summarized field.
6. To focus on the negative ratings, select **Low** in the **What influences Rating to be** drop-down box.

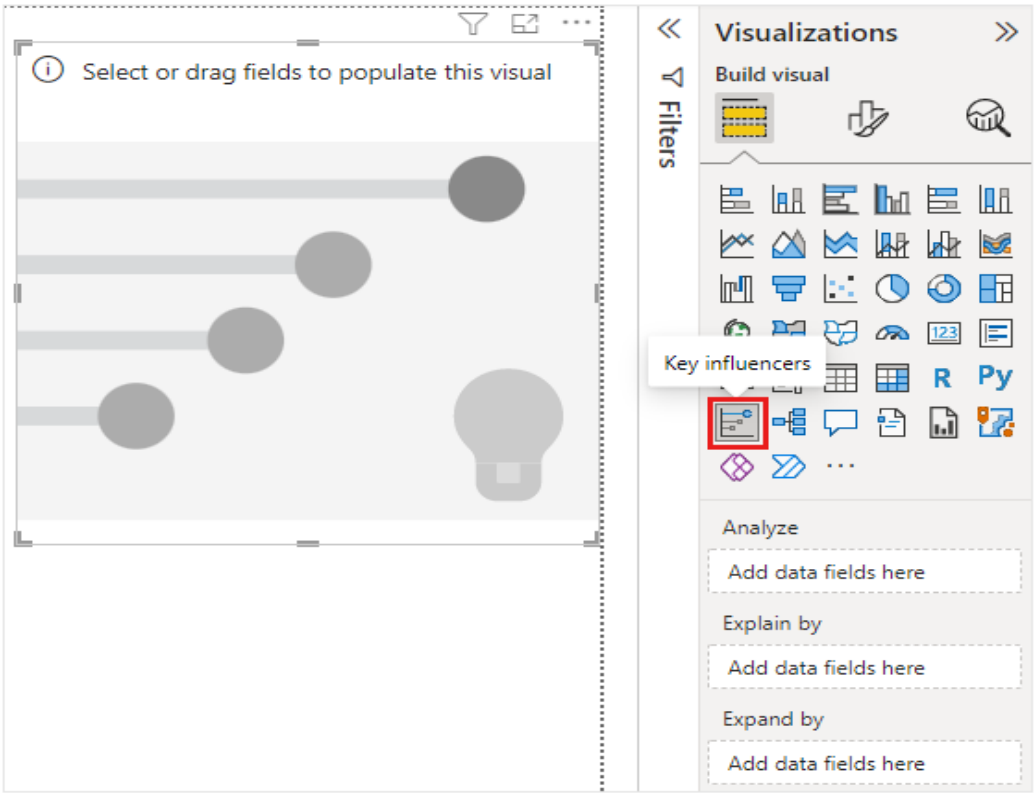

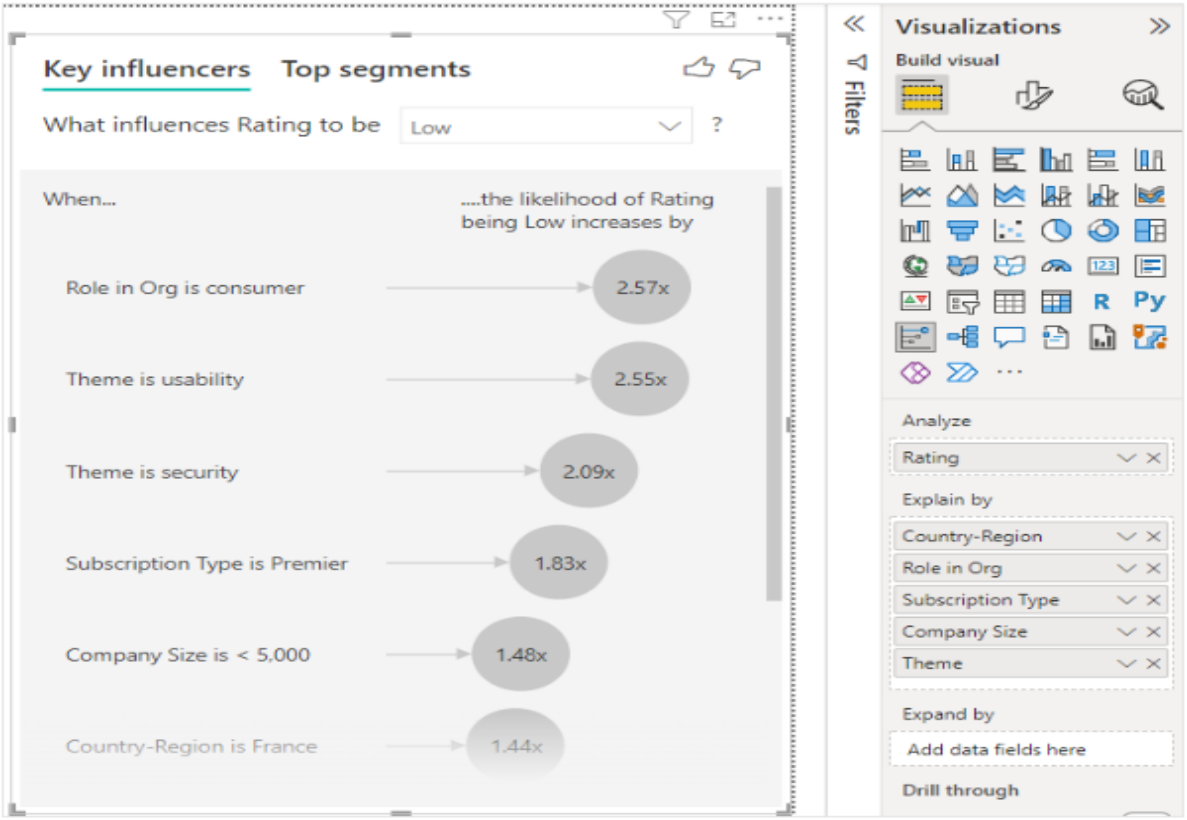

The analysis runs on the table level of the field that's being analysed. In this case, it's the Rating metric. This metric is defined at a customer level.

Each customer has given either a high score or a low score. All the explanatory factors must be defined at the customer level for the visual to make use of them.

In this case, each customer assigned a single theme to their rating. Similarly, customers come from one country, have one membership type, and hold one role in their organization. The explanatory factors are already attributes of a customer, and no transformations are needed. The visual can make immediate use of them.

Interpret categorical key influencers

Let's take a look at the key influencers for low ratings.

Top single factor that influences the likelihood of a low rating

The customer in this example can have three roles: consumer, administrator, and publisher. Being a consumer is the top factor that contributes to a low rating.

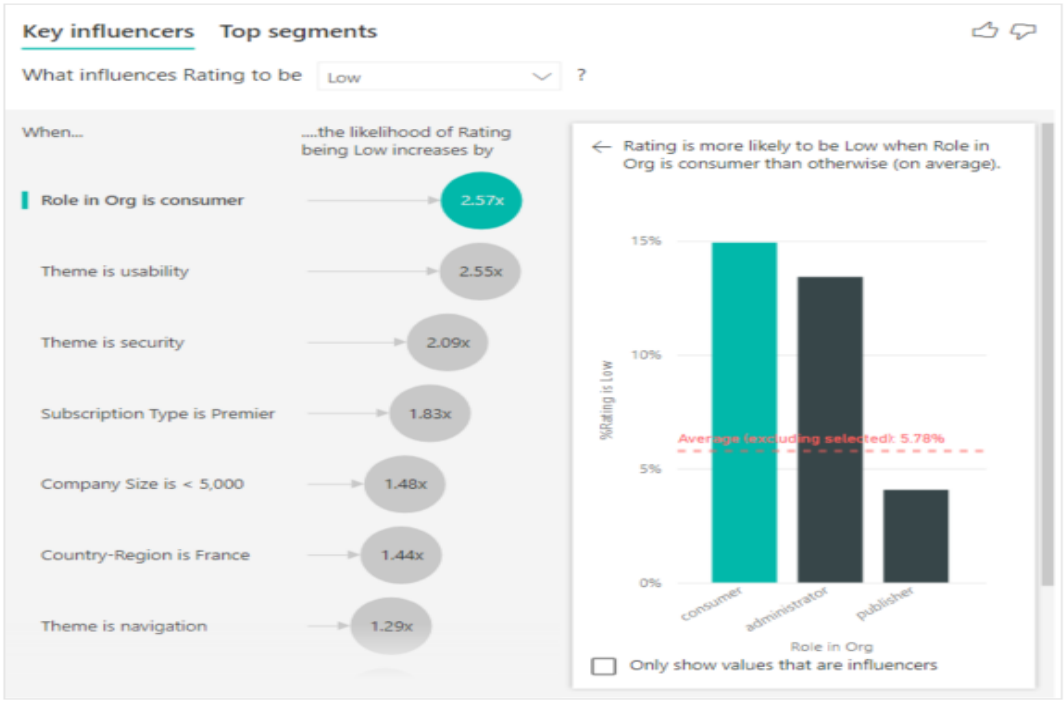

More precisely, your consumers are 2.57 times more likely to give your service a negative score. The key influencers chart lists **Role in Org is consumer** first

in the list on the left. By selecting **Role in Org is consumer**, Power BI shows more details in the right pane. The comparative effect of each role on the likelihood of a low rating is shown.

- 14.93% of consumers give a low score.
- On average, all other roles give a low score 5.78% of the time.
- Consumers are 2.57 times more likely to give a low score compared to all other roles. You can determine this score by dividing the green bar by the red dotted line.

Second single factor that influences the likelihood of a low rating

The key influencers visual compares and ranks factors from many different variables. The second influencer has nothing to do with **Role in Org**. Select the second influencer in the list, which is **Theme is usability**.

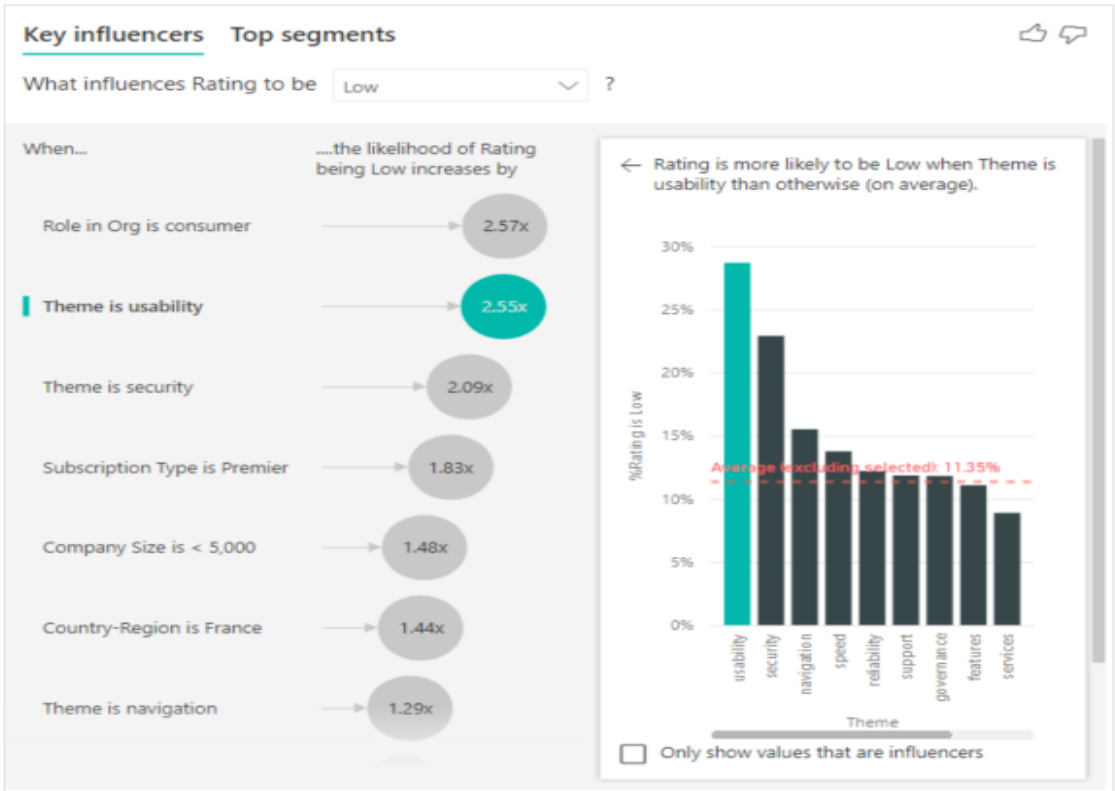

The second most important factor is related to the theme of the

customer's review. Customers who commented about the usability of the product were 2.55 times more likely to give a low score compared to customers who commented on other themes, such as reliability, design, or speed.

Between the visuals, the average, which is shown by the red dotted line, changed from 5.78% to 11.35%. The average is dynamic because it's based on the average of all other values. For the first influencer, the average excluded the customer role. For the second influencer, it excluded the usability theme.

Select the **Only show values that are influencers** check box to filter by using only the influential values. In this case, they're the roles that drive a low score. 12 themes are reduced to the four that Power BI identified as the themes that drive low ratings.

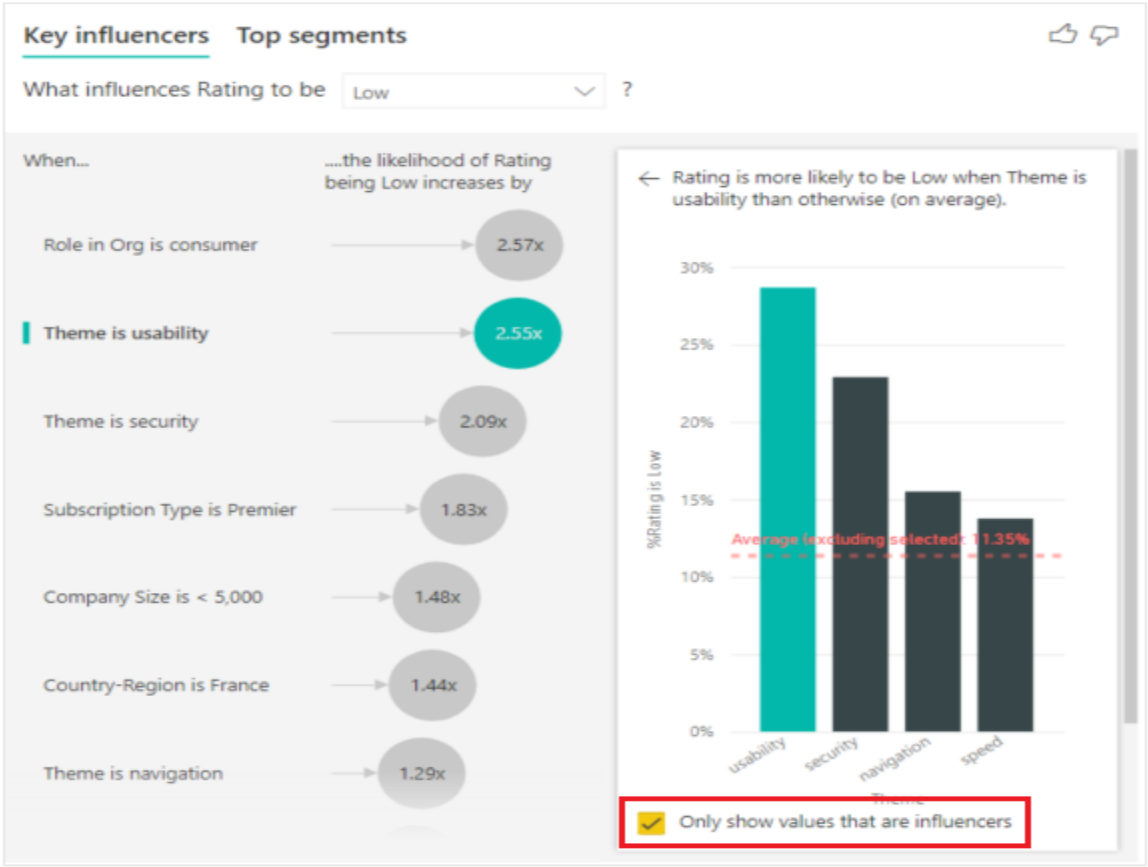

Interact with other visuals

Every time you select a slicer, filter, or other visual on the canvas, the

key influencers visual reruns its analysis on the new portion of data. For example, you can move **Company Size** into the report and use it as a slicer. Use it to see if the key influencers for your enterprise customers are different than the general population. An enterprise company size is larger than 50,000 employees.

Select **>50,000** to rerun the analysis, and you can see that the influencers changed. For large enterprise customers, the top influencer for low ratings has a theme related to security. You might want to investigate further to see if there are specific security features your large customers are unhappy about.

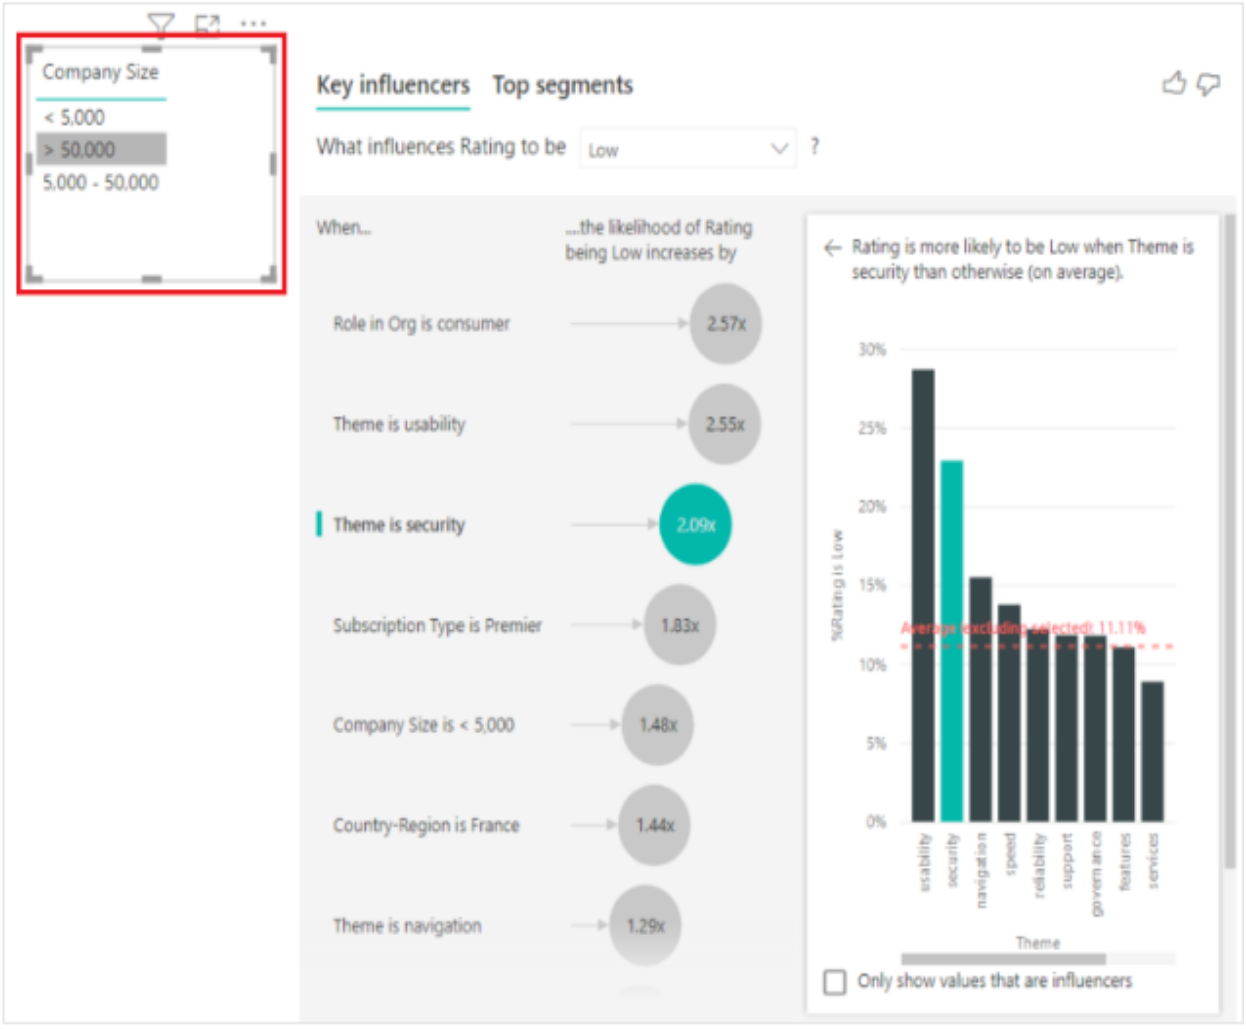

Interpret continuous key influencers

So far, you've seen how to use the visual to explore how different categorical fields influence low ratings. It's also possible to have continuous

factors such as age, height, and price in the **Explain by** field. Let's look at what happens when **Tenure** is moved from the customer table into **Explain by**. Tenure depicts how long a customer has used the service.

As tenure increases, the likelihood of receiving a lower rating also increases. This trend suggests that the longer-term customers are more likely to give a negative score. This insight is interesting, and one that you might want to follow up on later.

The visualization shows that every time tenure goes up by 13.44 months, on average the likelihood of a low rating increases by 1.23 times. In this case, 13.44 months depict the standard deviation of tenure. So, the insight you receive looks at how increasing tenure by a standard amount, which is the standard deviation of tenure, affects the likelihood of receiving a low rating.

The scatter plot in the right pane plots the average percentage of low ratings for each value of tenure. It highlights the slope with a trend line.

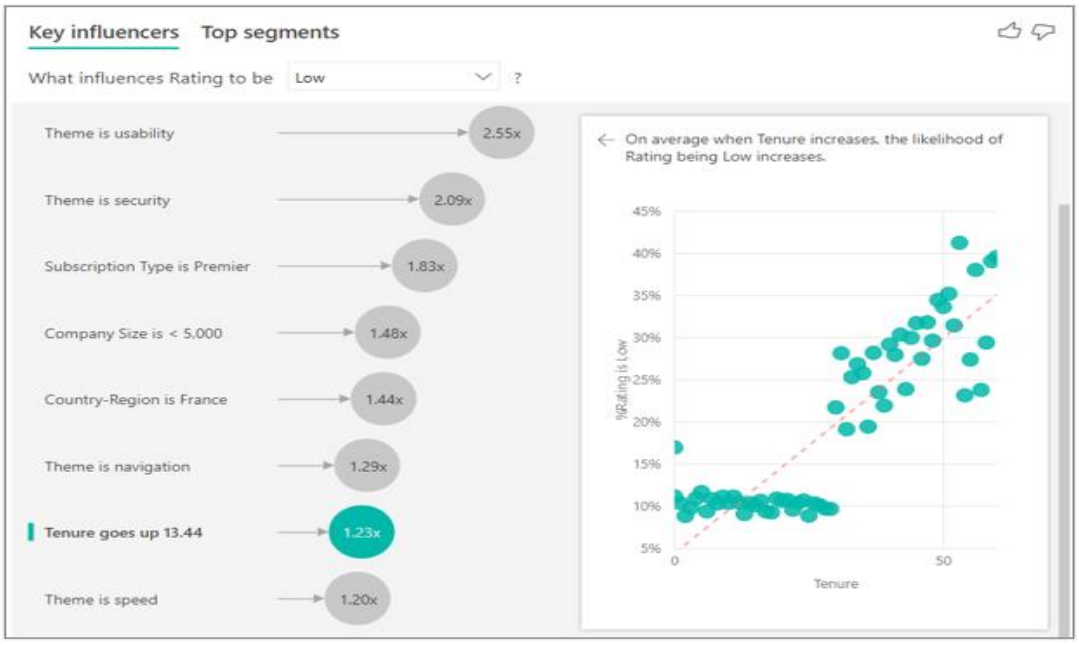

Binned continuous key influencers

In some cases, you may find that your continuous factors were automatically turned into categorical ones. If the relationship between the

variables isn't linear, we can't describe the relationship as simply increasing or decreasing.

We run correlation tests to determine how linear the influencer is with regard to the target. If the target is continuous, we run Pearson correlation and if the target is categorical, we run Point Biserial correlation tests. If we detect the relationship isn't sufficiently linear, we conduct supervised binning and generate a maximum of five bins. To figure out which bins make the most sense, we use a supervised binning method that looks at the relationship between the explanatory factor and the target being analysed.

Interpret measures and aggregates as key influencers

You can use measures and aggregates as explanatory factors inside your analysis. For example, you might want to see what effect the count of customer support tickets or the average duration of an open ticket has on the score you receive.

In this case, you want to see if the number of support tickets that a customer has influences the score they give. Now you bring in **Support Ticket ID** from the support ticket table. Because a customer can have multiple support tickets, you aggregate the ID to the customer level. Aggregation is important because the analysis runs on the customer level, so all drivers must be defined at that level of granularity.

Let's look at the count of IDs. Each customer row has a count of support tickets associated with it. In this case, as the count of support tickets increases, the likelihood of the rating being low goes up 4.08 times. The visual on the right shows the average number of support tickets by different **Rating** values evaluated at the customer level.

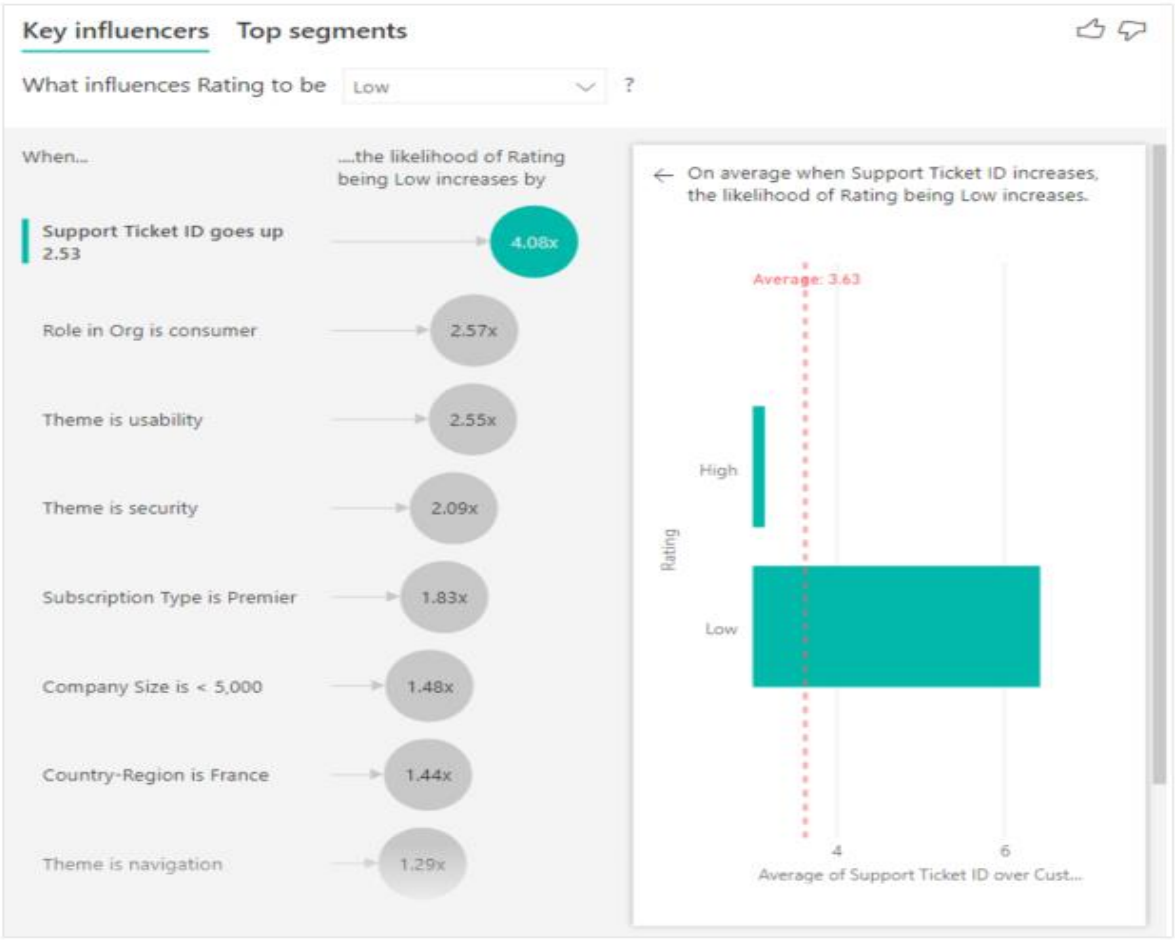

Interpret the results: Top segments

You can use the **Key influencers** tab to assess each factor individually. You also can use the **Top segments** tab to see how a combination of factors affects the metric that you're analysing.

Top segments initially show an overview of all the segments that Power BI discovered. The following example shows that six segments were found. These segments are ranked by the percentage of low ratings within the segment. Segment 1, for example, has 74.3% customer ratings that are low. The higher the bubble, the higher the proportion of low ratings. The size of the bubble represents how many customers are within the segment.

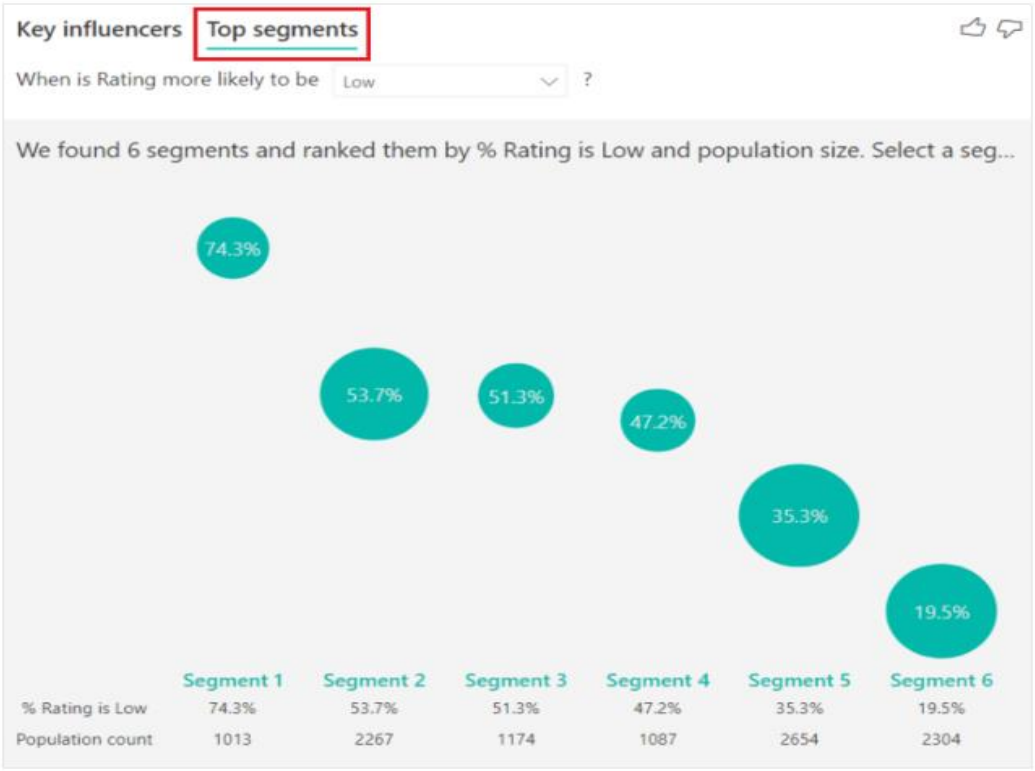

Selecting a bubble displays the details of that segment. If you select Segment 1, for example, you find that it's made up of relatively established customers. They've been customers for over 29 months and have more than four support tickets. Finally, they're not publishers, so they're either consumers or administrators.

In this group, 74.3% of the customers gave a low rating. The average customer gave a low rating 11.7% of the time, so this segment has a larger proportion of low ratings. It's 63 percentage points higher. Segment 1 also contains approximately 2.2% of the data, so it represents an addressable portion of the population.

Adding counts

Sometimes an influencer can have a significant effect but represent little of the data. For example, **Theme** is **usability** is the third biggest influencer for low ratings. However, there might have only been a handful of customers who complained about usability. Counts can help you prioritize which influencers you want to focus on.

You can turn on counts through the **Analysis card** of the formatting pane.

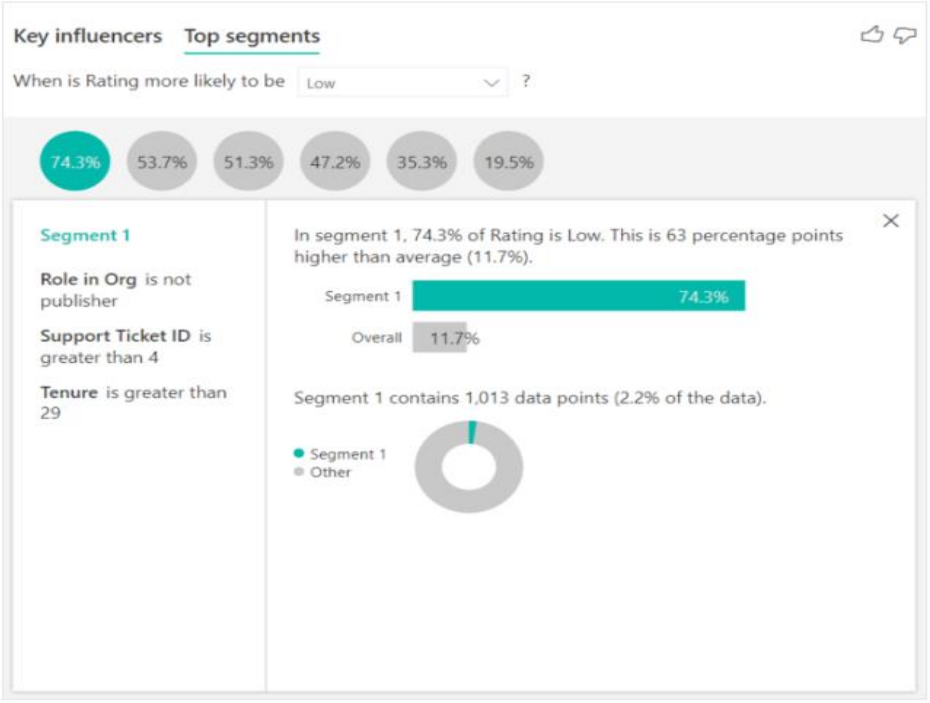

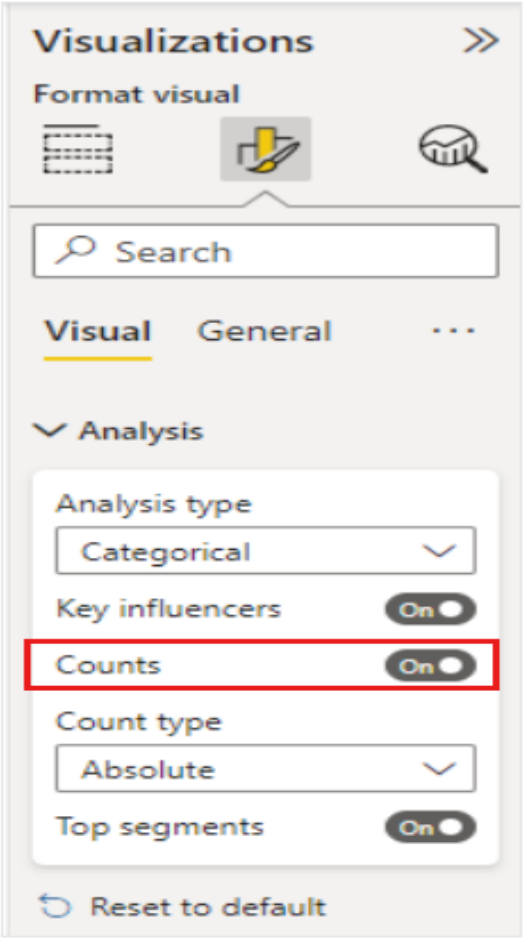

After counts are enabled, you'll see a ring around each influencer's bubble, which represents the approximate percentage of data that influencer contains. The more of the bubble the ring circles, the more data it contains. We can see that **Theme** is **usability** contains a small proportion of data.

You can also use the Sort by toggle in the bottom left of the visual to sort the bubbles by count first instead of impact. **Subscription Type** is **Premier** is the top influencer based on count.

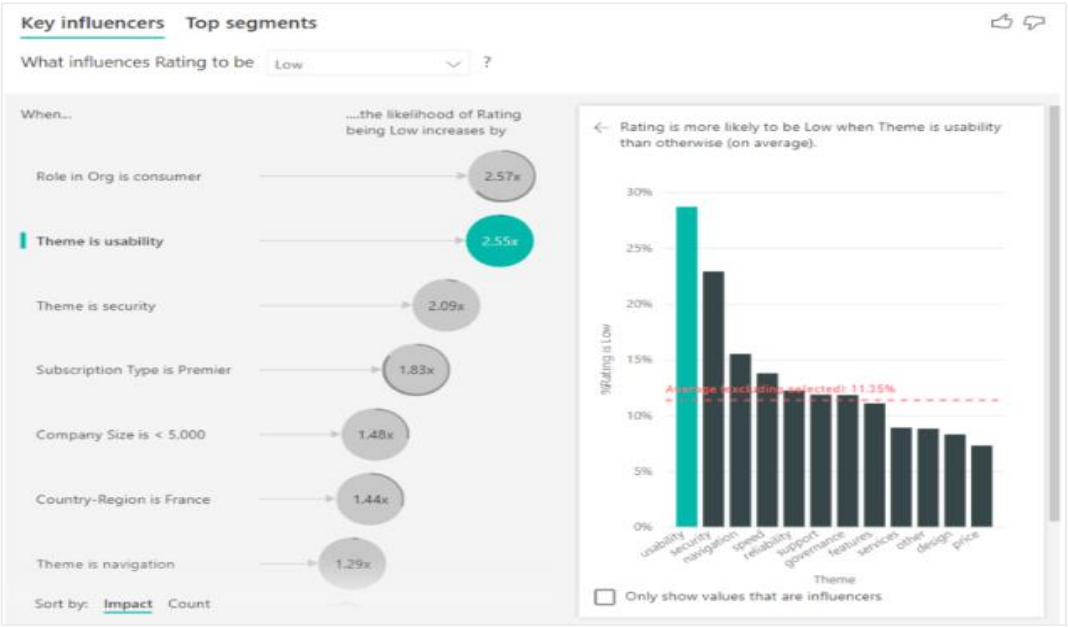

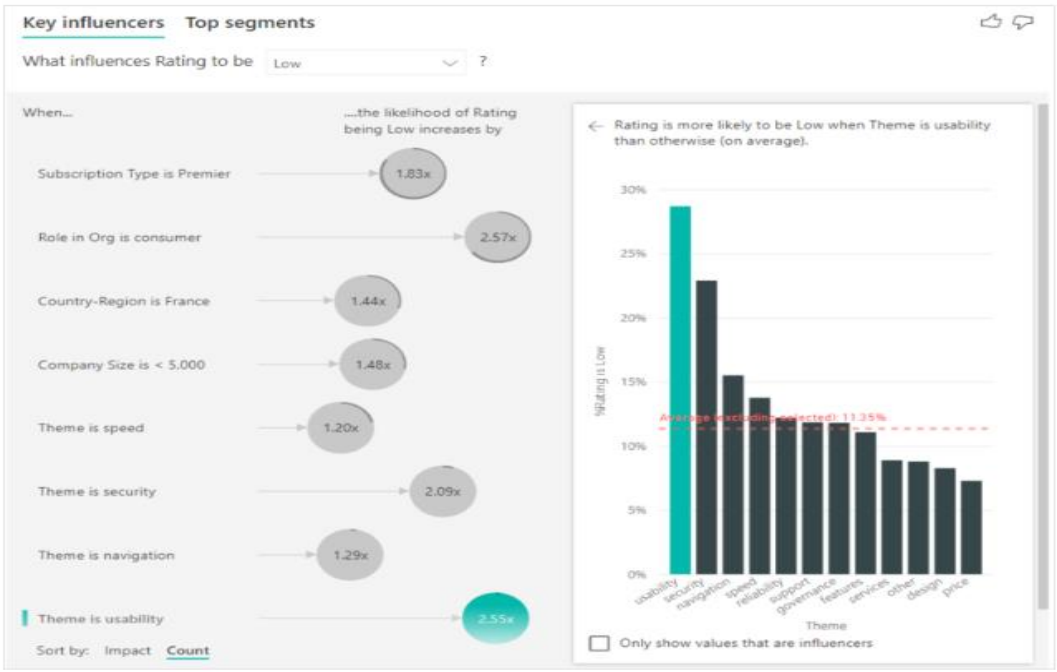

Having a full ring around the circle means the influencer contains 100% of the data. You can change the count type to be relative to the maximum influencer using the **Count type** dropdown in the **Analysis card** of the formatting pane. Now the influencer with the most amount of data will be represented by a full ring and all other counts will be relative to it.

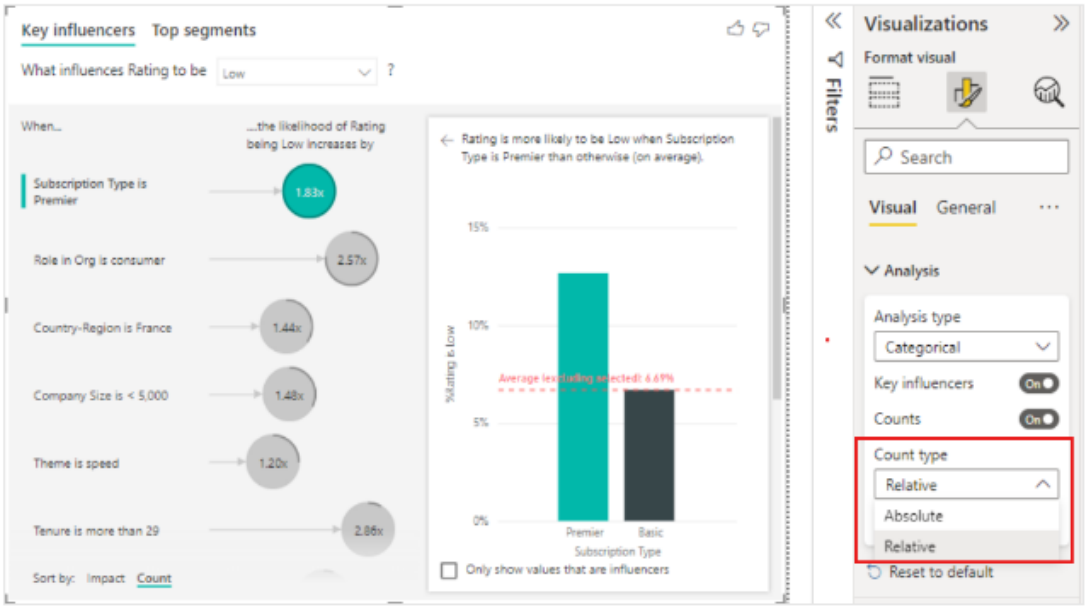

Analyse a metric that is numeric

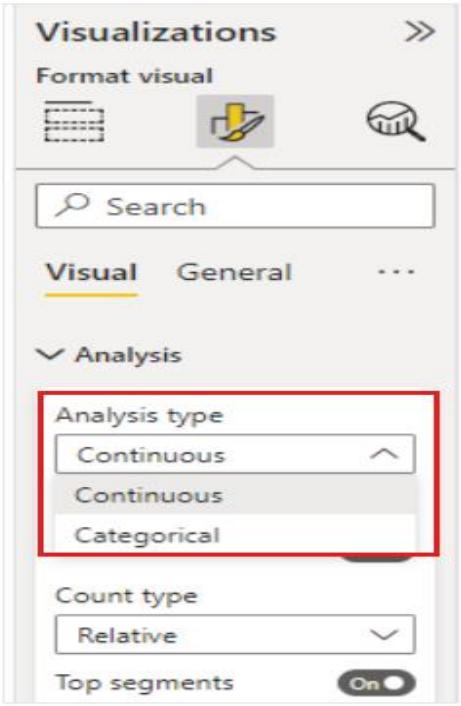

If you move an unsummarized numerical field into the **Analyse** field, you have a choice how to handle that scenario. You can change the behaviour of the visual by going into the **Formatting Pane** and switching between **Categorical Analysis Type** and **Continuous Analysis Type**.

A **Categorical Analysis Type** behaves as described above. For instance, if you were looking at survey scores ranging from 1 to 10, you could ask 'What influences Survey Scores to be 1?'

A **Continuous Analysis Type** changes the question to a continuous one. In the example above, our new question would be 'What influences Survey Scores to increase/decrease?'

This distinction is helpful when you have lots of unique values in the field, you're analysing. In the example below, we look at house prices. It isn't meaningful to ask 'What influences House Price to be 156,214?' as that is very specific and we're likely not to have enough data to infer a pattern.

Instead we may want to ask, 'What influences House Price to increase'? which allows us to treat house prices as a range rather than distinct values.

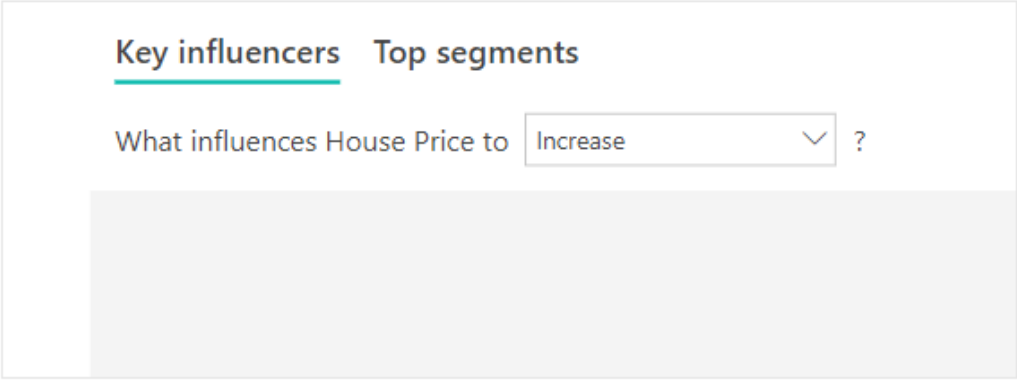

Interpret the results: Key influencers

In this scenario, we look at 'What influences House Price to increase'. A number of explanatory factors could impact a house price like **Year Built** (year the house was built), **Kitchen Qual** (kitchen quality), and **YeaRemodAdd** (year the house was remodelled).

In the example below, we look at our top influencer which is kitchen quality being Excellent. The results are similar to the ones we saw when we were analysing categorical metrics with a few important differences:

- The column chart on the right is looking at the averages rather than percentages. It therefore shows us what the average house price of a house with an excellent kitchen is (green bar) compared to the average house price of a house without an excellent kitchen (dotted line)
- The number in the bubble is still the difference between the red dotted line and green bar but it's expressed as a number ($158.49K) rather than likelihood (1.93x). So, on average, houses with excellent kitchens are almost $160K more expensive than houses without excellent kitchens.

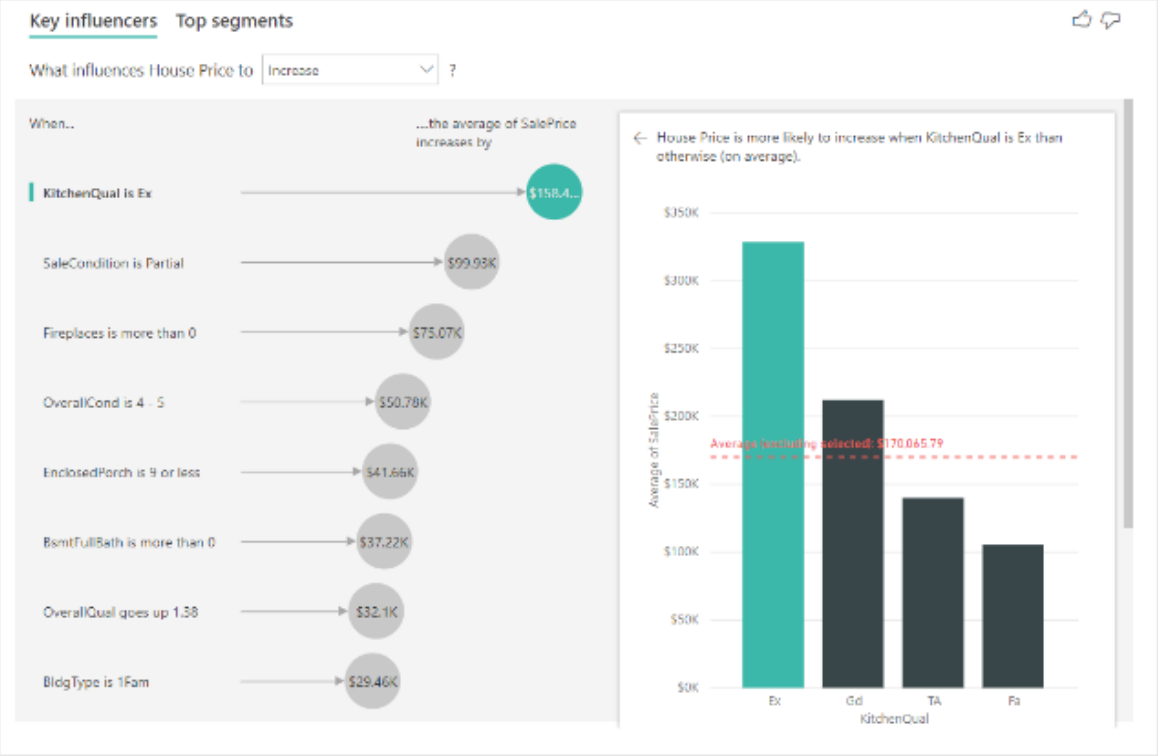

In the example below, we look at the impact a continuous factor (year house was remodelled) has on house price. The differences compared to how we analyse continuous influencers for categorical metrics are as follows:

- The scatter plot in the right pane plots the average house price for each distinct value of year remodelled.
- The value in the bubble shows by how much the average house price

increases (in this case $2.87k) when the year the house was remodelled increases by its standard deviation (in this case 20 years)

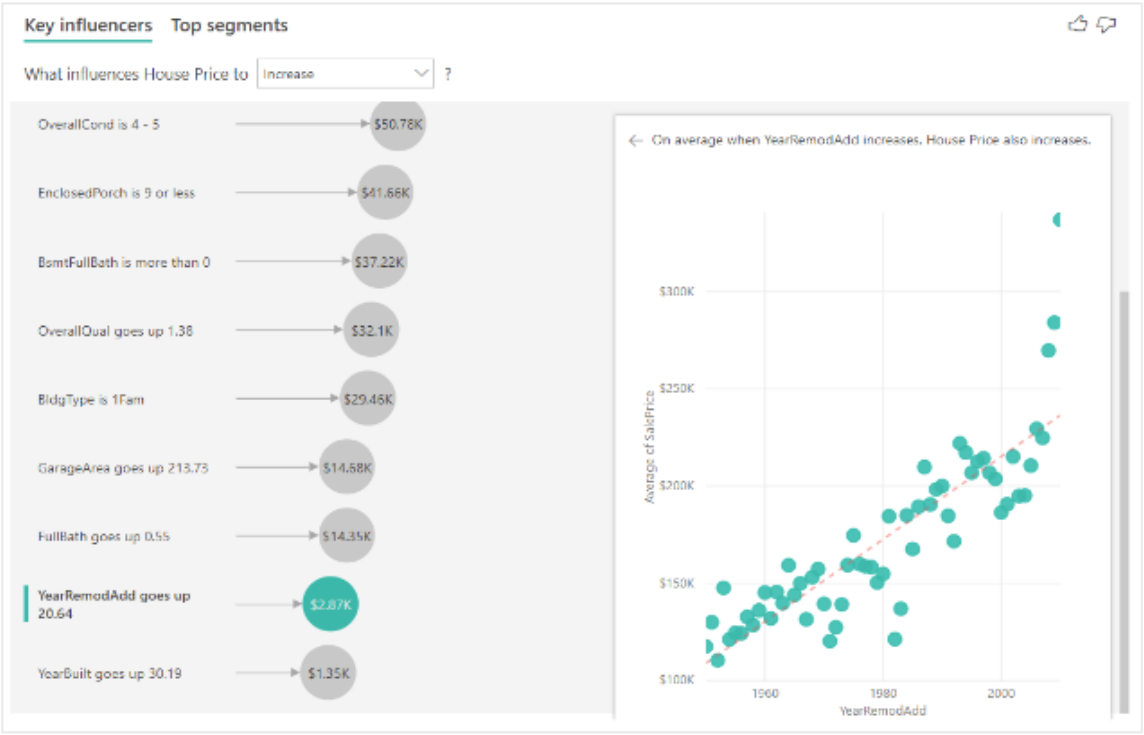

Finally, in the case of measures, we're looking at the average year a house was built. The analysis is as follows:

- The scatter plot in the right pane plots the average house price for each distinct value in the table
- The value in the bubble shows by how much the average house price increases (in this case $1.35K) when the average year increases by its standard deviation (in this case 30 years)

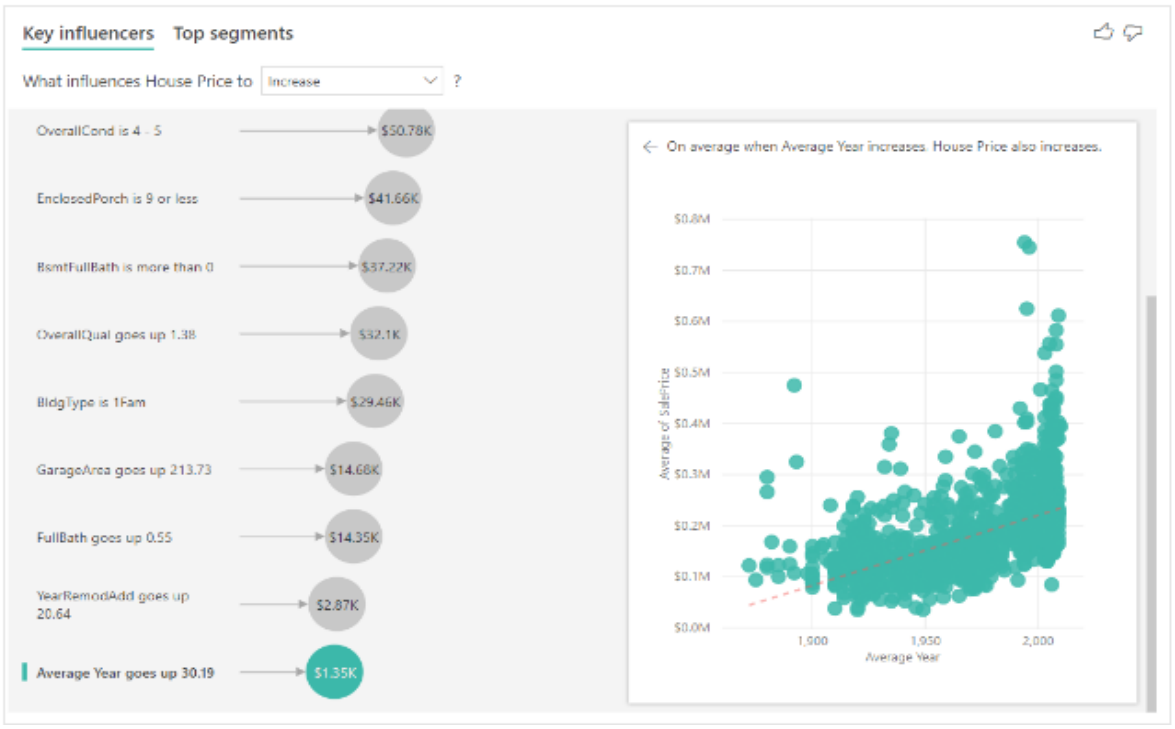

Interpret the results: Top Segments

Top segments for numerical targets show groups where the house prices on average are higher than in the overall dataset. For example, below we can see that **Segment 1** is made up of houses where **Garage Cars** (number of cars the garage can fit) is greater than 2 and the **Roof Style** is Hip. Houses with those characteristics have an average price of $355K compared to the overall average in the data which is $180K.

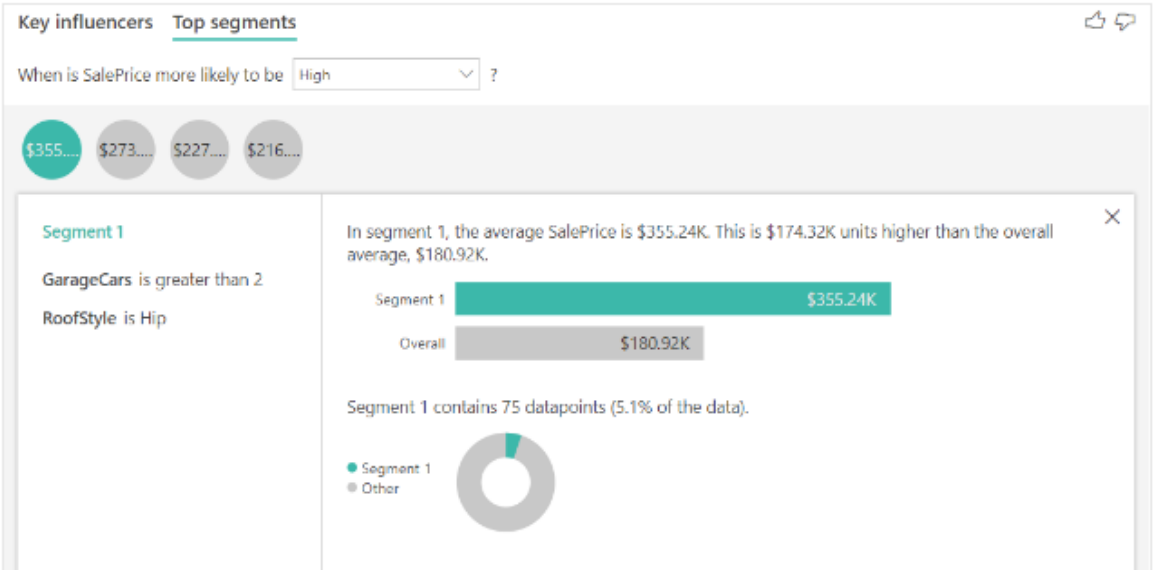

8 Smart Narratives Tutorial - Power BI | Microsoft Docs

The smart narrative visualization helps you quickly summarize visuals and reports. It provides relevant innovative insights that you can customize.

Use smart narrative summaries in your reports to address key takeaways, to point out trends, and to edit the language and format for a specific audience. In PowerPoint, instead of pasting a screenshot of your report's key takeaways, you can add narratives that are updated with every refresh. You can use the summaries to understand the data, get to key points faster, and explain the data to others.

Get started

How to use smart narratives is described below.

Choose the **Products** tab, then select the **Smart narrative** icon in the **Visualizations** pane to automatically generate a summary.

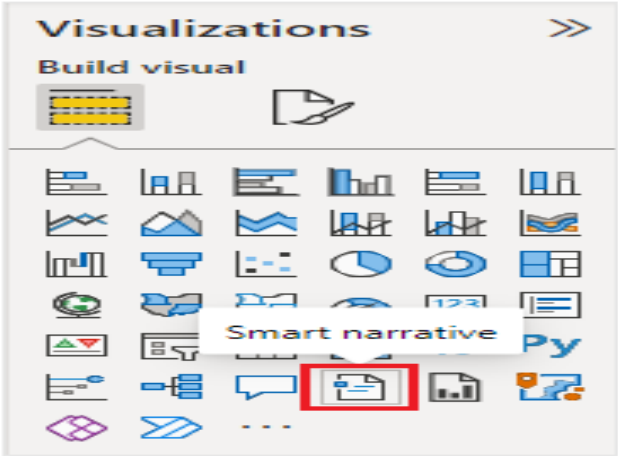

smart narratives can automatically generate a summary of the report's visuals that address revenue, website visits, and sales. To generate a smart narrative of a visualization, right-click it and then select **Summarize**. The smart narrative also shows the expected range of values for these metrics.

Edit the summary

The smart narrative summary is highly customizable. You can edit or add to the existing text by using the text box commands. For example, you can make the text bold or change its colour.

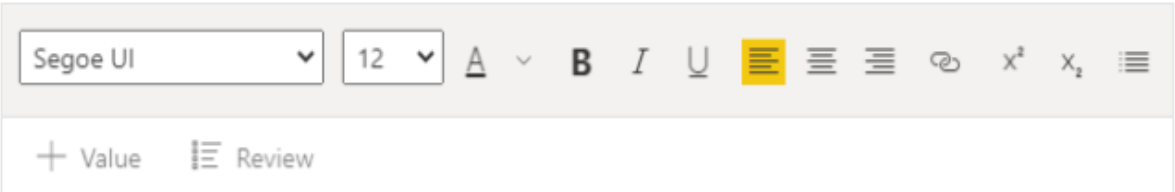

To customize the summary or add your own insights, use *dynamic values*.

You can map text to existing fields and measures or use natural language to define a new measure to map to text.

To format a dynamic value, select the value in the summary to see your editing options on the **Review** tab. Or in the text box, next to the value that you want to edit, select the edit button.

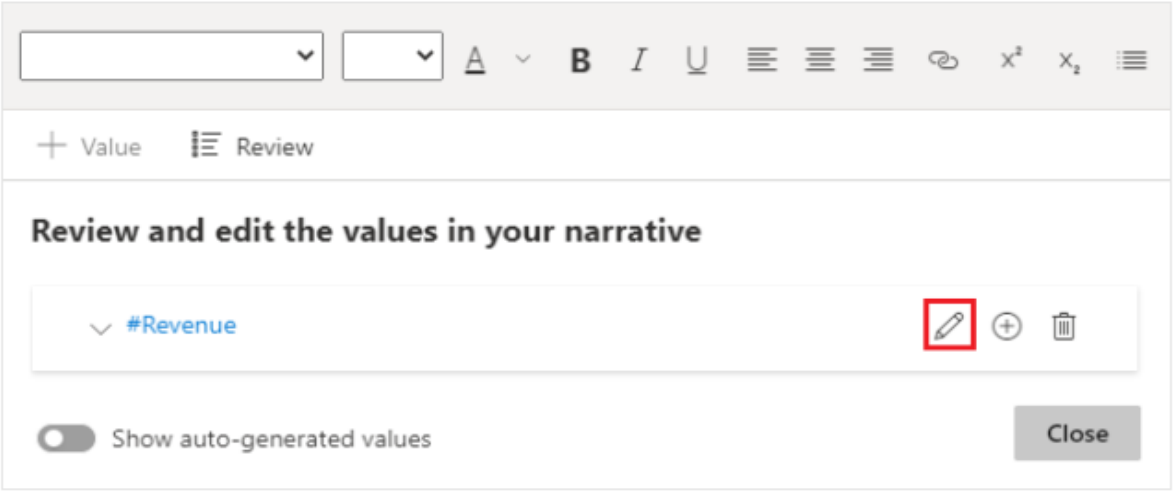

You can also use the **Review** tab to review, delete, or reuse previously defined values. Select the plus sign (+) to insert the value into the summary. You can also show automatically generated values by turning on the option at the bottom of the **Review** tab.

Sometimes a hidden-summary symbol appears in the smart narrative. It indicates that current data and filters produce no result for the value. A summary is empty when no insights are available. Hidden-summary symbols are visible only when you try to edit a summary.

Visual interactions

A summary is dynamic. It automatically updates the generated text and dynamic values when you cross-filter. For example, if you select electronics products in the sample file's donut chart, the rest of the report is cross-filtered, and the summary is also cross-filtered to focus on the electronics products.

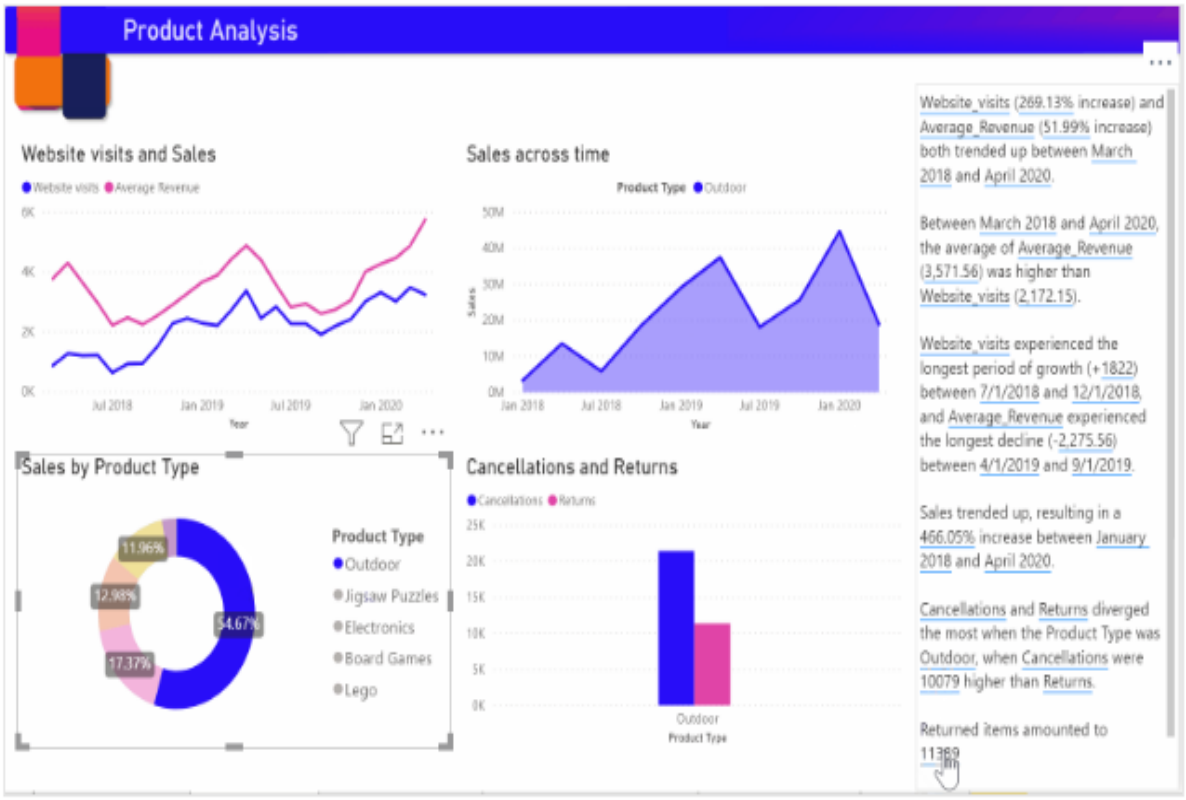

You can also do more advanced filtering. For example, in the sample file, look at the visual of trends for multiple products. If you're interested only in a trend for a certain quarter, then select the relevant data points to update the summary for that trend.

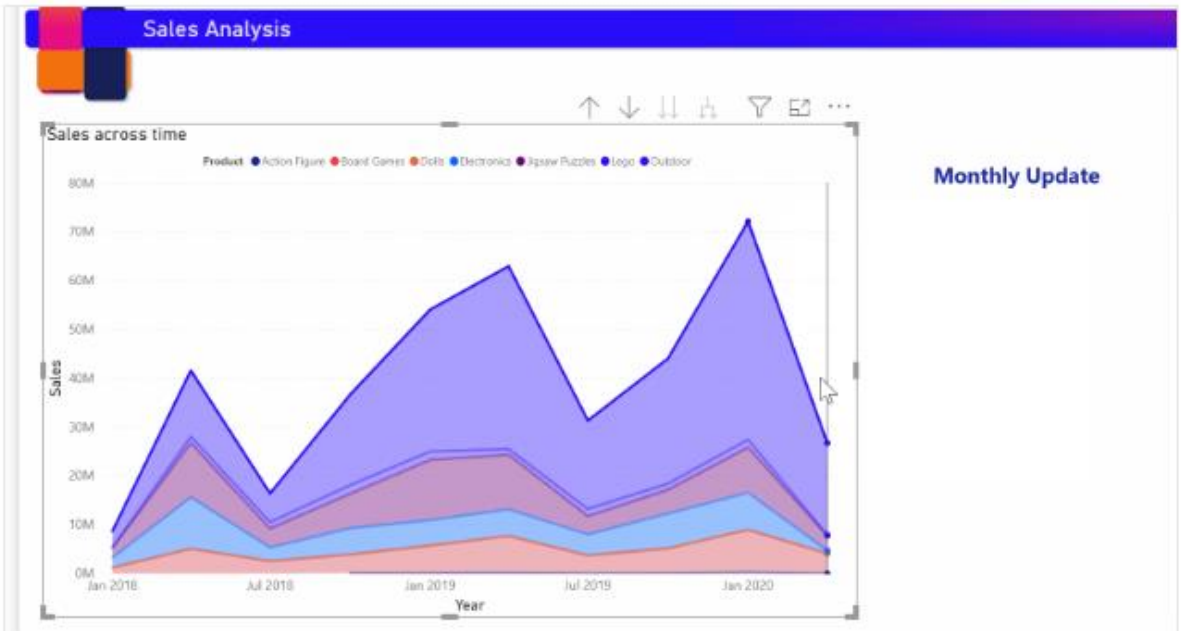

There's a limit to the number of summaries that can be generated so

Smart Narratives picks the most interesting things to summarize about the visual. Smart Narratives generates up to four summaries per visual and up to 16 per page. The summaries that are generated for a page depend on the location and size of visuals and it avoids generating the same kind of summaries for different visuals. Therefore summarizing just the visual can generate more summaries that aren't present while summarizing the whole page.

Considerations and limitations

The smart narrative feature doesn't support the following functionality:

- Pinning to a dashboard
- Using dynamic values and conditional formatting (for example, data bound title)
- Publish to Web
- Power BI Report Server
- On-premises Analysis Services
- Live Connection to Azure Analysis Services or SQL Server Analysis Services
- Multidimensional Analysis Services data sources
- Key influencers visual with a categorical metric or unsummarized numerical field as 'Analyse' field from a table:
  - that contains more than one primary key
  - without a primary key, and measures or aggregates as 'Explain by' fields
- Map visual with non-aggregated latitude or longitude
- Multi-row card with more than three categorical fields
- Cards with non-numeric measures
- Tables, matrices, R visuals or Python visuals, custom visuals
- Summaries of visuals whose columns are grouped by other columns and for visuals that are built on a data group field
- Cross-filtering out of a visual

- Renaming dynamic values or editing automatically generated dynamic values
- Summaries of visuals that contain on-the-fly calculations like QnA arithmetic, complex measures such as percentage of grand total and measures from extension schemas.

PRACTICAL

1. Describe Artificial intelligence workloads and considerations

Artificial Intelligence (AI) is computers thinking and acting in a way that simulates a human. AI is a technology that takes information from its environment and responds based on what it learns. The goal of AI is to create a machine that can mimic human behaviour.

AI is more than learning—it is knowledge representation, reasoning, and abstract thinking. Machine learning (ML) is the subset of AI that takes the approach of teaching computers to learn for themselves, rather than teaching computers all that they need to know. ML is the foundation for modern AI. ML focuses on identifying and making sense of the patterns and structures in data.

ML is about machines' reasoning and decision-making using software that learns from past experiences. ML allows computers to consistently perform repetitive and well-defined tasks that are difficult for humans to accomplish. Over the past few years, machine learning algorithms have proved that computers can learn tasks that are tremendously complicated for machines, demonstrating that ML can be employed in a wide range of scenarios and industries.

AI is now being embedded into the software you use today, sometimes without us realizing it. For example, Microsoft PowerPoint has a feature called Design Ideas that offers suggestions for themes and layouts for slides, and Microsoft Word offers suggestions to rewrite sentences to

improve clarity.

2. Identifying features of common AI workloads

Artificial Intelligence is software that mimics human behaviours and capabilities. Today, software can use AI to automatically detect and predict actions that machines, and humans, should take.

Microsoft Azure provides a set of services for Artificial Intelligence and machine learning that you can utilize to create your own intelligent solutions. Microsoft Azure AI Fundamentals is a certification that requires you to have entry-level knowledge of AI and ML concepts and knowledge of the related Microsoft Azure services.

3. This skill covers how to:

a) Describe Azure services for AI and ML
b) Understand Azure Machine Learning
c) Understand Azure Cognitive Services
d) Describe the Azure Bot Service
e) Identify common AI workloads

a) Describe Azure services for AI and ML

There is a wide and rapidly growing series of services in Azure for AI and ML

- Cognitive Services A set of prebuilt services that you can easily use in your applications.
- Azure Bot Service A service to help create and deploy chatbots and intelligence agents.
- Azure Machine Learning A broad range of tools and services that allow you to create your own custom AI.
- We will be exploring some of the features and capabilities of these three services. However, these services do not work in isolation; they utilize many other Azure services to deliver solutions such as the following:
- Storage

- Compute
- Web Apps
- HD Insights
- Data Factory
- Cosmos DB
- Azure Functions
- Azure Kubernetes Service (AKS)

To explain how Azure services support Azure Machine Learning, consider the scenario of a company that wants to provide recommendations to its users. By providing personalized targeted recommendations, users will more likely purchase more of their products and user satisfaction will increase.

Fig 2.1 shows an example of an ML architecture to support recommendations.

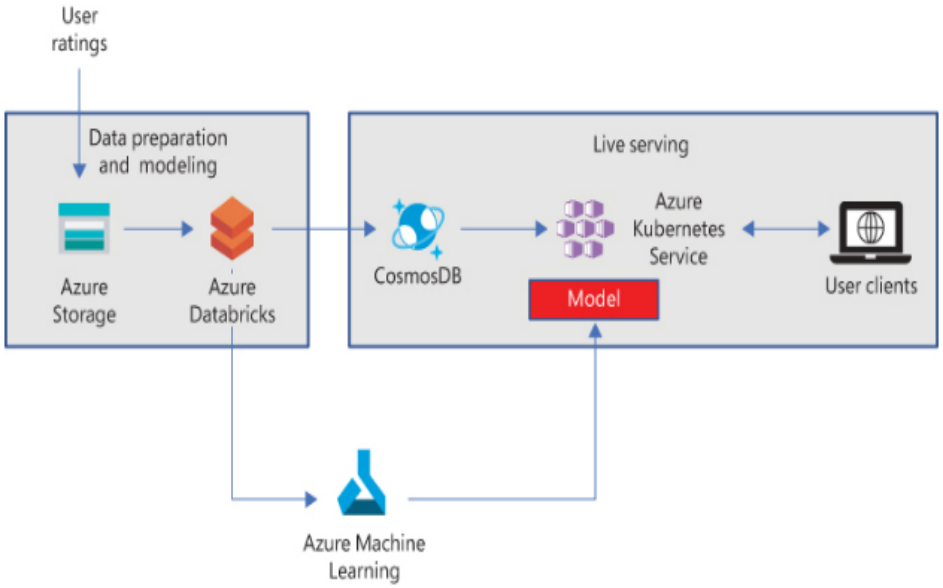

Fig 2.1: example of ML architecture

b) Understand Azure Machine Learning: Azure Machine Learning is the foundation for Azure AI. In Azure Machine Learning, you can build and train AI models to make predictions and inferences. Training a machine learning model requires lots of data and lots of computing resources. Azure provides many services for preparing data and then analysing the data.

cAzure ML is a platform for training, deploying, and managing machine learning models.

c) **Machine learning model types:** Machine learning makes use of algorithms to identify patterns in data and take action. The types of machine learning models created from the outputs of the algorithms are as follows:
    a. Anomaly Detection -Finds unusual occurrences.
    b. Classification- Classifies images or predicts between several categories.
    c. Clustering (including K-Means Clustering)-Discovers structures.
    d. Regression- Predicts values.

d) **Understand Azure Cognitive Services:** Cognitive Services is a suite of prebuilt AI services that developers can use to build AI solutions. Cognitive Services meets common AI requirements and allow you to add AI to your apps more quickly with less expertise. Cognitive Services are machine learning models trained by Microsoft with massive volumes of data. While you can build your own custom AI models to perform the same analyses, Cognitive Services allow you to meet many AI requirements easily around processing images and analyzing text. However, Cognitive Services only address a subset of AI requirements. You can create your own machine learning models to meet more complex and specific needs.
        a. Cognitive Services are available as a set of REST APIs for the following capabilities:
        b. Computer vision, including images and video
        c. Decision, including Anomaly Detector
        d. Language, including text analysis and Language Understanding
        e. Speech, including translation
        f. Intelligent search, including knowledge minin

Cognitive Services have a broad and growing range of features and capabilities. The Azure AI Fundamentals exam focuses on two of these capabilities:

a) Image processing
b) Natural Language Processing (NLP)

A great example of the use of Cognitive Services is the free Seeing AI app that uses these two capabilities. Designed for visually impaired people, this app turns the physical environment into an audible experience, locating faces, identifying objects, and reading documents.

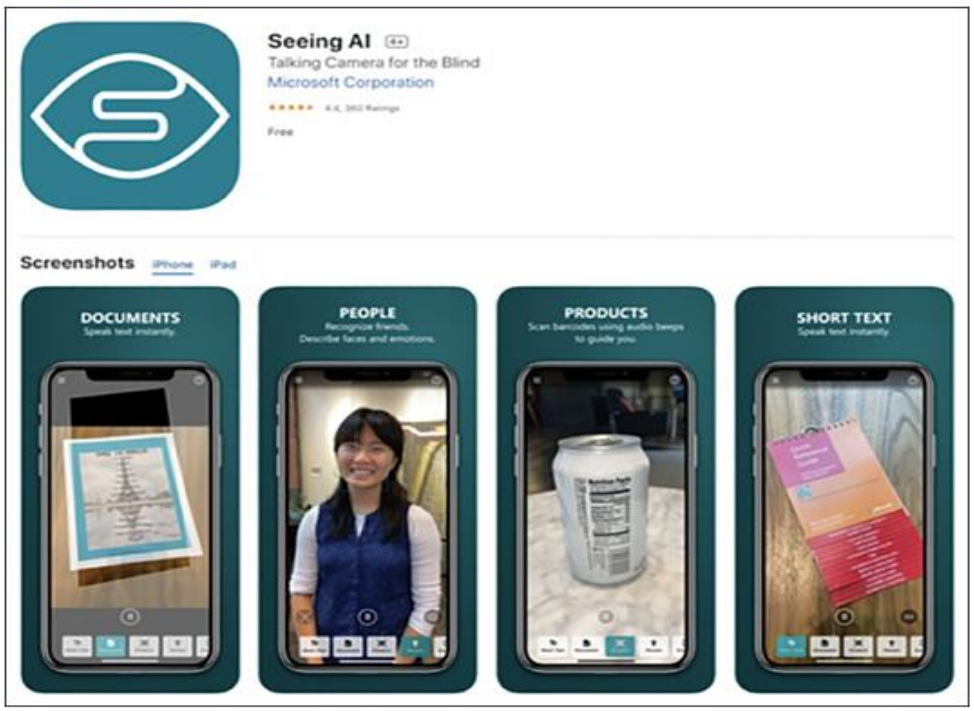

Fig: The Seeing AI app

e) **Describe the Azure Boot Service:** The Azure Bot Service is a cloud-based platform for developing, deploying, and managing bots. Azure Bot Services provide the capabilities known as conversational AI. Conversational AI is where the computer simulates a conversation with a user or customer. Conversational AI has extended beyond simple chatbots to intelligence agents and virtual assistants like Cortana.

There are two conversational AI services included in the Azure AI

- QnA Maker- A tool to build a bot using existing support and other documentation.
- Azure Bot Service Tools to build, test, deploy, and manage bots.

Both QnA Maker and the Azure Bot Service leverage the Language Understanding (LUIS) service in Cognitive Services.

Identify common AI workloads

There are many use cases for AI. Here we will look at some common AI workloads, describing their features and providing some examples for their use.

f) **Conversational AI:** Conversational AI is the process of building AI agents to take part in conversations with humans. Conversational AI is commonly experienced by humans as chatbots on websites and other systems. AI agents (bots) engage in conversations (dialogs) with human users. Bots use natural language processing to make sense of human input, identify the actions the human wants to perform, and identify the entity on which the actions are to be performed. Bots can prompt the human for the information required to complete a transaction.

There are three common types of bot that you may encounter:

- Webchat
- Telephone voice menus (IVR)
- Personal Digital Assistants

3. Describe fundamental principles of machine learning on Azure

Machine learning (ML) is the current focus of AI in computer science. Machine learning focuses on identifying and making sense of the patterns and structures in data and using those patterns in software for reasoning and decision making. ML uses past experiences to make future predictions.

ML allows computers to consistently perform repetitive and well-defined tasks that are difficult to accomplish for humans. Over the past few years, machine learning algorithms have proved that computers can learn tasks that are tremendously complicated for machines and have demonstrated that ML can be employed in a wide range of scenarios and industries.

This chapter explains machine learning algorithms such as clustering, classification, and regression. The chapter then explains how machine

learning works in terms of organizing datasets and applying algorithms to train a machine learning model. The chapter then looks at the process of building a machine learning model and the tools available in Azure.

4. **Describe features of computer vision workloads on azure :** Computer vision is an area of artificial intelligence where software systems are designed to *perceive the world visually* using cameras, images, and video. The challenge here is that humans and computers see different things when they look at the same object. Where a human sees an apple (object), a machine sees an array of pixel values (image *color* data). To give machines a higher-level understanding of what the image data *represents*, we use pixel values as *numeric features* to train a machine learning model. This model behaves like a pattern-detection function, mapping computer-friendly features (pixel values) into human-friendly labels (objects, attributes) in a probabilistic manner. When we feed an input image to this model, it can now *predict* a relevant label with an associated *confidence* value. In some sense, we have taught the computer to "see" the image the way humans would.
5. **Describe features of natural language processing (NLP) workloads on Azure:** Natural language processing (NLP) has many uses: sentiment analysis, topic detection, language detection, key phrase extraction, and document categorization.

   Specifically, you can use NLP to:

   - Classify documents. For instance, you can label documents as sensitive or spam.
   - Do subsequent processing or searches. You can use NLP output for these purposes.
   - Summarize text by identifying the entities that are present in the document.
   - Tag documents with keywords. For the keywords, NLP can use identified entities.
   - Do content-based search and retrieval. Tagging makes this functionality possible.

- Summarize a document's important topics. NLP can combine identified entities into topics.
- Categorize documents for navigation. For this purpose, NLP uses detected topics.
- Enumerate related documents based on a selected topic. For this purpose, NLP uses detected topics.
- Score text for sentiment. By using this functionality, you can assess the positive or negative tone of a document.