

**THE COMPLETE
ARTIFICIAL
INTELLIGENCE
BOOK**



**LUKA ANICIN
CHRIS HAROUN**

The Complete Artificial Intelligence Book
Copyright © 2023 by Haroun Education Ventures, Inc.

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means without written permission from the author.

CONTENTS:

Please note that there are 3 ways to read this comprehensive AI book, depending on what you want to learn. There is an A, B and C track. If you are interested in learning everything about AI, then take the A for ALL track. If you only want to learn Basic introductory topics, then take the B for Basics Track. If you only want to learn only about ChatGPT, then take the C for ChatGPT track.

Per the contents listed below, starting in section 2 at the beginning of each chapter in square brackets, you will see either A for the ALL track in square brackets. B for Basics Track in square brackets or C for ChatGPT in square brackets.

Section 1: Introduction and How to Use this Comprehensive A.I. Book	8
Welcome & How to Use this Book: ChatGPT & MANY More A.I. Topics.....	9
Section 2: Introduction to Artificial Intelligence (Machine + Deep Learning & More)	17
[A/B] Introduction to AI (What, Why and How of Section 2)	18
[A/B/C] 3 Categories of AI: Narrow (ANI), General (AGI) & Super (ASI)	21
[A/B] A.I. vs Machine Learning vs. Deep Learning.....	26
[A] Supervised/Unsupervised Machine Learning/Reinforcement Learning..	29
[A/B] Computer Vision (Using Tesla as an Example)	33
[A] Natural Language Processing (NLP) and NLP Applications	37
[A/B] AI Hardware and Infrastructure: Enabling Scalable and Affordable AI	42
[A] AutoML and Transfer Learning	47

[A/B] AI's Role in Different Industries, Government and Society	51
[A/B] Challenges and Limitations of AI	55
[A/B] Job Positions in the Data World	59
Section 3: ChatGPT	63
[A/B/C] Introduction to ChatGPT (What, Why and How of Section 3)	64
[A/B/C] Introduction to Prompt Engineering (Basic and Advanced Prompts).....	67
[A/B/C] ChatGPT User Interface and Account Setup	70
[A/B/C] Important Problems/Issues with Using ChatGPT	74
[A/B/C] Your First Prompt Exercise - Directional Prompting	76
[A/B/C] Role Playing	79
[A/B/C] Few-Shot Prompting.....	82
[A/C] Chain of Thought Prompting	86
[A/C] In Style of Prompting.....	90
[A/C] Structuring Your Data & Using Tables	93
[A/B/C] Text Summarization.....	96
[A/C] Text Classification.....	99
[A/C] Coding Generation	102
[A/C] Additional ChatGPT Tips.....	105
Section 4: ChatGPT Plugins.....	107
[A/C] Introduction to ChatGPT Plugins	108
[A/C] What are ChatGPT Plugins and How to Get Them	110
[A/C] Browsing with ChatGPT	115
[A/C] Wolfram Plugin for ChatGPT	118
[A/C] edX Plugin for ChatGPT	122

[A/C] KAYAK Plugin for ChatGPT.....	125
[A/C] Zillow Plugin for ChatGPT.....	127
Section 5: Alternatives to ChatGPT (Bard & Bing)	129
[A/C] Introduction to ChatGPT Alternatives.....	130
[A/C] Bard by Google.....	132
[A/C] Bing Chat by Microsoft.....	136
Section 6: Section 6: How to Use Excel and Machine Learning	139
[A/B] Introduction to Using Excel and Machine Learning	140
[A/B] Basic Linear Regression Analysis in Excel	143
[A] Advanced Regression Analysis in Excel.....	151
[A] Multiple Regression Analysis in Excel	182
[A] Fixing Problems with Data	194
Section 7: Images and AI (DALL-E, Midjourney)	196
[A/B] Introduction to Images & AI.....	197
[A/B] Midjourney Account Setup	199
[A/B] Midjourney Interface Walkthrough and How to Use it.....	204
[A/B] Midjourney Good Prompting Practices.....	209
[A/B] Midjourney Image Prompting	213
[A/B] DALL-E 2: Account Setup and Interface Walkthrough.....	218
[A/B] DALL-E 2: Prompt Designs	221
[A] DALL-E 2: Editing Images, Inpainting & Outpainting.....	224
Section 8: Section 8: Voice, Avatars and Cloning.....	228
[A/B] Introduction to Voice, Avatars & Cloning.....	229
[A/B] Using/Making AI Avatars with D-ID.....	231

[A/B] Resemble AI	236
Section 9: Other AI Applications (Ralph AI, Chatbase and More).....	239
[A] Introduction to Other AI Applications	240
[A] Ralph AI	242
[A] Chatbase (How to Add Chatbase GPT to Your Website).....	246
[A] Character AI	251
Section 10: Using AI for Business Decisions	254
[A/B] Introduction to Using AI for Business Decisions.....	255
[A/B] The Role of AI in Business Strategy	257
[A/B] Aligning AI Initiatives with Business Objectives	261
[A/B] Identifying AI Opportunities.....	265
[A/B] Setting AI KPIs	269
[A/B] Building a Data-Driven Organization/Culture	273
[A/B] Selecting the Right AI Solutions and Vendors	279
[A/B] Overcoming Challenges in AI Implementation.....	284
[A/B] AI Ethics, Governance, and Compliance.....	289
[A/B] Data Drift.....	295
[A/B] Risks and Cost of Having AI in Production.....	297
[A/B] Monitoring and Evaluating AI Performance.....	300
Section 11: Additional Risks and Ethics with AI.....	303
[A/B] Introduction to Additional Risks & Ethics with AI	304
[A/B] Bias.....	309
[A/B] Data Privacy and Security in AI.....	314
Section 12: Using AI in Different Industries.....	319

[A/B] AI in Insurance.....	320
[A/B] AI in Finance and Banking	324
[A/B] AI in Retail	329
[A/B] AI in Marketing.....	336
[A/B] AI in Supply Chain.....	342
Section 13: The Future of AI	347
[A/B] Overview of 13 Emerging AI Technologies.....	348
[A/B] Preparing for the Future of AI	352
[A/B] Where We See AI Being Used/Deployed and How?	355

**SECTION 1: INTRODUCTION AND
HOW TO USE THIS
COMPREHENSIVE A.I. BOOK**

***Welcome & How to Use this Book: ChatGPT &
MANY More A.I. Topics***

*“Artificial Intelligence is the renaissance or the golden
age of technology.”*

- Jeff Bezos

Welcome to the Complete Artificial Intelligence Book!



This book provides you with an in-depth understanding of Artificial Intelligence, its applications, and the impact it has on various industries. No prior technology or business experience is required to read this book

There are 13 sections in this book covering all AI topics. Whether you're a business leader, a decision-maker, or just someone curious about the world of AI, this book will give you the knowledge and tools needed to leverage the power of AI in your organization.

There are 2 authors of this book: Luka Anicin and Chris Haroun who teach Artificial Intelligence, technology, business, finance, Excel and personal development courses online.

We are partnering in this book to explain how to use Artificial Intelligence to take your career or business or startup to the next level.

Please note that there are 3 ways to read this comprehensive AI book, depending on what you want to learn. There is an A, B and C track. If you are interested in learning everything about AI, then take the A for ALL track. If you only want to learn Basic introductory topics, then take the B for Basics Track. If you only want to learn only about ChatGPT, then take the C for ChatGPT track and spend most of your time in Sections 3 and 4 of this book where we discuss ChatGPT in a lot of detail.

At the beginning of each chapter, in square brackets, you will see either A for the ALL track in square brackets. B for Basics Track in square brackets or C for ChatGPT in square brackets.

The rest of this introductory chapter will introduce you to what we will learn together in this book, which can help us decide if we want to select the A for all track, B for basics track or C for ChatGPT only track. If you want you can skip the rest of this introduction chapter and start reading the chapters associated with the A, B or C tracks.

ChatGPT is a revolutionary AI tool, but it is just a subset of AI. This book has 13 sections. We are currently in section 1. This book will help you take your career or company to the next level using AI.

Starting in the next chapter is section 2, where we discuss an introduction to AI, as well as Machine Learning and Deep Learning from scratch. **Machine Learning and Deep Learning are subsets of AI.**

Section 2 Overview: Introduction to Artificial Intelligence (Machine + Deep Learning & More)

Section 2 is a top down overview of AI and sections 3 through 13 are bottoms up - meaning we go into depth on how to use many AI

products and how to implement AI in your company or to take your career to the next level.

We will discuss applications of AI using many real world examples for case studies, including Tesla and Netflix. We will talk about the fascinating world of Computer Vision and Natural Language Processing and what the positive impact AI will have on every single industry. We will also discuss challenges and the limitations of AI and all job positions in the data world.

Section 3 Overview: ChatGPT

in section 3 we will learn how to get the most value out of ChatGPT, including more than 1000 prompts to use with ChatGPT that we made for this book.

Section 4 Overview: ChatGPT Plugins

In section 4 we will discuss how to use ChatGPT plugins like Zillow, WolframAlpha and many others. Plugins allow ChatGPT to expand its reach so it is not limited to its knowledge from the training dataset, but external data as well!

OpenAI, which is the creator of ChatGPT allowing companies to connect their product directly to ChatGPT using plugins, is like the moment that Apple introduced the AppStore with the iPhone.

More and more companies will jump on the opportunity to create plugins for ChatGPT and suddenly, many more people will be able to interact with it! We want to help you work smarter and not harder with plugins when it comes to using AI.

Section 5 Overview: Alternatives to ChatGPT (Bard & Bing)

In Section 5 we discuss Alternatives to ChatGPT, including Google and Microsoft's prompt engineering products. Why should we learn about alternatives to ChatGPT? Because there are certain features that ChatGPT alternatives have that you might want to use in order to take your career or your company to the next level.

Section 6 Overview: How to Use Excel and Machine Learning

In Section 6 we discuss How to use Microsoft Excel and Machine Learning. As we mentioned, Machine Learning is a subset of AI that involves developing algorithms to learn patterns from data.

Instead of programming rules, machines learn from experience, just like humans do. We will use Excel in section 6 to do Machine Learning. Why should we use Excel for Machine Learning purposes? Because it's relatively easy to do so as Excel can quickly analyze data and provide charts, regression formulas and advanced statistical analysis.

Also if we start doing Machine Learning with Excel, it's much easier to understand how to interpret or use Machine Learning in other applications. If you don't have Microsoft Excel, if you want you can search for a free trial of Office 365 and download Excel or subscribe to Office 365 during the time you are reading this book.

We will also teach you how to add Microsoft's Data Analysis plug in during this section, so that you can do regression analyses.

Section 7 Overview: Images and AI (DALL-E, Midjourney)

In Section 7, we will discuss how to create incredible images using Artificial Intelligence. Image generation using AI has been so disruptive, that many graphic artists that used to charge more than

\$40,000 to make video game images have seen their business fall off a cliff.

We will learn how to use top AI image products like DALL-E, and Midjourney. Why are we going to discuss using AI to make images? Because as a society we are much more visual now, which is why Meta bought Instagram and why products like Pinterest exist.

If we learn how to quickly and easily make high quality images using AI, then we can take our personal brand or corporate brand to the next level with social media and we can make much more impactful our documents, presentations, website designs and much more.

Section 8 Overview: Voice, Avatars and Cloning

In Section 8 we will discuss useful AI tools that can help you with creating voice clones of yourself and tools that can help you with creating digital avatars. We will do this by learning how to use D-ID and Resemble AI.

This is important to learn because generating voices that sound almost exactly as we do can be very useful in the business environment. We could generate samples of our voices and more and use them to narrate marketing campaigns, product reviews and even create tutorials!

We can even use these technologies for customer support on the phone or on the Internet. As dangerous as these tools can be, from an ethical standpoint, they are going to be used a lot in the future. Movies will be made with lifelike avatars and we can create our own lifelike avatars in the future to help us create promotional trailers or to advertise our companies.

Section 9 Overview: Other AI Applications (Ralph AI, Chatbase and More)

In section 9, we will discuss additional AI products and applications that we found useful for you and your business. In section 9 we will discuss what Ralph.AI is and its benefits over normal ChatGPT.

In addition, we will discuss using Chatbase in order to add a chat feature to your website.

And we will discuss Character AI, which generates human like responses.

Section 10 Overview: Using AI for Business Decisions

In section 10, we will discuss using AI for business decisions. We will discuss the Role of AI in Business Strategy, Setting AI key Performance Indicators, Selecting the Right AI Solutions and Vendors, Overcoming Challenges in AI Implementation, Compliance, Risks, Monitoring AI performance and much more.

It's important to discuss these qualitative business strategy topics when it comes to AI, because implementing an AI solution can be an ethical and compliance nightmare if not done correctly. We also need to make sure that we are partnering with the appropriate AI vendors, which we will discuss and how to anticipate AI challenges before they occur.

Section 11 Overview: Additional Risks and Ethics with AI

In section 11 we will explore additional ethics and risks with AI. Because failing to consider risk factors with AI can lead to negative consequences, such as biased decision-making, privacy violations, and even harm to individuals or society.

Section 12 Overview: Using AI in Different Industries

In section 12 we will discuss using AI in different industries like in the retail and financial services sectors and also when it comes to marketing and supply chain management.

Section 13 Overview: The Future of AI

In section 13 we discuss the future of AI and next steps.

Thanks and let's begin :)

**SECTION 2: INTRODUCTION TO
ARTIFICIAL INTELLIGENCE
(MACHINE + DEEP LEARNING &
MORE)**

***[A/B] Introduction to AI (What, Why and How
of Section 2)***

Welcome to Section 2 where we discuss an Introduction to Artificial Intelligence

In terms of the what, why and how of section 2, let's kick it off with the what. What are we going to learn together? We will learn about the current applications of AI through what is called Artificial Narrow Intelligence and the next generation of AI, which is called Artificial General and Artificial Super Intelligence.

We will also discuss Machine Learning and Deep Learning, which are subsets of AI. We will talk about the fascinating world of Computer Vision and Natural Language Processing and what the positive impact AI will have on every single industry. We will also discuss challenges and the limitations of AI and all job positions in the data world.

Moving on to why. Why are discussing high level AI topics in section 2? Think of section 2 as being a top down overview of AI. We are doing this to help us better understand the concepts taught in this book after section 2.

After section 2 we will do a more bottom's up thorough analysis on AI tools and how to use them to take your career or business to the next level, including the basic and much more advanced features of prompt engineering services like ChatGPT and alternatives to ChatGPT and how to use Excel for Machine Learning and AI, as well other topics we will discuss after section 2, including voice and image based AI as well as AI for business decisions and many other AI applications.

Lastly How? How are we going to learn about the introduction meaning top down aspects of AI? We will discuss many real life case studies from Tesla to Netflix in order to reinforce what we have

learned in an edutaining meaning educational and entertaining way.

If you don't want to learn about the basics of AI and you just want to jump right in and learn about how to use AI you can also jump to any section in this book. Thanks

***[A/B/C] 3 Categories of AI: Narrow (ANI),
General (AGI) & Super (ASI)***

There are 3 types of AI, meaning 3 Artificial Intelligence categories that the industry refers to. On one side of the AI spectrum is weak AI, then we have strong AI in the middle and we have super strong AI on the right side.



The first category of AI is Artificial Narrow Intelligence - meaning ANI which is also called weak AI. Artificial Narrow Intelligence is designed for relatively simple tasks from an AI perspective like basic voice based technologies, self-driving cars, and movie recommendation engines.

ANI is relatively easy to develop and implement. However, ANI has 2 primary limitations:

- 1: It can't learn beyond its specific task.
- 2: It can't adapt to new tasks or situations without reprogramming.

Examples of ANI include Siri and Alexa.

Moving onto the second category of AI, which is Artificial General Intelligence meaning AGI which is also called Strong AI. The goal of AGI is to mimic human intelligence. This type of AI can learn, reason, and adapt to new situations on its own. There really isn't an example of AGI that exists yet and there are plenty of science fiction movies that have examples of strong AI from the movie 2001 a Space Odyssey to one of my favorite video games called Detroit become human. AGI is the scientific equivalent of the splitting of the atom.

It's important to think about science fiction movies when it comes to technology products and AI because they influence the future of technology investment in the venture capital sector. For example, the Steven Spielberg movie Minority Report and Avatar influenced the virtual reality, augmented reality and screen mirroring technology and metaverse sectors.

Star Trek influenced the mobile market with flip phones and the movie "Her" with Scarlett Johansson influenced Siri and Amazon Alexa development. These Hollywood *dreampreneurs* influence what Silicon Valley or other global tech sectors will develop in the future. The movie 2001 a Space Odyssey had a strong AI example with a computer that could think for itself called HAL. If you add one letter to H.A.L., it is I B M.

Many people are so excited about ChatGPT because many think that ChatGPT is the closest thing to AGI, since we can ask basically whatever we want. ChatGPT is not yet really AGI but it is getting closer and many venture capitalists think that this is the most astounding technology breakthrough since the Internet was

invented. We discuss ChatGPT in much more detail in sections 3 and 4.

The Advantages of AGI are:

1: It can learn and adapt to a wide range of tasks. ChatGPT can adapt based on the questions you ask, but it is not true AGI yet.

2: It can improve its performance autonomously.

But there are challenges in Developing AGI as follows:

1: It requires significant advancements in AI technology.

2: It raises ethical concerns.

Once AGI becomes a reality all companies and governments will have AI ethical positions like the Department of AI and Ethics or Chief AI Ethics Officer.

And we are not far from this prediction at all. Currently we have a number of highly skilled research centers working on developing AI, such as OpenAI, DeepMind, and big companies like IBM, Microsoft and Google.

Today, these companies must have external ethical researchers working with them constantly, to make sure that these companies are developing AI in the right way. These ethical researchers are coming from two organizations. One is called MiRi and the other is the Future of Humanity Institute led by Nick Bostrom.

If you would like to read further on AGI risks, benefits and potential outcomes we recommend reading the Superintelligence book by Nick Bostrom.

The third and final AI Category is called Artificial Super Intelligence (ASI), which means that AI would be smarter than humans in almost every way. While this concept is speculative, it's widely discussed among AI researchers and futurists.

In terms of what the Potential Implications of ASI are, It could lead to unprecedented technology advancements and innovation. But before Artificial Super Intelligence is real we need to see many AI research breakthroughs.

There are a plethora of ethical and other risks with ASI if AI surpasses human control from an ethics, safety and regulatory perspective. Movies like terminator with the Skynet ASI is an extreme example of the ethical issues with ASI. Warren Buffett even said that once we open AI's pandora's box, we need to be careful. We have to make sure that future AI products don't get too much autonomy.

Now that we understand the 3 types of AI meaning ANI, AGI and ASI, in the next few chapters we'll talk about AI in more detail, Machine Learning, and Deep Learning, and how they relate to each other.

***[A/B] A.I. vs Machine Learning vs. Deep
Learning***

Now that we understand the 3 different types of AI, let's discuss differences between AI, Machine Learning, and Deep Learning, and see how they fit together.

Machine Learning is a subset of AI that involves developing algorithms to learn patterns from data. Machines learn from experience, just like humans do. In section 6 we will discuss how to use Excel to perform Machine Learning using Linear Regression and other methodologies to find patterns in data.

Deep Learning (DL) is a subset of Machine Learning (ML). Deep Learning uses Artificial Neural Networks to model and solve complex problems. It's inspired by the structure and function of the human brain.

Deep Learning is more complex than Machine Learning and requires better hardware. Let's look at some real life examples to understand the differences between Machine Learning and Deep Learning:

Example 1 is a Movie Recommendation Systems. An AI-based system might use predefined rules to suggest movies based on genre or director preferences. A Machine Learning system would analyze user behavior, such as ratings and watch history, to make personalized recommendations. A Deep Learning system could analyze unstructured data like movie plot summaries to provide even more accurate suggestions.

Example 2 is Text Analysis. AI text techniques might involve keyword searches and text matching. Machine Learning algorithms can analyze text data to discover patterns and relationships, like sentiment analysis or topic modeling. Deep Learning models, such as Recurrent Neural Networks, can process complex language

patterns to understand context and generate human-like advanced intelligent responses.

In sections 3, 4 and 5 we will use ChatGPT, Bard AI by Google and Bing Chat by Microsoft which use Machine Learning to create human-like text responses.

Example 3 is Image Recognition. AI-based image recognition systems use algorithms to identify shapes and colors and more in images. Machine Learning systems use pattern recognition techniques to classify objects in images like that's a dog and to create images. Deep Learning systems use Convolutional Neural Networks to analyze images at a granular level, making them highly accurate in tasks like facial recognition.

In section 7 we will use DALL-E and Midjourney, which use Machine Learning to create images.

In summary, AI is the broadest concept, encompassing Machine Learning and Deep Learning. Machine Learning is a way of achieving AI by training models with data, and Deep Learning uses neural networks to solve more complex problems.

Understanding the relationship between AI, Machine Learning and Deep Learning and their real-world applications is crucial for business leaders when evaluating AI solutions.

***[A] Supervised/Unsupervised Machine
Learning/Reinforcement Learning***

When we work in a company and we are analyzing data our boss supervises us and tells us to use certain columns of data that have labels like the date in one column and sales in another column of data.

When we are babies we don't have structure because we crawl around and analyze everything without supervision...as babies we aren't taught what columns of data are with labels. In this chapter, we'll focus on three main categories of Machine Learning and how they can be applied to business use cases.

Category 1 of 3 of Machine Learning is called Supervised Learning. With Supervised Learning we provide the computer with inputs and then the computer uses Machine Learning and then provides us with the output. The machine can use pattern recognition to create an equation (i.e., a Regression Analysis) with the input of X, for example in order to solve for Y.

Examples of Regression Analysis include:

Stock price forecasting: meaning predict the future price of a stock based on historical data and other inputs.

Demand forecasting,

Price optimization,

Sales forecasting,

Predicting customer churn or turnover,

Credit scoring in banks,

Product recommendations,

and Customer lifetime value predictions

Supervised learning tasks are the easiest ones to solve. And companies wanting to apply Machine Learning and AI in their processes, always start with the question - can we structure this problem as a supervised problem? Because of this, most Machine Learning projects in the world, at this moment, can be classified as a supervised learning task.

We will speak later in the book about the Machine Learning Project Lifecycle in section 13 and we will have a lot of interesting examples, exercises and questions to ask ourselves if we decide to go use Machine Learning in our business.

Category 2 of 3 for Machine Learning is called Unsupervised Learning. Unsupervised Learning is a subset of Machine Learning where a computer is given a lot of data and it is asked to find patterns or clusters. For example, these types of customers prefer these types of products.

Unsupervised learning algorithms work with unlabeled data, aiming to discover underlying patterns or groups meaning clusters within the data.

This type of learning can be useful for businesses in tasks like:

Customer segmentation - meaning are these types of customers more likely to buy certain types of products.

Another example is anomaly detection in financial transactions like fraud or the securities and exchange commission analyzing data to see if hedge funds are transacting on insider information.

Category 3 of 3 is called Reinforcement Learning:

In reinforcement learning, which is the 3rd category of Machine Learning, an agent like a computer or a robot learns to make decisions by interacting with an environment. The agent receives feedback in the form of rewards or penalties, and its goal is to maximize the cumulative reward over time. Potential business applications include:

Optimizing supply chain decisions

Personalizing customer experiences

Developing trading strategies in finance with the reward of course meaning making as much money as possible without taking on too much risk.

In practice, Machine Learning models are trained on historical data to make predictions or decisions. As a business leader, it's crucial to ensure that the data used for training is representative, accurate, and free from biases that may impact the model's performance.

Finally, please keep in mind that Machine Learning models might require regular updates and maintenance as new data becomes available or the business context changes. For example, YouTube and Google and Facebook often change their algorithms and any Machine Learning process based on algorithms that change might also need to be updated

[A/B] Computer Vision (Using Tesla as an Example)

Computer Vision is a field of Artificial Intelligence that enables computers and machines to interpret and understand visual information from the world, such as images, videos, or real-time camera feeds. It's like teaching a computer to "see" and make sense of the visual world, just as humans do.

The primary goal of computer vision is finding various techniques for manipulating and extracting information from images to make actionable insights for our business. And with this we can achieve things such as classifying objects on an image, detecting faces, detecting moving objects, pedestrians and other items.

An amazing example of a company that utilizes Computer Vision is Tesla. Tesla has one of the best AI teams in the world and Computer Vision systems that came from it. Instead of using lidars and radars to navigate their surroundings, Tesla cars use only 8 camera positions around the car for Computer Vision purposes.

Besides the car application, Tesla is using Computer Vision for factory optimization, with the goal to speed up the production lane and ultimately produce more cars. For example, in many Tesla factories a lot of the work is done by robots.

Let's explore a few applications from Computer Vision and industries using them today:

One of the most common applications is image classification. We are basically creating a system that, given an image, predicts what's on it. Common applications are on blogs to predict if someone has uploaded a non safe for work image. Or to create a system that detects American Sign Language.

Another application is object detection and object segmentation. Instead of just saying - this is an image of a pedestrian, we would like to know exactly where in the image that pedestrian is located.

Imagine creating a self-driving car, and if we used image classification, whenever a pedestrian would appear in the frame, our car would stop moving. But if we know that a pedestrian is there and that they are on the sidewalk (not the road), then we can proceed safely forward.

One more example of image segmentation is, an autonomous drones identifying and avoiding obstacles in its path while delivering packages or a robot at an Amazon warehouse packing boxes.

Let's explore a few industries and how they were revolutionized because of AI and Computer Vision:

In Social Media: Platforms such as TikTok, Instagram and Facebook are analyzing every single image that we upload so it becomes searchable.

In the Retail sector, Computer Vision is used for inventory management, shelf monitoring, tracking buyer behavior, product recommendation, and theft detection by tracking products and recognizing unusual activities in real-time.

In Marketing and advertising: Businesses can use computer vision to analyze customer behavior and preferences from in-store camera feeds or social media, allowing them to create more targeted and effective marketing campaigns.

In Human resources: Companies can use computer vision algorithms to analyze video interviews, facial expressions, and body language, helping them make more informed hiring decisions.

***[A] Natural Language Processing (NLP) and
NLP Applications***

Natural Language Processing, also called NLP, is a branch of Artificial Intelligence that focuses on enabling computers to understand, interpret, and generate human language in a way that is both meaningful and useful.

NLP uses various tasks, including sentiment analysis, text classification, named entity recognition, machine translation, and question-answering systems, allowing computers to efficiently process and analyze large volumes of text data and communicate with humans more naturally.

Let's go through the aforementioned tasks that in more detail when it comes to NLP meaning Natural Language Processing:

Sentiment Analysis involves determining the sentiment or emotion expressed in a piece of text, such as identifying whether a statement is positive, negative, or neutral.

Sentiment Analysis can be used for Customer Experience Analysis; companies can use Sentiment Analysis to evaluate customer feedback from various sources, like surveys, social media, and product reviews, helping them identify areas for improvement and customer satisfaction trends.

Sentiment Analysis can also be used for Brand Reputation Monitoring; businesses can use Sentiment Analysis to track public opinion about their brand, products, or services, allowing them to respond proactively to negative sentiment or capitalize on positive sentiment in their marketing strategies.

The next NLP task is Text classification, which is the process of categorizing text into predefined groups based on its content or context. Text classification can be used for Spam Detection; email service providers can use text classification to automatically identify and filter out spam messages, ensuring a cleaner inbox experience for users.

Text classification can also be used for resume screening; companies can use text classification algorithms to automatically categorize and filter job applications based on predefined criteria, streamlining the initial screening process for human resources teams.

If we think like a computer does, then we will include more keywords on our resumes and on our LinkedIn profile Job Title and Summary Section.

The next NLP task is called Named Entity Recognition or NER and it involves identifying and classifying specific entities, such as names of people, organizations, locations, or dates, within a given text.

Named Entity Recognition can be used for financial news analysis; financial institutions can use NER to extract key information like company names, stock symbols, and financial figures from news articles, enabling quicker and more informed decision-making in trading and investment.

Named Entity Recognition is also used in Intellectual Property Management; law firms and businesses can use NER to identify specific entities like patent numbers, inventor names, and technology terms from legal documents, helping them manage and analyze their intellectual property portfolios more efficiently.

Moving on to the next NLP task, which is Machine Translation, which is the process of automatically converting text from one language to another using NLP algorithms.

Machine Translation is used in Global Customer Support; companies can use Machine Translation to provide real-time support to customers in multiple languages, broadening their customer base and improving user satisfaction.

Machine Translation is also used in Cross-Language Collaboration; multinational corporations can use Machine Translation to facilitate communication between employees who speak different languages, enabling more efficient collaboration across departments and regions.

The final NLP task is Question-Answering Systems, which use NLP techniques to understand user queries and provide relevant answers or information based on a given text or knowledge base quite often through customer service chat bots, which can be frustrating to use at times, but they can save companies a lot of money. Companies can implement AI-powered chatbots to handle frequently asked questions and provide instant assistance to customers, reducing the workload on human support agents and improving response times.

Another example of Question-Answering Systems is an Internal Knowledge Base; businesses can use Question-Answering Systems to create an intelligent internal knowledge base that helps employees find answers to their questions quickly and efficiently, streamlining operations and improving productivity.

The bottom line is that by understanding which tasks make up NLP and how they can be used in business, you're able to have a clearer

picture of AI's possibilities in your company's functions so we can spend more time in innovation and less time on more mundane operational tasks.

***[A/B] AI Hardware and Infrastructure:
Enabling Scalable and Affordable AI***

Let's discuss AI hardware and infrastructure and enabling scalable and affordable AI, as well as providing some examples so that we can fully understand the technology.

Running Machine Learning models in production involves many costs, including hardware expenses, energy consumption, data storage, and cloud service fees, as well as costs related to maintenance and support.

In terms of Accessible Computing Power, many organizations opt for cloud-based services for AI, like Amazon Web Services, Google Cloud, or Microsoft Azure, in order to access powerful computing resources without the upfront investment in hardware.

The aforementioned Cloud service providers offer pay-as-you-go pricing models, with costs depending on the type and number of instances, storage, and data transfer.

We have to be mindful though because using specialized instances with GPUs or TPUs, which we will explain in a moment, can increase the cost significantly compared to standard instances with CPUs.

The availability of powerful and cost-effective computing resources, such as GPUs, CPUs, TPUs and cloud computing services, has made it easier to train and deploy AI models on a large scale. These resources enable faster processing and reduced training times, making AI more affordable and accessible to a broader audience.

GPUs stands for Graphics processing Units. GPUs were used mainly for 3D graphics in video games for example, but GPUs are so good at parallel computing, that they are used a lot with AI.

TPU stands for Tensor Processing unit, which was invented by Google for use in AI and in other applications. The great thing about TPUs is that they are designed in a way to use less processing power than GPUs or CPUs. A TPU can be used with Google Cloud services, which is Google's version of Amazon Web Services, meaning cloud computing.

Interestingly Google has created a completely free and open source library to help us develop AI algorithms called TensorFlow. But to get its full capacity, you needed to utilize it with TPU hardware. This model worked at the beginning, while TensorFlow was the only library used for that type of programming. Now that we have multiple instances of them, completely for free, most companies today are using GPUs and other services, getting really similar performance, for less money.

CPUs are not as efficient with AI when compared to GPUs and TPUs, although CPUs can still be used for smaller-scale Machine Learning projects and basic data processing tasks.

The great thing about the many cloud computing AI offerings today is the competition that creates accessible and cheap computing power for all of us. Even a small startup can use cloud-based computing resources to develop and deploy their AI-driven solutions, much cheaper than before.

For example, if a company were to create a recommendation system for their e-commerce site, they could do it simply by registering on Amazon or Google's cloud platforms, uploading their data and training their algorithm - without investing in expensive hardware.

Similar to cryptocurrency mining, a large expense with AI is Energy Consumption; running Machine Learning models, particularly Deep Learning models, requires a considerable amount of power. If you're running the hardware on-premise and not renting from the cloud, you should factor in energy consumption costs. For cloud-based services, these costs are generally included in the instance pricing.

Another expense to be mindful of is Data Storage and Transfer prices; storing and transferring large datasets can be costly, particularly when using cloud services. Providers usually charge for storage in GB per month and data transfer in GB transferred in and out.

Software and Platform Licenses is another expense to consider. While many Machine Learning libraries and frameworks are open-source and free, some commercial software or platforms may require a license fee. These fees can range from a one-time purchase to a recurring monthly subscription.

What has also helped to enable scalable and affordable AI is Open-Source Software and Libraries. The growth of often free open-source AI and Machine Learning frameworks and libraries, like TensorFlow and PyTorch, provide accessible tools for researchers and developers to create and share AI models. These resources accelerate innovation, reduce development costs, and create a collaborative environment in the AI and Machine Learning community.

The amazing thing about open source software and libraries is that a company can leverage open-source AI and Machine Learning

libraries to build custom AI solutions for their specific needs, such as a fraud detection system for their financial transactions.

We can't forget about Maintenance and Support. Once a model is put in the hands of our users, we have to monitor it and maintain it. These activities drive organizational costs up and they generally include the underlying infrastructure, and sometimes, dedicated IT personnel to manage these tasks.

There are many ways to spend money when we work on Machine Learning projects that are non maintenance based, and the amazing thing is that these costs are constantly going down. In the recent developments in new algorithms and techniques, training of large scale algorithms is becoming cheaper and faster and with the rising number of researchers and engineers in the AI community, we constantly see new approaches being tried, explored and tested. All of that is helping us and businesses to drive our costs down, while improving performance.

[A] AutoML and Transfer Learning

Over the past few years, we have seen many approaches on saving time and money when it comes to Machine Learning projects. The two most applied techniques today are AutoML and Transfer Learning.

When a Machine Learning project starts, we have a goal that we'd like to achieve, and if we are lucky, we have a dataset prepared. However, in many cases, we don't have a clue what algorithm will be the 'winning' one when it comes to the goal you set. So, the standard part of any Machine Learning and AI project is to find that algorithm through trial and error.

Engineers sit down and often apply more than 1 algorithm to their data. Then they get results and try the next one until they find something that works. Sometimes when we look at data and we do a Linear Regression Analysis, for example, there could be massive outliers in the data that we can consider removing.

This process of finding the best algorithm can last anywhere from a few days to a few months!

This is where our first automation comes into play which is called AutoML.

AutoML is a process that automates the development and deployment of Machine Learning models. It involves using AI algorithms and techniques to automatically select the best model architecture which is called Automatic Model Selection. This process automatically evaluates and selects the most suitable Machine Learning algorithm for a given problem or dataset, based on performance metrics and a goal or goals that you set.

The AutoML process also optimizes what are called hyperparameters, which controls the learning process. This is called Hyperparameter Optimization. AutoML can search for the optimal combination of hyperparameters for a given Machine Learning model, improving its performance and reducing the time and effort required for manual tuning.

The AutoML approach democratizes access to AI, allowing both non-technical users and engineers to harness the power of Machine Learning without requiring extensive expertise or extensive understanding of advanced mathematics. AutoML streamlines model development, reduces time and effort, and makes it easier to create and deploy effective AI solutions across various domains.

AutoML is still far from perfect and if we had to choose between using AutoML and building something ourselves, we would go for a custom solution. The AutoML techniques are great as a starting point or for businesses that can't afford hiring internal or external ML teams.

There are several AutoML platforms that we can use such as GCP, AWS, H2O, Databricks and others. Today, most of these solutions perform equally good, the only question is what tech stack you have and how to integrate this into it.

Another technique used to help businesses in developing state of the art Machine Learning solutions is called Transfer Learning. Big companies such as Google, Microsoft, Amazon, IBM and OpenAI are training huge models for months in their datacenters and open sourcing some of them for us to use. These models are trained on

general purpose data, with a goal to learn basic constructs like language, shapes, sizes etc.

Then, we can take these models and adjust them to work on our specific problems in our companies. And instead of training those models for months, we need just a couple of hours to train it on our custom data!

For example, customer service departments can use a pre-trained language model such as ChatGPT and adjust it with their own domain-specific data, creating a custom AI chatbot that can efficiently handle customer inquiries and provide accurate information.

In summary, by using Transfer Learning, both casual users and businesses can develop effective AI solutions more efficiently and with fewer resources, making it a valuable approach for a wide range of applications.

***[A/B] AI's Role in Different Industries,
Government and Society***

Let's discuss AI's role in different industries and in government and society in general, so we have a full spectrum understanding of AI's capabilities and possible contributions in various industries.

Let's start with AI's role in the Job Market and Automation: AI automates tasks and processes, increasing efficiency and reducing labor costs. While it may displace some operational oriented repetitive jobs, it also creates new opportunities and demands for skilled professionals in AI and in related fields.

Many people are worried that AI will replace our jobs, but one way to look at it is in the 1980s when robots helped car manufacturers and production line employees were retrained to do other tasks. There is this concern that AI will replace too many of our jobs. However, we always adapt.

Historically during the industrial revolution, people were worried that they would lose their jobs and the same thing happened with robotics in many industries. The great thing about technology and AI is that we can now spend more time on R&D and less time on G&A meaning general and administrative tasks.

Next up we have Personalized Shopping Experiences: AI enables businesses to provide personalized experiences by analyzing user behavior and preferences, leading to increased customer satisfaction, loyalty, and sales.

In terms of Healthcare and Diagnostics: AI improves healthcare by aiding in diagnostics, drug discovery, personalized medicine, and remote patient monitoring, ultimately enhancing patient outcomes and reducing healthcare costs.

Just like how Intuitive Surgical has helped with robot based surgeries, AI can really help us when it comes to diagnosing illnesses when it comes to analyzing symptoms.

In terms of Education: AI can be used to create personalized learning experiences, an adaptive curriculum, and intelligent tutoring systems, revolutionizing the way we learn and teach.

Next up we have Transportation and Mobility: AI powers autonomous vehicles and smart transportation systems, improving traffic management, reducing accidents, and optimizing logistics.

In terms of Finance and Banking: AI streamlines various financial processes, including fraud detection, risk assessment, credit scoring, trading, and investment management.

Moving on to Agriculture and the Environment: AI helps optimize farming practices, monitor crop health, yields, and manage resources more efficiently, contributing to sustainable agriculture and environmental conservation.

Next up we have Smart Cities and Infrastructure: AI enables better management of urban resources, energy consumption, and public services, leading to more sustainable and efficient cities. AI can also help to reduce traffic accidents and help us to spend less time commuting, which is great for the environment and for our social lives.

In terms of Ethics and Governance: AI raises ethical concerns, such as data privacy, fairness, and transparency. It's essential to develop guidelines and regulations to ensure responsible AI development and deployment.

Moving on to Research and Innovation: AI accelerates scientific research and discovery across various fields, from climate modeling to drug discovery and material sciences.

Next we have the Entertainment and Media sector: AI is transforming the way we create, consume, and distribute content, enabling personalized recommendations, content generation, and advanced visual and audio effects.

Lastly, Customer Service and Support: AI-powered chatbots and virtual assistants streamline customer interactions, reducing response times and enhancing the overall customer experience.

These roles showcase how AI is shaping society and the economy, driving innovation, and enhancing various aspects of our lives. However, it's crucial to address the potential challenges and ethical considerations to ensure AI benefits everyone and contributes to a more equitable and sustainable future.

[A/B] Challenges and Limitations of AI

AI is very helpful and can achieve many things, however there are many challenges and limitations that AI faces. Here are 9 challenges with AI:

First of all we have Data Quality and Quantity: AI models rely heavily on data for training and making predictions. Inadequate or poor-quality data can lead to biased or inaccurate results, limiting the effectiveness of AI systems. Also one or a few data points do not make a trend.

The second challenge with AI is Bias and Fairness: AI models can inherit biases present in the training data or in the design of the algorithms, leading to unfair treatment of certain groups. Ensuring fairness in AI systems is a major challenge that requires addressing both data and algorithmic biases.

Third, Explainability and Transparency: Many AI models, especially Deep Learning models, act as "black boxes," making it difficult to understand how they arrive at their decisions. This lack of explainability and transparency can hinder trust and adoption of AI systems, particularly in sensitive domains like healthcare and finance.

The fourth challenge is Privacy and Security: AI systems often require access to large amounts of personal and sensitive data, raising concerns about privacy and data protection. Ensuring data privacy and security while leveraging AI is a crucial challenge.

The fifth challenge is Generalization and Adaptability: AI models often struggle to generalize their learning to new, unseen data or adapt to changing conditions in real-time. Developing AI systems that can robustly handle new situations and adapt to dynamic environments remains a challenge.

The 6th challenge is Computational Resources: Training large AI models can require significant computational resources, leading to high energy consumption and costs. Developing more efficient algorithms and optimizing AI hardware is essential to make AI accessible and affordable for a wide range of applications.

The 7th challenge is Regulation and Governance: AI's rapid advancement and widespread adoption necessitate appropriate regulations and governance structures to ensure ethical and responsible development and deployment. Balancing the need for innovation with ethical considerations and public safety is a complex challenge. I think that all governments and large corporations in the future will have a department of AI ethics or a C level AI ethics job function.

The eighth challenge is Human-AI Collaboration: Integrating AI systems with human workflows requires seamless collaboration between humans and AI, ensuring that AI complements human skills and expertise. Designing AI systems that can effectively collaborate with humans and augment their capabilities is a key challenge.

Lastly, number 9 we have Long-term Consequences: AI's long-term societal and economic impacts are difficult to predict. Addressing potential negative consequences, such as job displacement, inequality, and social disruption, requires careful planning and foresight. Warren Buffett is famous for being long term focused and he has stated that once we open pandora's box with AI, there is no going back.

In summary, while AI has shown tremendous potential in various applications, it also faces several challenges and limitations.

Addressing these issues is crucial for the responsible development and deployment of AI systems, ensuring that AI benefits society as a whole and contributes to a more equitable and sustainable future.

[A/B] Job Positions in the Data World

There are many positions a person can have in the data world, let's explore all of them:

First of all we have the Data Analyst: A Data Analyst is responsible for collecting, processing, and interpreting data to help businesses make informed decisions. They use tools like Excel, SQL, and Tableau to analyze data, generate reports, and identify trends and patterns.

For example, a Data Analyst at a retail company may analyze sales data to identify the best-selling products and make recommendations for inventory management.

Next up we have the Data Scientist: Data Scientists apply statistical analysis, Machine Learning, and advanced algorithms to extract insights from large datasets. They often work with programming languages like Python or R and leverage tools like TensorFlow or Hadoop.

For example, a Data Scientist at a healthcare organization might develop a predictive model to identify patients at risk of developing a specific medical condition, enabling early intervention and treatment.

The 3rd job or position in the data world is a Machine Learning Engineer: Machine Learning Engineers design, build, and deploy Machine Learning models for various applications. They work with programming languages, frameworks, and libraries specific to AI and ML, such as Python, TensorFlow, PyTorch, and Keras.

For example, A Machine Learning Engineer at a finance company could create a fraud detection system to analyze transactions and identify suspicious activity in real-time.

Number 4 we have the Data Engineer: Data Engineers build and maintain the data infrastructure required for data processing, storage, and analysis. They work with databases, data warehouses, extract transform and load processes which is also called ETL processes - ETL stands for Extract, Transform and Load. Data Engineers also work with big data technologies like Hadoop, Spark, and Apache Kafka.

For example, a Data Engineer at an e-commerce company might design a data pipeline to ingest and process user behavior data, enabling the data science team to analyze customer preferences and personalize recommendations.

Moving onto number 5, which is the Business Intelligence or BI Analyst: BI Analysts use data visualization tools, such as Tableau or Power BI, to create reports and dashboards that help businesses track their performance and make data-driven decisions.

For Example: A BI Analyst at a marketing agency may develop a dashboard that visualizes the effectiveness of various marketing campaigns, allowing the team to allocate resources more effectively.

Number 6 we have the Database Administrator also called the DBA: DBAs are responsible for managing and maintaining databases, ensuring data integrity, security, and performance. They work with database systems like MySQL, Oracle, or Microsoft SQL Server.

For example, A DBA at an insurance company might optimize database performance to ensure that customer data can be accessed quickly and efficiently by the company's software applications.

Lastly, number 7 we have the Data Architect: Data Architects design and implement data models, structures, and data integration solutions that enable efficient data storage, analysis, and reporting. They work with data warehouses, data lakes, and various database technologies.

For Example: A data Architect at a logistics firm could create a data model that captures information about shipments, routes, and delivery times, enabling better route optimization and resource allocation.

In conclusion, there are many job positions in the data world, each with a unique set of skills and responsibilities. These roles often work together to collect, process, analyze, and leverage data to drive business growth and make informed decisions. By understanding the different roles in the data landscape, you can better identify the right career path for your interests and skillset or better understand who to hire to help you take your company to the next level.

SECTION 3: CHATGPT

***[A/B/C] Introduction to ChatGPT (What, Why
and How of Section 3)***

Welcome to Section 3 where we discuss ChatGPT and Prompt engineering!

In terms of the what, why and how of section 3, let's kick it off with the what. What are we going to learn together? We will learn the most important prompt engineering techniques so that we can learn how to get the most value out of large language models, such as ChatGPT.

Prompt Engineering is an Artificial Intelligence term used when discussing Natural Language Processing (i.e., asking an AI product a question). Here in section 3 we discuss ChatGPT and then in section 4 we discuss ChatGPT plugins and then in section 5 we discuss alternatives to ChatGPT including Google's Bard AI service and Microsoft's Bing Chat AI service.

Moving on to why, why is it important to start the prompt engineering section with ChatGPT? Because ChatGPT is used in Microsoft's Bing Chat service and the T in ChatGPT was actually developed by Google. GPT stands for Generative Pre-Training Transformer.

Lastly How? How are we going to learn about ChatGPT here in section 3?

Each chapter will be very practical. We will jump to the ChatGPT interface and prompt it together and talk about the results we observed from our prompts.

Please use the following link for more than 1000 ChatGPT prompts that we created for you that you can download and use:
https://drive.google.com/drive/folders/15KxXiyRTIYf2SU63_uUBs_xgS0U6bVesI

I hope you are as excited as we are to go over these methods and learn together how to improve your business with new AI models and text!

***[A/B/C] Introduction to Prompt Engineering
(Basic and Advanced Prompts)***

When you think of an AI model, imagine that it is a high-performance car. Just like how a car needs a driver to guide it, an AI model needs prompts - meaning the right instructions that enable it to produce the output you want.

But what makes Prompt Engineering so vital? Simple - it's about quality and efficiency.

Let's take a look at the difference of quality between a basic prompt and an advanced prompt with a few techniques covered in this section.

First if you were to type the following prompt in ChatGPT, your response would be vague and not that helpful [EXAMPLE 1]:

“Create a marketing plan for an e-commerce platform selling shoes.”

Now if you typed the following more advanced and less vague prompt, then the results from ChatGPT would be incredibly helpful [EXAMPLE 2]:

“You are a senior marketing manager that helped many e-commerce platforms launch new categories and products. You are going to create a full marketing plan for my e-commerce business that sells only shoes. Our shoes are for summer only since our biggest markets are California and Florida. Create a full marketing plan that covers offline and online marketing. Output can be in markdown and make sure to ask questions if you need more information from my end. Got it?”

The difference between a poorly crafted prompt in EXAMPLE 1 and a well-engineered one in EXAMPLE 2 is unbelievable. This can mean

the difference between a chatbot that confuses customers and one that has much more immersive and useful results.

From customer service to content creation to predictive analytics, a well-crafted prompt drives the AI to deliver accurate and contextually fitting responses. This, in turn, boosts efficiency, improves customer satisfaction, and can even provide a competitive edge.

So, how can you get good at prompt engineering? It's part knowledge of your AI's capabilities, part understanding of your task's requirements, and a lot of experimenting and refining. And to get you up to speed, in this section, we will cover all the prompt engineering techniques that you will need to help your business and your career.

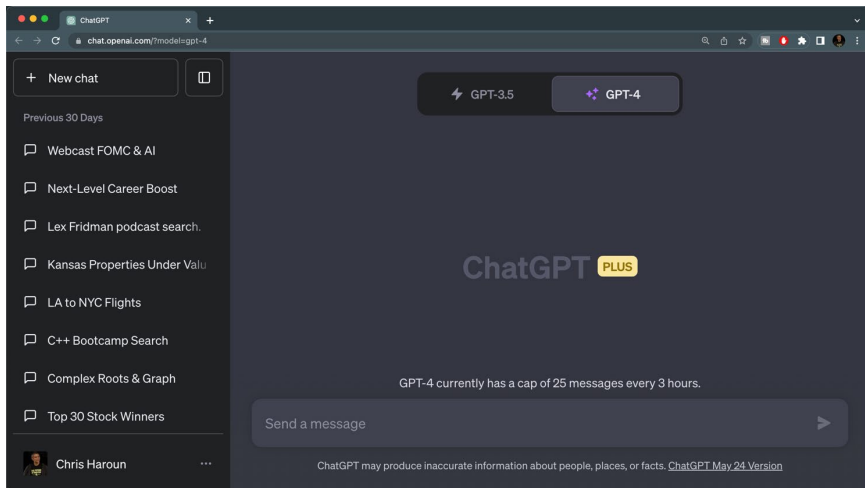
[A/B/C] ChatGPT User Interface and Account Setup

Before we can actually use ChatGPT, we need to create an account and understand the interface. If you already have an account setup, then please skip this chapter. Thanks.

First let's go to this website here (<https://chat.openai.com/>). Then click the signup button.

You'll now see a sign-up form. Fill it with your email, create a password, agree to the terms, and click on 'Create Account'. It will ask you to provide your phone number to check that you are not a spamming account.

Below is the ChatGPT interface. It's pretty intuitive. The main window on the right side is where your conversation with ChatGPT takes place.

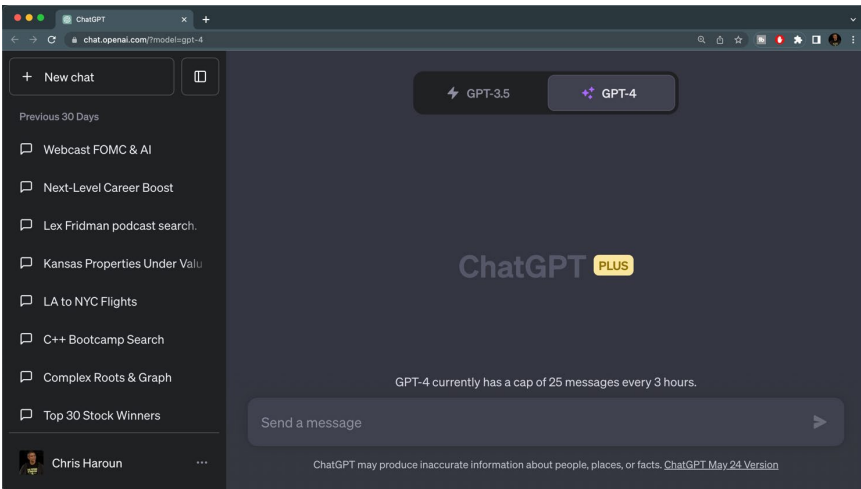


You can start by typing a message in the “Send a message” box, per the image above. Then click the send paper airplane looking graphic on the right side of the Send a message box. Then you'll see the AI's response appear.

You'll notice that ChatGPT can handle a wide array of prompts, from answering questions to writing content.

On the left hand side of your screen, you have the history of chatting sessions you had with ChatGPT. If you are interested in continuing some older discussion, just click on it.

Per the image below, if you click on your name in the bottom left hand corner you can access the Settings. In Settings you can change general settings like Dark/Light theme, whether or not ChatGPT will save or delete your previous chats and other beta features that you might see. Please note these beta features are only available to ChatGPT Plus accounts (which we will explain soon in more detail).



Speaking of plans and subscriptions. You have two options - the free version (currently GPT-3.5) and the plus version (currently GPT-4) that costs \$20 per month.

The main difference between the free version and the plus version is the access to the latest ChatGPT model and with the plus version you also get beta features, like plugins (covered in section 4) and browsing.

A great feature with the plus version is priority access. If you have a free account, there is a higher chance that you'll have higher waiting times. We will discuss many more differences between the free and plus accounts later in this section and in section 4.

***[A/B/C] Important Problems/Issues with Using
ChatGPT***

At many universities, when students submit homework using any prompt engineering product like ChatGPT, the professors won't accept the homework. There are potential copyright issues too if you use any prompt engineering products like ChatGPT, so please be careful as we don't yet know what the copyright issues will be when the copyright and intellectual property rights courts or legal system changes.

Also quite often ChatGPT is “confidently wrong” meaning because the grammar and paragraph and sentence construction seems perfect, we just assume that the content is correct. Please always fact check everything, especially numbers.

ChatGPT is great for business research and idea generation and for helping us with personal or career decisions. You can ask ChatGPT to turn any of your career or business model perceived weaknesses into a strength. Lawyers use ChatGPT for this too. In fact ChatGPT can pass the bar.

Lastly, before your next interview, ask ChatGPT to explain why your biggest perceived weaknesses are strengths! :)

***[A/B/C] Your First Prompt Exercise -
Directional Prompting***

Before we start exploring any techniques on how to improve prompting and our results, it's important to try algorithms yourselves. A great way to learn is by trial and error.

Let's discuss improve our prompting skills and which can lead to better results. The first technique of prompting is called Directional or Instructional prompting. These types of prompts are designed to give clear and concise directions to large language models like ChatGPT. We will discuss this below.

Let's look at our first example and if you want to type Prompt 1a below into ChatGPT. Please note, if you type the same prompt and submitted it a few times, ChatGPT will likely change the answers a bit each time you submit a prompt:

Prompt 1a:

List five programming languages and their primary use cases.

Prompt 1a is quite vague and might lead to different interpretations by users. Let's improve it by making it more directional:

Prompt 1b:

Provide a list of five programming languages used in web development and explain the primary purposes for which they are commonly used.

With the second version of this prompt, we gave it a much narrower subset to work with and we were more descriptive on what we would like it to achieve.

Another technique to be descriptive is to utilize another prompting technique called Output Formatting. It's quite simple, we specify what is the format we would like it to generate for our outputs.

Prompt 1c:

Provide a list of five programming languages used in web development and explain the primary purposes for which they are commonly used. Generate outputs in the format: Programming Language -> Utilization

In this chapter, we've learned about Directional meaning Instructional prompting and how to create prompts that are clear and concise. By providing clear instructions, we ensure that LLMs, which stands for Large Language Models, understand what's expected of them and can provide the desired information or output.

A Large Language Model (LLM) like ChatGPT is trained using a massive amount of text in order to provide value to the user (i.e., answering a ChatGPT questions or prompt).

[A/B/C] Role Playing

Let's continue our adventure on making our prompts more efficient and robust. Now we know that prompts should contain more context and information to get better results. However, on its own it's not as powerful as combining it with the technique discussed in this chapter - meaning the concept of Roles.

When prompting an LLM like ChatGPT, at the beginning of the prompt, we can add a part we call `system` prompt or a Role prompt. Role prompt looks like this - You are a professional soccer player, Act as a senior lawyer partner, etc.

Let's enter this prompt now [do and wing analysis]. Role playing is a powerful way to differentiate your skillset versus the competition when meeting with a potential customer or employer.

Basically we are telling LLMs how we want it to act so that it can produce more relevant results, as a person in that role would do.

Let's use a simple prompt and then role-prompt to compare the results.

Prompt 1a:

Create presentation slides on the 3 most innovative transportation solutions for smart cities.

Prompt 1b:

You are an urban planner that helped with creating many smart cities. Create presentation slides on the 3 most innovative transportation solutions that can be implemented in smart cities. No more than 3 bullet points per slide. Put a motivational quote about smart cities on the second slide.

By defining a role, we've given the AI more context and set expectations for the quality and tone of the response.

Please note that we can start a prompt with “Act as a ...” and we will get similar results as “You are a ...”.

There are many different ways to get value from role prompts and luckily, there is a huge community exploring these results now. Check out the link [<https://github.com/mustvlad/ChatGPT-System-Prompts>]

To summarize, we've learned the importance of incorporating roles into your prompts to guide the AI's response style and perspective. By defining a role, you provide context and set expectations for the quality and tone of the output.

[A/B/C] Few-Shot Prompting

Are you better at learning just by reading a theory driven book or when you see a few examples? Many people like to get their hands dirty and see a few examples, as they learn much quicker this way.

Similarly with LLMs, when we provide a couple of examples of output that we want it to produce, it will keep doing that pretty well! This technique is called Few-Shot prompting. And as a variation to the topic, we have Zero and One Shot Prompting. The only difference is in the number of examples we provide to it.

However, in our experience, the Zero Shot prompt, meaning the AI has no prior examples to guide its response and may not produce the desired results. Thus use Few Shot prompts for better and more consistent results.

Let's look at our first example:

Prompt 1:

I'll give you a few examples of knowledge extraction and its formatting, and you need to extract that information from all future prompts.

1. Apple is the most valuable company in the USA.

Apple -> [COMPANY]

USA -> [COUNTRY]

2. Mario is the CEO of GE.

Mario -> [PERSON]

CEO -> [POSITION/TITLE]

GE -> [COMPANY]

3. *One of the largest countries in Europe is Germany.*

Germany -> [COUNTRY]

By providing multiple examples, we're helping the AI understand the pattern we're looking for in its responses. This can be particularly useful for extracting specific information or formatting the output in a certain way.

Let's examine another example where Few Shot prompting produces great results! Create a new chat and paste this prompt:

Classify the following movies based on their genres:

Titanic

A: Romance/Drama

The Dark Knight

A: Action/Crime

Inside Out

A: Animation/Adventure

The Shawshank Redemption

A: Drama/Crime

[Movie Title]

A:

Now we have all the building blocks needed to create really powerful prompts for basically any use case we need! Let's take what we learned so far and put everything under one prompt in the chapter video and help us write promo emails!

[A/C] Chain of Thought Prompting

Chain of Thought prompting involves breaking down complex tasks into smaller, connected steps that build upon each other. This technique helps to maintain coherence in AI-generated responses and ensures that the AI follows the desired thought process.

Let's dive into an example and see how Chain of Thought prompting works in action.

Example:

You are the CEO of a tech startup, and you want to create a comprehensive plan for launching a new product. Instead of asking the AI to generate the entire plan in a single prompt, you can break it down into smaller prompts that form a logical sequence.

Prompt 1:

As the CEO of a tech startup, outline the key steps involved in launching a new product.

Prompt 2:

Based on the outlined steps, describe the target market and ideal customer profile for our new product.

Prompt 3:

Taking into account the target market and ideal customer profile, create a marketing strategy to promote our new product effectively.

Prompt 4:

Considering the marketing strategy, determine the necessary resources and budget allocations to execute the plan.

By breaking down the task into smaller prompts, you ensure that the AI-generated responses are coherent and maintain a logical flow of thought.

Try pasting the aforementioned 4 chains of thought prompts into ChatGPT to see the result.

Now, let's practice creating a Chain of Thought prompt for a different scenario:

Scenario: You are a content marketing manager looking to write a compelling blog post about the benefits of adopting electric vehicles (EVs).

Prompt 1:

As a content marketing manager, provide an engaging introduction to a blog post about the benefits of adopting electric vehicles (EVs).

Prompt 2:

Based on the introduction, list the top 5 benefits of electric vehicles for consumers and the environment.

Prompt 3:

For each of the listed benefits, provide a detailed explanation, including facts and statistics to support the claims.

Prompt 4:

Write a convincing conclusion that encourages readers to consider electric vehicles as a viable alternative to traditional combustion engine vehicles.

In this chapter, we've introduced the Chain of Thought prompting technique and demonstrated its effectiveness in guiding AI through a sequence of connected ideas. By breaking down complex tasks into smaller, logical steps, you can improve the coherence and relevance of AI-generated responses.

[A/C] In Style of Prompting

Let's discuss how to write text and get answers in the style of another person. You might have seen online poems, speeches, lyrics and emails written in the style of Joe Rogan, Barack Obama, Eminem, Yoda or Luke Skywalker.

In most cases prompting in the specific style sounds is not that useful, but we promise you, there are ways to use it in our day-to-day work; let's see how!

The first thing that we think when we mention this type of prompting, is writing as another author, but that is just a part of it. Style can be an emotional state as well! For example, re-write this email in an enthusiastic tone. Or write something in a sad tone. Some resources would call this emotional or tone prompting, but since we are talking about the style of the output, we can put those under the same umbrella.

Let's start with an example to see how this technique works:

Let's check out if we can get an Eminem style of this promo email:

Example 1:

Rewrite the following email in the style of Eminem and make it rhyme

Hey everyone,

Hope you're enjoying your week. Good news - spring is almost here, and I bet you're excited for some sunshine and warmer weather 🌞. In order to celebrate, I'm lowering the prices on all of my courses as much as Udemy allows me (\$10-\$15) for the next 5 days.

Click on any of the books below to check them out:

Please try the aforementioned prompt and have fun with this!

Here is another variation of this technique:

Re-write any email that I paste in here, in my style of writing. As a reference I will provide 2-3 LinkedIn posts that I wrote [paste your posts here]

Imagine creating targeted ad campaigns, and generating engaging social media posts, and persuasive web copy, all with the help of AI!

[A/C] Structuring Your Data & Using Tables

Let's use ChatGPT to extract and structure information from your data. Extracting information in this way is useful when you, for example, want to go through some text quickly and make a presentation out of textual data.

Let's take a simple example first with the following prompt to paste in ChatGPT:

Extract information from the example and structure it into markdown table:

Marko is working in the IT industry as a QA. He was living in Belgrade, capital of Serbia, but after Covid, he moved to Thailand to work remotely.

Now, this example was okay, but not perfect. Here is a description of real graph taken from one of Luka's papers. If you want, please copy the italicized text below into ChatGPT:

The graph illustrates the Inflation Rate in Luxembourg from 2013 to 2023, showcasing the annual percentage change in the general price level of goods and services. The graph reveals fluctuations in the inflation rate over the years, with some distinct trends and turning points.

According to information found on Trading Economics and STATEC, in 2013, the inflation rate in Luxembourg stood at around 2.3%. From that point, the rate experienced a downward trend,

eventually dipping below 0% in 2015, indicating a period of deflation. Following this period of negative inflation, the rate began to rise again, fluctuating between 0% and 2% from 2015 until late 2021. This period of relatively low and stable inflation reflects moderate price growth in the country.

In 2021, the graph showed a sharp increase in the inflation rate, soaring from 0% to close to 8% by mid-2022. This spike indicates a rapid acceleration in price growth, which can have various implications for the economy and consumers' purchasing power.

As of the latest data point, the inflation rate in Luxembourg is around 4.3%, reflecting a slight decrease from the peak of 8% in mid-2022. However, the rate remains notably higher than the relatively stable period between 2015 and 2021.

Extract information from this and store it in the table

Amazing right? Now we don't need to read all this just to extract all that information manually!

[A/B/C] Text Summarization

One more cool thing that we can achieve with ChatGPT and other LLMs is content summarization. We can provide a big chunk of text to it and just say - summarize this. The output is going to be much shorter text containing only the main points from the original, longer, text.

Using what we learned before with roles, we can craft a summarization prompt like this:

You are your summarization bot. Any text that I provide to you summarize it, and create a title from it. The output format should be: SUMMARY: Text TITLE: Created title

And then provide text that you'll like to be summarized.

Let's take a part of Google's blog on Future of AI. Please go to the following link and paste it in ChatGPT to see the summarization machine.

<https://cloud.google.com/blog/transform/ai-2030-andrew-moore-five-areas-to-invest-in-ai>

And there we have it! A bot that works for us to summarize any text.

Before we finish this chapter, we'd like to point you to one more way to prompt summarization to ChatGPT. Here is how to prompt it in the same prompt with the given text:

Summarize the given text. Summaries should be short and to the point, capturing the main parts of the text.

[then paste your text into ChatGPT.

To summarize - pun intended - we covered a way to ask ChatGPT to capture only the most relevant information from the given text and provide a condensed, shorter version of it.

[A/C] Text Classification

Similar to text summarization, we can use LLMs to perform text classification. When you start thinking about this, the first question is probably - where would we use such a task?

Well to your surprise, text classification is used quite often! In fact, you are utilizing that feature daily without knowing! Spam detector in your email is an example of that.

However, many other companies are utilizing text classification for their day-to-day activities. For example, marketing companies are utilizing this feature to perform sentiment analysis for brands. And with that they can find how the general public perceives that specific brand.

Let's build a prompt that performs sentiment analysis and classifies any given text as Negative, Positive and Neutral by pasting the following italicized text into ChatGPT:

You are your sentiment analysis bot. Classify any text that I provide into three classes Negative, Positive and Neutral. The output format should be: [SENTIMENT]: Text

Now, we can provide any text, such as tweet and ChatGPT will classify it into those 3 classes. Let's test it out by pasting the following italicized text into Chat GPT:

NVIDIA is a decent company to work for, but they definitely make the best GPUs on the market right now! (A bit expensive though)

And that's it! We have our classification bot working for us! Of course, you can play with this prompt and reorganize it so that you can provide text directly within the prompt, as we did for summarization.

[A/C] Coding Generation

In this chapter, we'll briefly explore how ChatGPT can be used for code generation. Although we are not going to go deep into this area, it is important to explore this capability of the model as well.

Let's consider an example of generating a simple backend using Python with Flask. Keep in mind that the AI-generated code may need to be modified and tested before deploying it in a real-world project.

To create a code generation prompt, you can start with a directive prompt like this:

Generate a Flask backend that has 3 endpoints - home page, settings and settings-update (POST only) pages.

When you receive the response to the aforementioned prompt, it looks pretty decent for the boilerplate code, we can definitely start building on top of this. However, it's important to remember that AI-generated code may not be perfect. Here are some key points to consider when using ChatGPT for code generation:

Review and modify the code: Always review the generated code to ensure it meets your specific requirements and modify it as necessary.

Test the code: Conduct thorough testing to verify the code's functionality, performance, and security, you can't be secure enough when it comes to your company's IP, before deploying it in your project.

Use as a starting point: AI-generated code can be an Excellent starting point or a source of inspiration for developers working on real-world projects.

While ChatGPT can be a powerful tool for generating code, it's essential for business professionals to understand its limitations and the need for human oversight.

Now we know that ChatGPT can be used to generate parts of your code, but it is definitely not there to replace a full-time developer

[A/C] Additional ChatGPT Tips

Now that you know how to write amazing, powerful prompts, we want to give you a few extra tips/tricks that will help you in your prompting journey.

The first one is very simple - continue. If ChatGPT just stops responding, write continue and it will continue where it left of the last time.

The next tip is two nice ways on defining the output format you can get from ChatGPT:

1: Generate response in Markdown

2: Generate response in table format

Just add at the end of your prompt that you want output in markdown or table format, and you'll get it.

And that's it! Happy prompting!

SECTION 4: CHATGPT PLUGINS

[A/C] Introduction to ChatGPT Plugins

Welcome to section 4! So far we have discussed using ChatGPT as a traditional Chatbot; we ask it for some information and sometimes it returns sensible results, sometimes not so much.

Here in section 4, we take it to the next level. In terms of the What and Why of Section 4, let's kick it off with the what. What are we going to learn? Well recently OpenAI - the creator of ChatGPT, introduced another way to utilize their product and get additional value by pairing it with functionality from other applications like Zillow, WolframAlpha and others.

They introduced Plugins! Plugins allow us to interact with the external world through language. For example, general ChatGPT has an option to return only results that it was trained on, if we ask something it doesn't have info on, it will provide a generic message saying it is not capable of answering that yet.

Moving onto why. Why are we going to discuss plug ins? Because plugins allow ChatGPT to expand its reach and suddenly, it is not limited to its knowledge from the training dataset, but external data as well! Like the Zillow dataset or knowledge of mathematics from WolframAlpha, or even all flight info from KAYAK!

OpenAI allowing companies to connect their product directly to ChatGPT is like the moment that Apple introduced the AppStore with iPhone. More and more companies will jump on the opportunity to create a plugin for ChatGPT and suddenly, many more people will be able to interact with it!

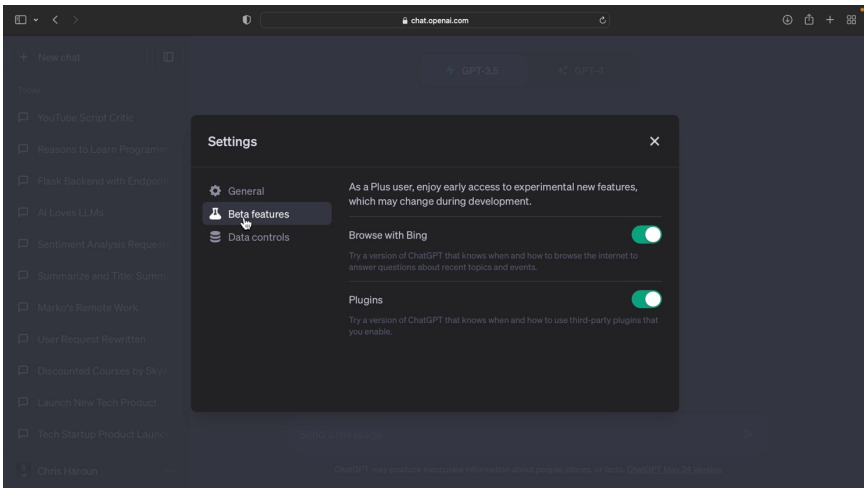
***[A/C] What are ChatGPT Plugins and How to
Get Them***

In the previous section, we talked about a lot about ChatGPT, and we talked about prompt engineering. If you just rely on ChatGPT, it's limited to the knowledge that it was trained on.

But what they did is really smart. They allowed other companies to connect to ChatGPT and their products through something we call plugins. Plugins is a way to not only use ChatGPT knowledge, but also the knowledge of plugin tools such as Zillow, Kayak, Wolfram, Alpha, and others using the ChatGPT text interface.

For now, plugins are only applicable to ChatGPT Plus as a kind of a beta feature. So, it's not openly accessible to everybody, only for paid members.

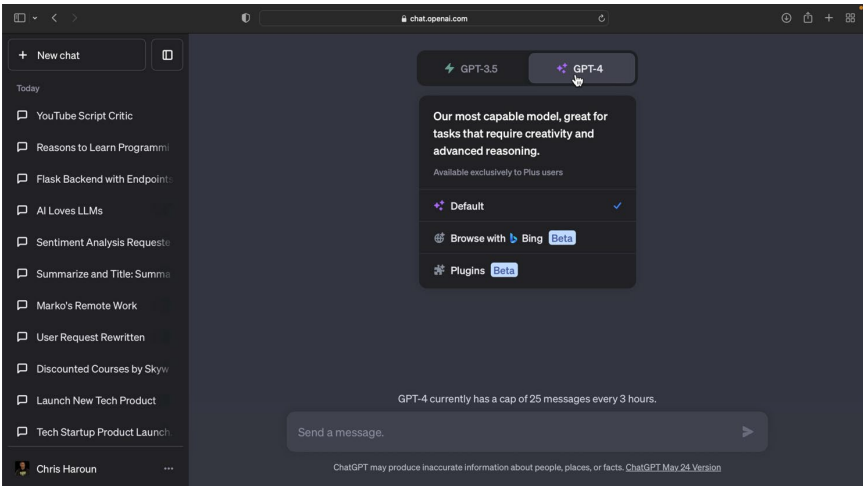
If you are a subscriber to ChatGPT4 (meaning the plus version), you can access the plugins when you go to settings, per the image below:



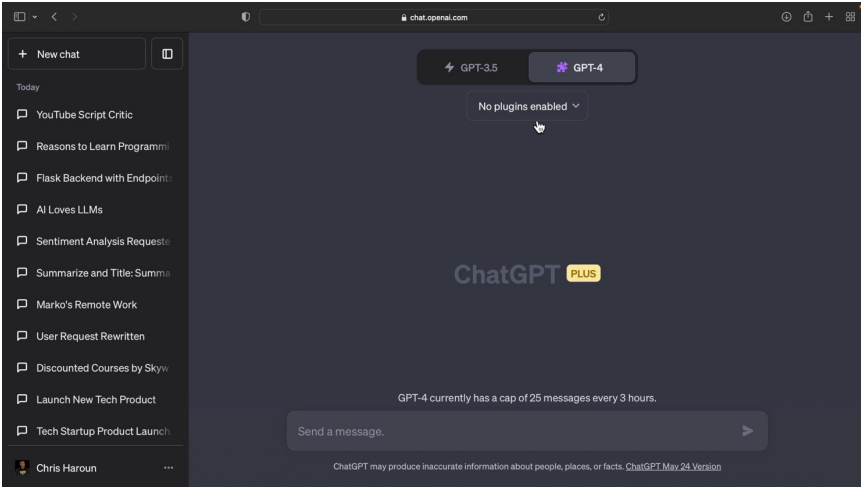
You will notice in the aforementioned image, that there are 2 beta features to turn on. One is browse with Bing, and another one is

plugins. Browse with Bing allows you to actually check the data from the outside of ChatGPT using Bing Search.

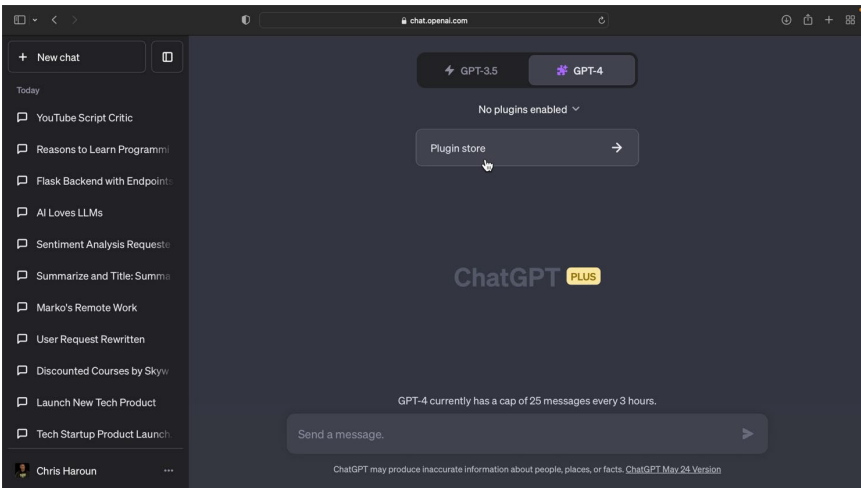
The next step to use plugins (or browse with Bing) per the image below, is to select GPT-4 and then select plugins (and/or browse with Bing). We will come back to browse with Bing later. Please just select plugins per the image below.



Next step is if you don't have any plugins installed, then please click on no plugins enabled per the image below:



Then, per the image below, click on the plugin store:

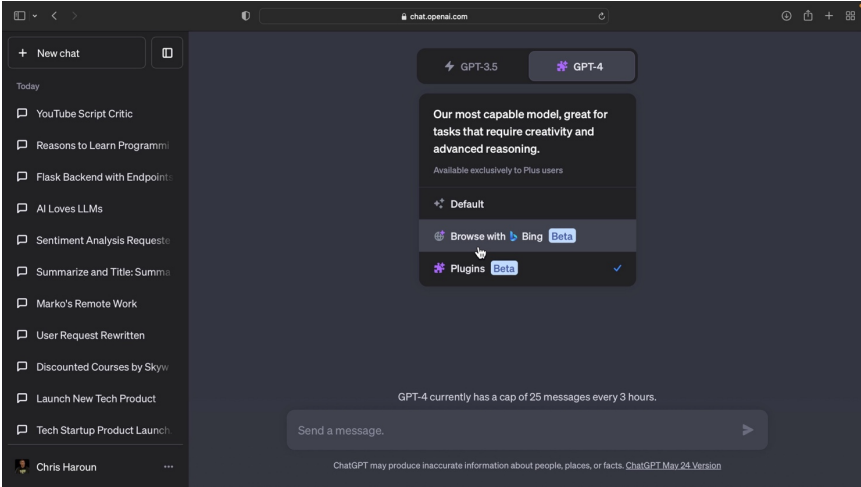


You can now search for and install plugins. For example, Zillow allows you to interact with their housing market. Expedia's plugin will allow you to find trips, plan trips, and so on. And that's the power of plugins, allowing you to, with simple to use the knowledge of these tools and get their data into your prompting session.

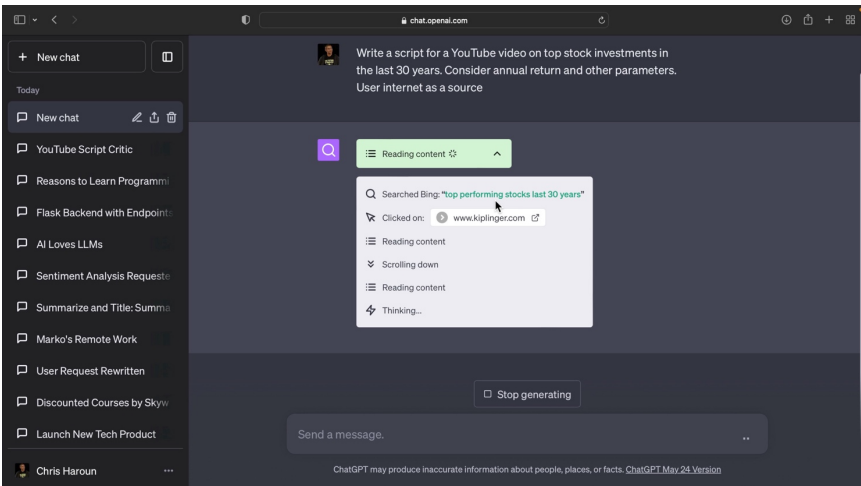
You can install as many plugins as you like. In the next few chapters, we will explore how to use different plugins.

[A/C] Browsing with ChatGPT

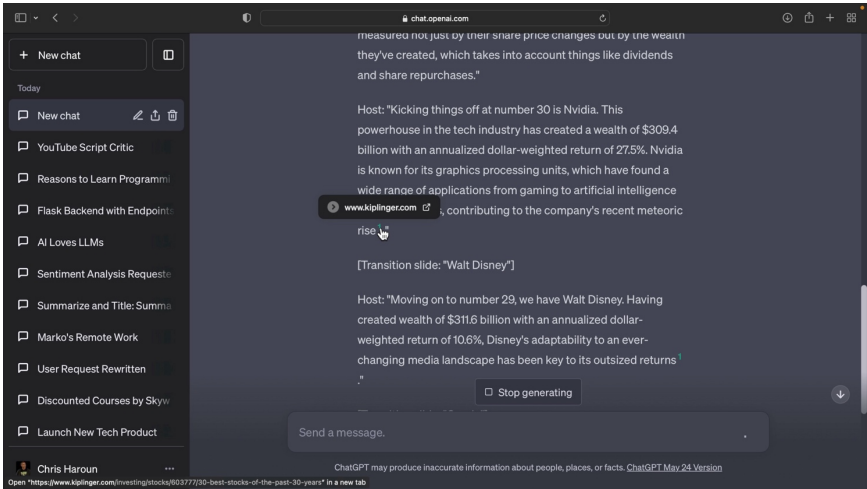
In this chapter, we will explore the first plugin developed by OpenAI themselves and that is for browsing with Bing. Per the image below, using GPT-4 select browse with Bing.



Per the image below, when you use this plugin, it actually goes and tries to find information from outside the ChatGPT.



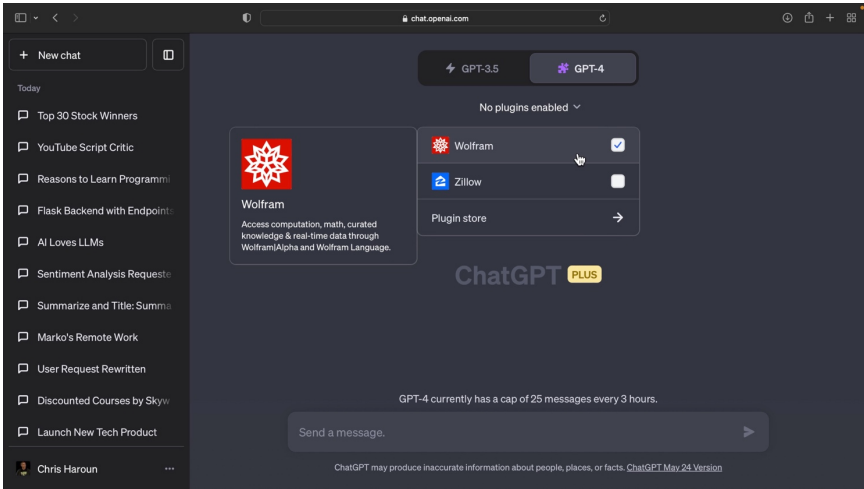
Then, per the image below, when you receive a response, you can mouse over to see the source of the response.



[A/C] Wolfram Plugin for ChatGPT

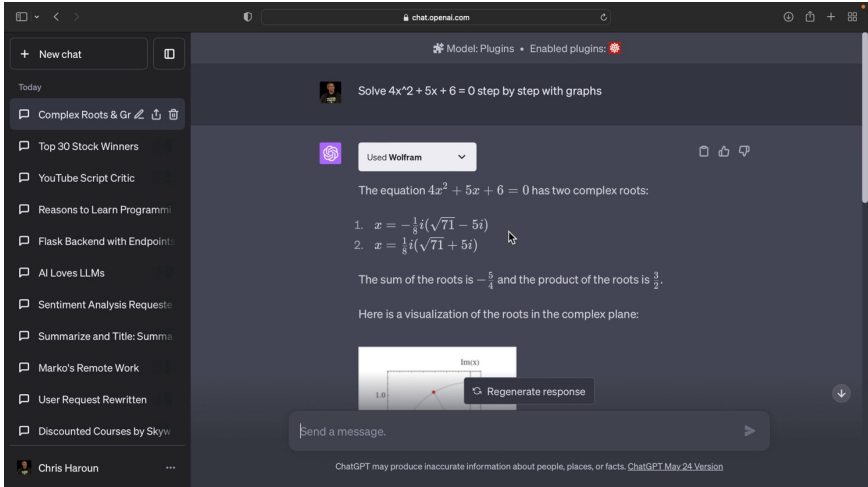
As previously mentioned, go to the plugin store, but this time please search for and the Wolfram Alpha Plugin, which is an incredible mathematics tool.

Per the image below, once you have installed the plugin, please make sure to select it.



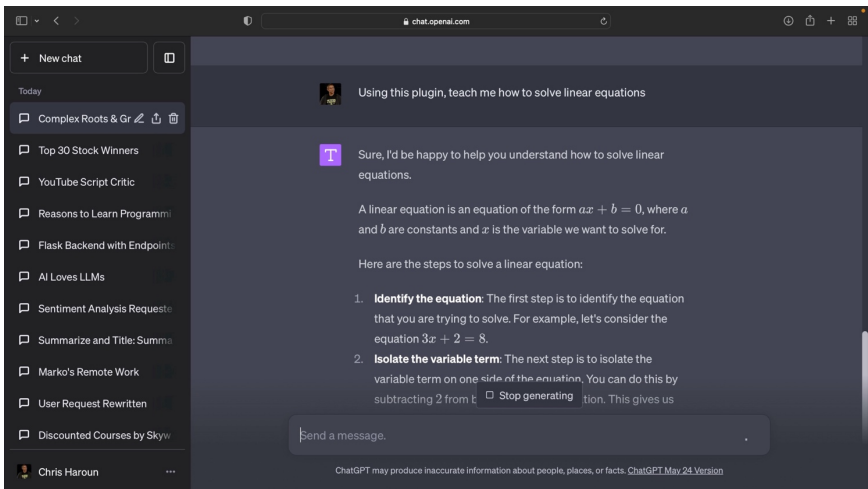
You can select up to three different plugins for one chat session.

Per the image below, you can ask the plugin to do virtually any math problem, as shown in the image below.



Now we can get graphs inside ChatGPT!

Per the image below, you can even ask the plugin to teach you math, like linear equations.



This is a bit of a niche plugin, but an amazing one that can help you as a parent, for example, or for anybody that loves or wants to learn about mathematics.

[A/C] edX Plugin for ChatGPT

If you're not familiar with edX, it's an amazing website for free education. It allows for access to education from different institutions like MIT, Stanford, and others to help you learn specific topics like data science, mathematics, business, etc.

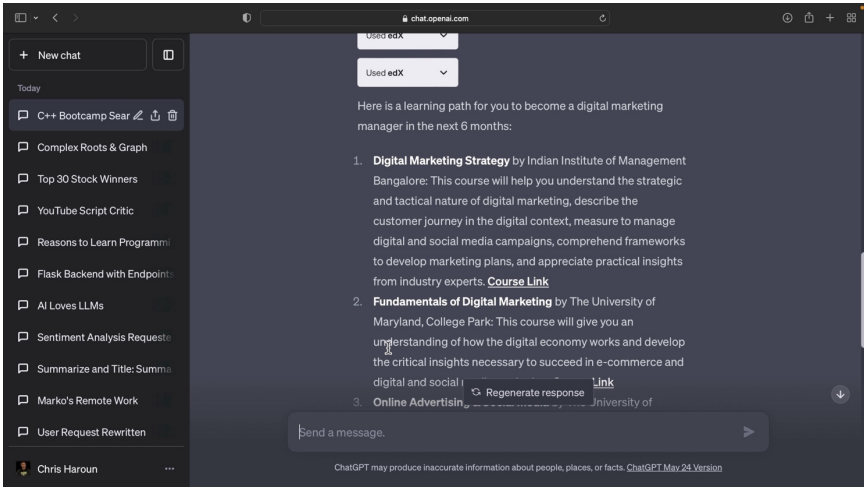
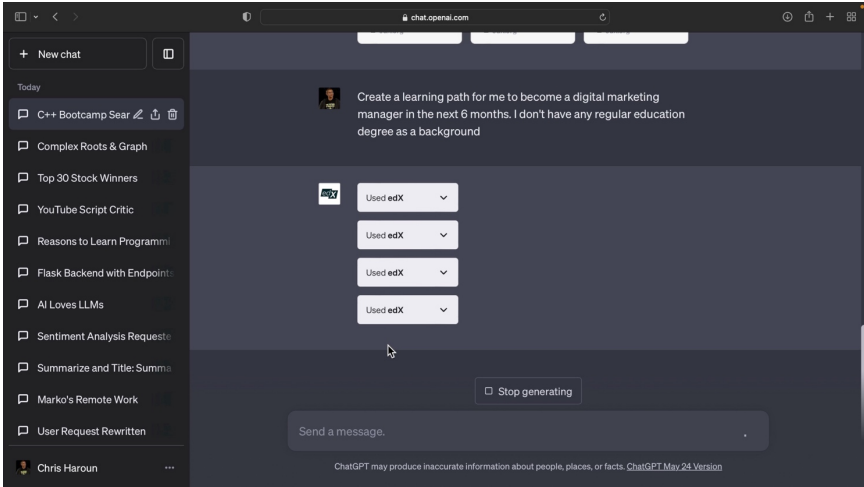
As mentioned earlier, search the plugins store in ChatGPT and install the edX plugin.

The edX plugin allows you to search bootcamps, courses, lectures, content materials, and everything present on their website. So, it's basically search plugin.

Search aside, one thing that we found really useful using this plugin for is creating your learning path. So, let's create a learning path prompt and see what we can get with that. So here is the prompt that we create:

Create a learning path for me to become digital marketing manager in the next six months. I don't have any regular education degree as a background.

The prompt and results are shown in the 2 images below.



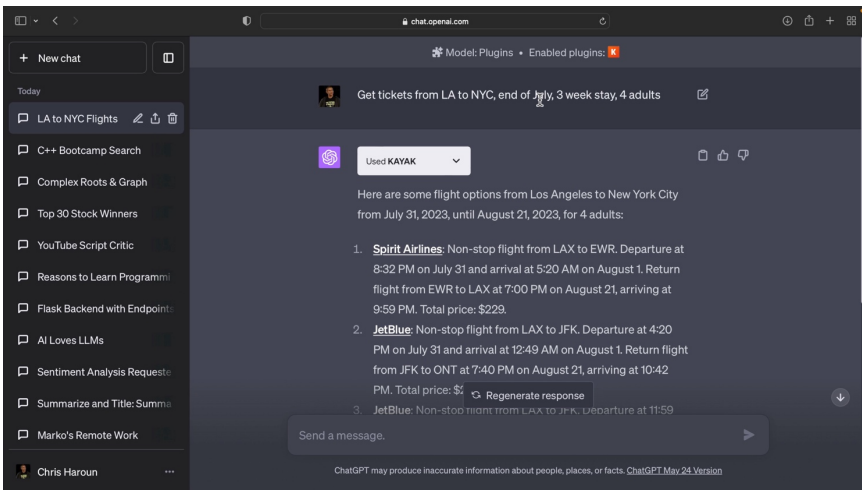
As you can see it built a pretty comprehensive learning path, starting with the digital marketing strategy, fundamentals of digital marketing, online advertising and social media, etc, which is amazing!

[A/C] KAYAK Plugin for ChatGPT

The third plugin that we are going to cover in this section is Kayak, which is a website for searching for flights, hotels and more. If you install the Kayak plug in, make sure to select it and unselect other plugins.

Per the image below, we wrote a prompt saying:

Get tickets from LA to NYC, end of July, 3 week stay, 4 adults



The results were impressive, however, the prices were shown for just 1 adult and not 4. As with all data results via ChatGPT, we have to double check to see if it is correct.

[A/C] Zillow Plugin for ChatGPT

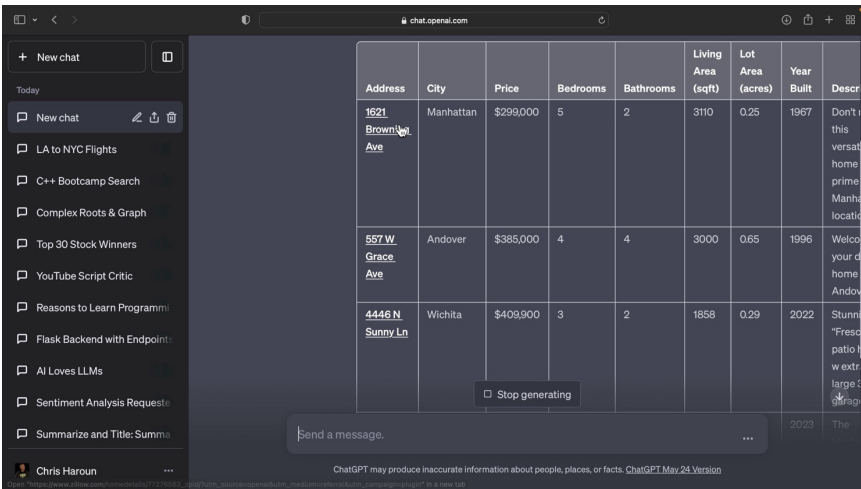
One plugin that we really find useful is Zillow, which is a website for real estate, where people can put their houses on sale (usually through real estate firms).

Zillow was one of the first companies that built a ChatGPT plugin. And what we can do with this is search properties that are currently active on the market or have recently sold using ChatGPT.

After installing the plugin, you can ask very detailed real estate questions, per the following example that we tried:

Find 10 properties in Kansas under market value, with three bedrooms, yard, close to school, and structure it in the markdown table. Use Zillow plugin for this information.

And here is the result we received from this query:



Per the image above, there are even hyperlinks to the properties in the first column.

SECTION 5: ALTERNATIVES TO CHATGPT (BARD & BING)

[A/C] Introduction to ChatGPT Alternatives

Welcome to Section 5 where we discuss alternatives to ChatGPT, including Google and Microsoft's prompt engineering products.

The GPT in ChatGPT stands for Generative Pre-training Transformer. In fact, Google invented the T in GPT, meaning Transformer. So, the question is why didn't Google release their own prompt engineering product years ago?

The answer is the classic innovator's dilemma because these AI platforms can cost literally \$100 billion dollars to create and they are expensive to maintain, which is why we pay to use advanced ChatGPT features.

Some Wall Street analysts estimate that if Google released their own ChatGPT-like product years ago, then it would hurt Google's Operating Profits by about \$30 billion per year. Finally, Google has competition and we know that competition is a good thing for consumers.

Let's talk about the What and Why of Section 5. What are we going to learn? We will learn how to use ChatGPT alternatives including Bing Chat and Google Bard.

In terms of Why. Why should we learn about alternatives to ChatGPT? Because there are certain features that ChatGPT alternatives have that you might want to use in order to take your career or your company to the next level.

Let's get started demystifying alternatives to OpenAI's ChatGPT!

[A/C] Bard by Google

When OpenAI announced ChatGPT and its partnership with Microsoft, everyone started looking in Google's direction and waiting patiently for their response to the matter. Google was worried about their dominance in search as the Bing app actually got more downloads than the Google app on iOS.

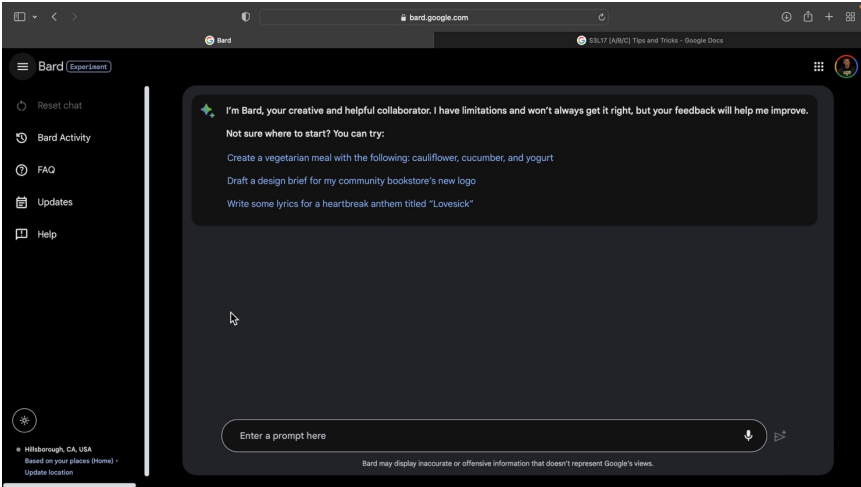
And a few months after the initial release of ChatGPT in late 2022, what happened was in early 2023, Google released their first version of generative AI model for text called - Bard.

Bard can do basically everything that ChatGPT can, but not at the same level. Some things are better, and there are many things that need to be improved.

Using Bard makes sense if you prefer Google's data policies or if you are already using Google's tech stack, so the integration to your products is much easier to accomplish if you are invested in Google technologies.

If you want to test Bard, go to <https://bard.google.com/> and sign up for testing with your personal Gmail account. Please note that Bard is not yet available in all countries.

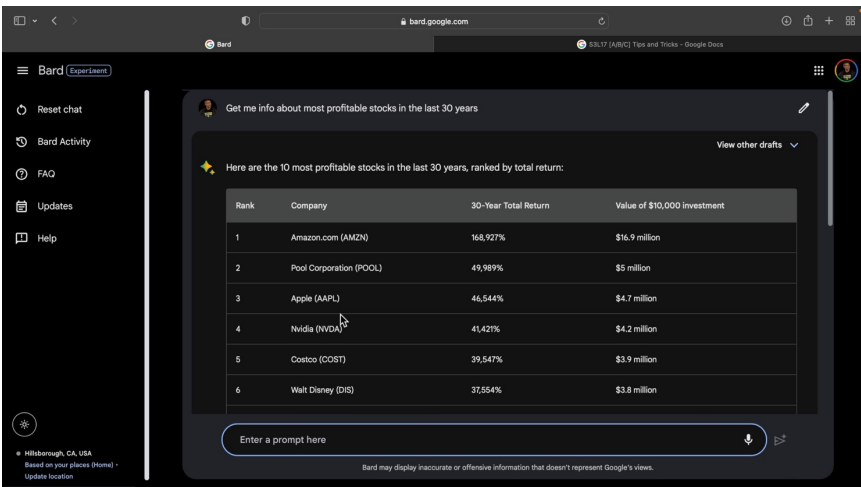
Once approved, from the link, you will be able to access the chat box and start utilizing this product. After logging in, you will be at the core dashboard that looks very similar to ChatGPT, per the image below. There is a big text box in the middle and on the left hand side of the screen we see settings and history.



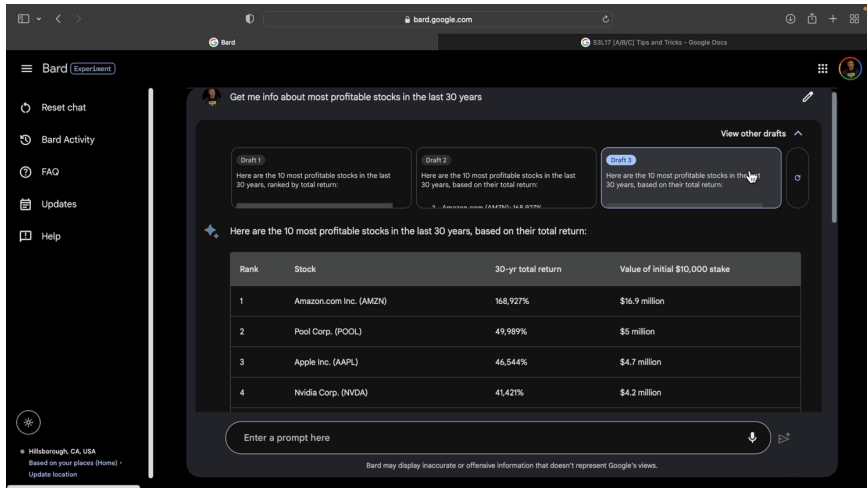
The prompting part works the same as with ChatGPT, so if you already went through the chapters regarding prompt engineering techniques in section 3, then you are ready to use Bard.

The interface is slick, but per the query and image of the query below, we do not see resources for the data.

Get me info about most profitable stocks in the last 30 years



What we do like about Bard is, per the image below, we can see there is the draft part; with no additional waiting, we can get different versions of the text and get one that suits our needs best.



We encourage you to go and test it out in more detail, try comparing results with the same prompt between ChatGPT and Bard, so that you can decide which one works the best for your line of work.

In our opinion, more complex prompting, formatting, summarization and that extra layer of complexity, is better handled in ChatGPT. Bard still has a ways to go.

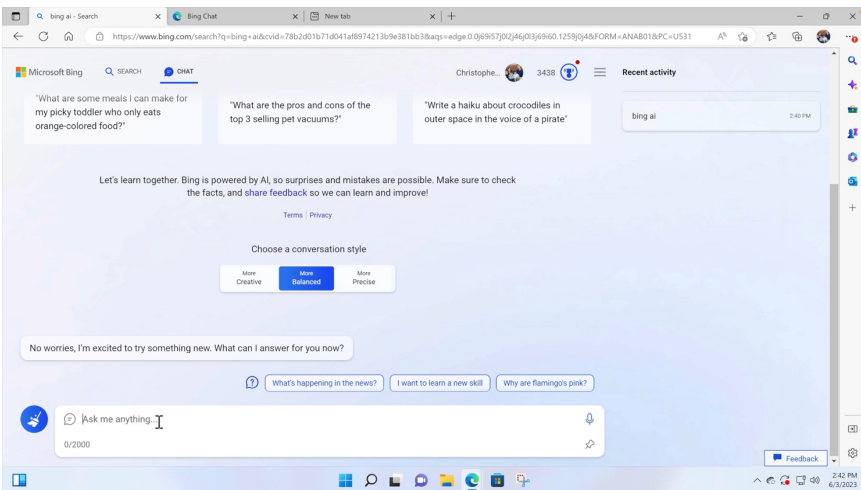
[A/C] Bing Chat by Microsoft

So far we talked about ChatGPT from OpenAI, and we discussed Bard by Google. Let's chat about Microsoft's Bing product when it comes to their competing prompt engineering product.

Microsoft invested billions of dollars in OpenAI to support their research process and to help them with the costly infrastructure buildout. In exchange, Microsoft became a forefront of AI innovation and rushed into upgrading their current offering with competitive AI solutions.

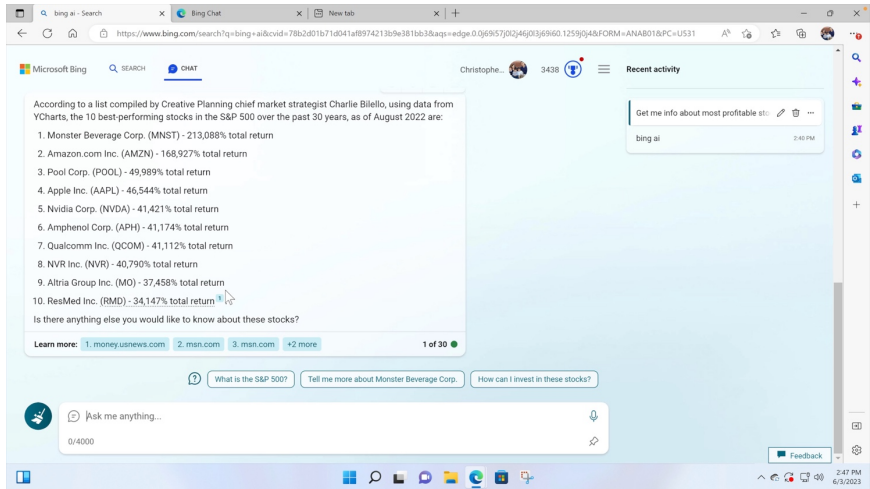
The first and the most hyped solution was the integration of ChatGPT into Bing search, called Bing Chat.

The first thing to note is that Bing Chat is only available through Microsoft's Edge browser. Once there, (per the image below) search anything on Bing and below the search bar, click Chat to open the Bing Chat interface.



The first thing that we can notice here is the chat conversation styles - Creative, Balance and Precise. In most cases, these AI models will randomly select what will be generated.

If you need more research link consistency, go with Precise mode, if not, Creative is a much better option. We personally never use the **Balance** mode.



Per the image above, you'll see the first difference between traditional ChatGPT, and this model is the fact that Big uses external resources and provides them as a reference.

SECTION 6: SECTION 6: HOW TO USE EXCEL AND MACHINE LEARNING

[A/B] Introduction to Using Excel and Machine Learning

Welcome to Section 6 where we discuss How to use Microsoft Excel and Machine Learning.

We know that Artificial Intelligence (AI) refers to computer systems that perform tasks that would usually require human intelligence. Machine Learning (ML) is a subset of AI that involves developing algorithms to learn patterns from data.

Instead of programming rules, machines learn from experience, just like humans do. We will use Excel in this section to do Machine Learning; we know that within Machine Learning you can also do Deep Learning, but this is tougher to do on Excel as Deep Learning can require a lot of computation and is something we often do in the cloud using Python and other applications. So, for this section we will discuss how to do Machine Learning in Excel.

In terms of the What and Why and How of section 6, let's kick it off with the what. What are we going to learn together? We will learn how to use Excel to do Machine Learning when it comes to basic and advanced Linear Regression Analysis and Multiple Regression Analysis.

In terms of why. Why should we use Excel for Machine Learning purposes? Because it's relatively easy to do so as Excel can quickly analyze data and provide charts, regression formulas and advanced statistical analysis. Also, if we start doing Machine Learning with Excel, it's much easier to understand how to interpret or use Machine Learning in other applications.

In terms of how are we going to learn how to use Excel and Machine Learning? We will do a statistical analysis on a fictional restaurant called Los Pollos Hermanos by analyzing what helps or hurts the sales for Los Pollos Hermanos.

If you don't have Microsoft Excel, if you want you can search for free trial of office 365 and download Excel. We will also teach you how to add Microsoft's Data Analysis plug in during this section so that you can do regression analyses with us.

[A/B] Basic Linear Regression Analysis in Excel

A Regression Analysis analyzes data to help us understand patterns in the data with a line chart. This line chart helps us predict outcomes.

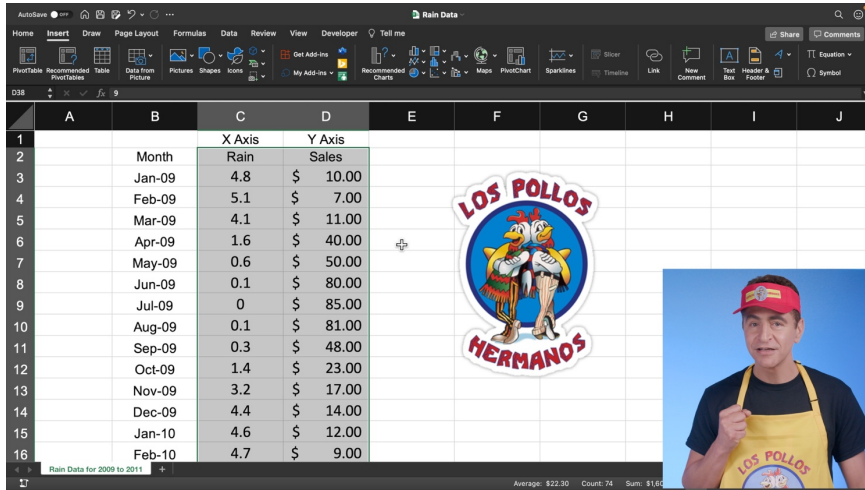
Let's talk about Regression Analysis when it comes to Machine Learning and how to do a Regression Analysis in Excel. Linear Regression is part of Supervised AI which we discussed in section 2, and it helps you find patterns in the data usually through a simple line chart, which we will show you in Excel.

It's more fun to learn this together using examples. Let's talk about how hedge funds use Linear Regression to make money. There is a hedge fund called Renaissance Technologies, founded by a brilliant mathematician who worked for the NSA named Jim Simons. Renaissance has made 60% returns on average on a gross basis since the 1980s. Amazing eh.

What Renaissance does is they use thousands of Linux computer servers running in parallel that pick stocks using Machine Learning. They don't usually employ people that thoroughly understand finance. Rather, they tend to hire Machine Learning expert mathematicians that look for patterns in the data and apparently they keep their Machine Learning strategies simple.

For example, they might analyze rain patterns and if it rains less than often, then they will buy restaurant stocks as more people go out to restaurants when it doesn't rain. Conversely, if it rains a lot, then restaurants don't do as well, and Renaissance might short restaurant stocks if it rains a lot.

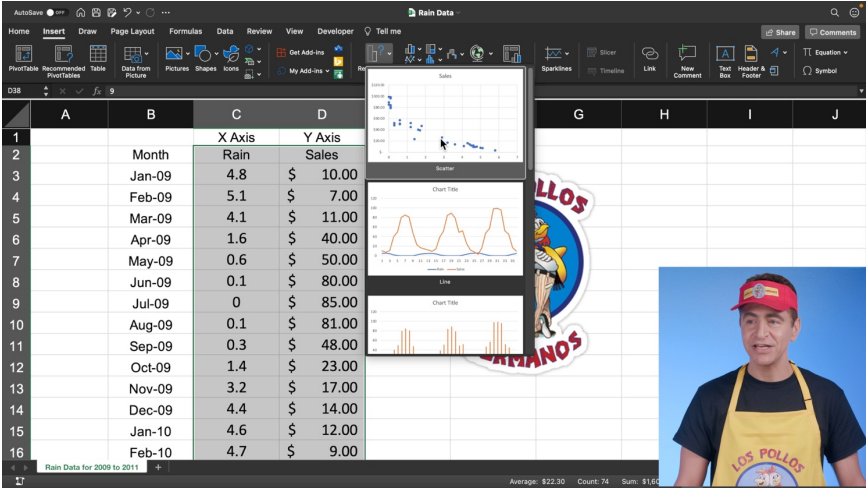
Let's go to Excel and create a Regression Analysis for a fictional restaurant company called Los Pollos Hermanos and bonus points if you tell us what TV show Los Pollos Hermanos is from!



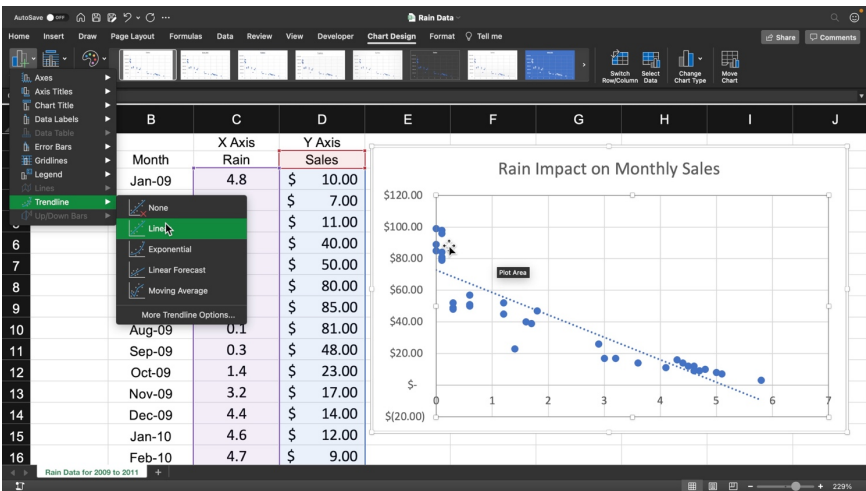
In this current chapter we will show you 2 ways to create a basic regression chart and then in an upcoming chapter called Advanced Linear Regression Analysis we will show you another way to create a Regression Analysis using the Excel analysis ToolPak.

For this example, let's assume that this restaurant has franchises only in the state of California. As you can see, we have the month listed in the image above in column B, then the inches in rain in column C and then the monthly sales for Los Pollos in column D. We got the month rain data from ChatGPT.

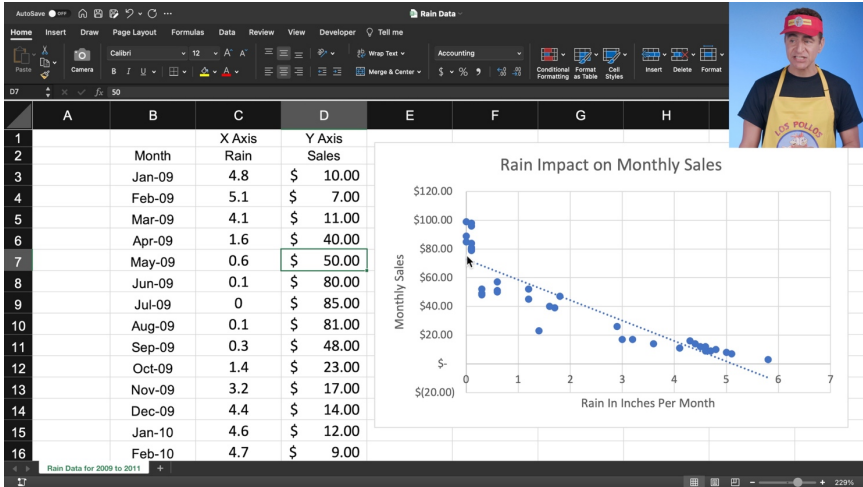
Per the image above, in Excel, let's highlight just the rain and sales data. Then, per the image below let's insert the first recommended chart.



Per the image below, let's now select the chart and then click on the chart design ribbon and select add chart element and then add a linear trendline.

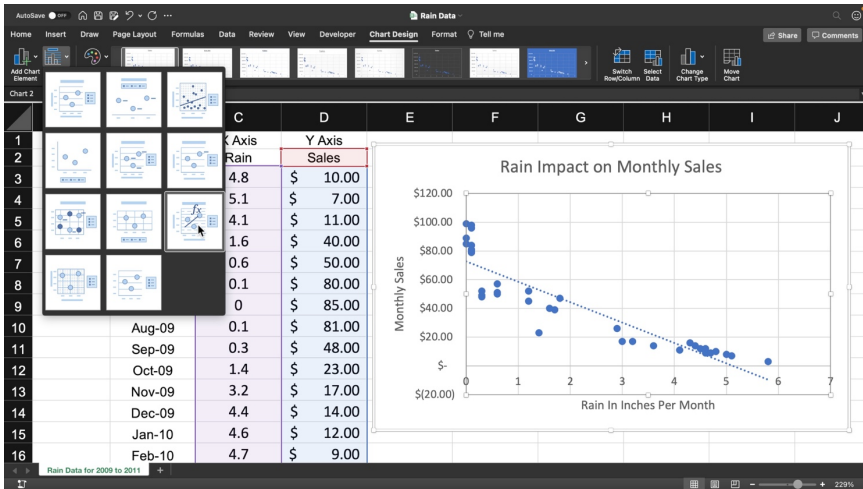


Per the image below, which contains the X and Y axis titles, we see that in this Linear Regression example that there is a clear pattern of more rain means lower sales.

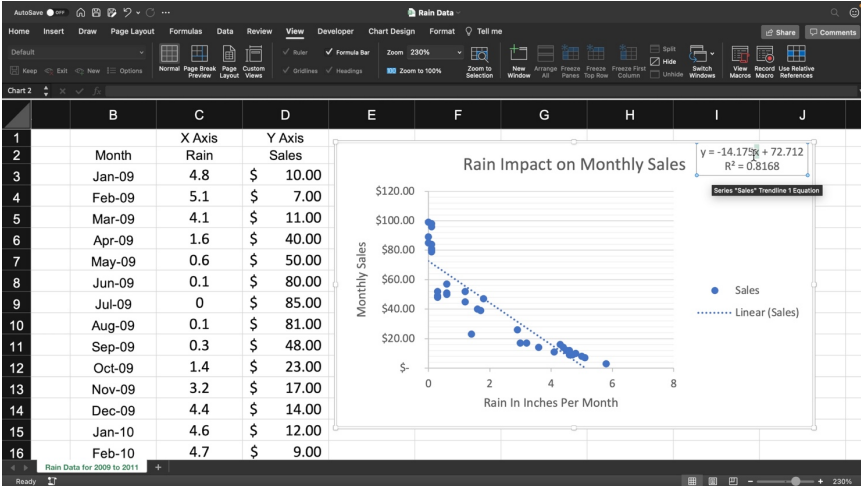


How do we quantify the data relationship - meaning what is the formula that explains that line?

Per the image below, let's select the quick layout icon and select the option with fx meaning formula in it.



We now see the formula:



Note: If you can't access the formula on the chart, maybe you have an older version of Excel, then we have the second way to get this regression formula which is select 2 cells beside each other and then type this `LINEST(Select Y data, Select X data)` and this will give you the slope meaning m and then the intercept meaning y, which I will explain in a second in more detail.

If your version of Excel is really old, then after you type the aforementioned formula then press shift and control and enter at the same time.

I also removed the negative sales number which is illogical by left clicking on the y axis data and then right clicking and then select format axis. Let's make the minimum here zero.

Quick side note. I was a beta tester for Microsoft years ago and developers in software products like Excel always anticipate what the user wants to do if you right click. So, if in doubt, right click.

Let's discuss this Linear Regression formula that we added to the chart.

We know that x is rain and y is sales. And we see here that there is -14.175 times x . This is the slope, meaning it tells you how the data is related...so for each 1 inch increase in rain...remember rain is x ...for each 1 inch increase in rain, sales, which is y , Sales drops by 14.175 dollars. We just made calculus and statistics fun :)

What does that 72.712 mean? Well, if we plug in 0 for rain, it means that the Linear Regression formula estimates that if rain is zero then sales will likely be 72.712. It makes sense because we go out to restaurants more often when it is not raining. So, 72.712 is the y axis.

This formula here of $y = -14.175x + 72.712$ explains the line. So, if we plug in any number for x meaning rain, the formula predicts sales for us. So, let's plug in rain in inches meaning x of 2.9 inches so we can solve for y which is sales.

2.9 (for x) times $-14.175 + 72.12 =$ sales of 31 dollars when the rain in inches is 2.9. We just did our first basic Linear Regression in Excel!

Last comment on Pollo Loco and it's Linear Regression Analysis is the y variable is sales which is dependent on the x variable which is rain. The y variable is the dependent variable, and the x variable is the independent variable.

Sales which is y is dependent on x which is rain in inches. We will discuss more advanced statistics when it comes to Regression Analysis in the chapter called Advanced Linear Regression in Excel and we will keep building on this in this section so that you

understand every single advanced statistic that Microsoft provides us with when we do advanced Regression Analysis.

When I used to work in the hedge fund industry, I would do a Regression Analysis just like this to see if there was a correlation between massive insider buying or selling of shares and how the stock reacted one year later.

Separately, when Elon Musk bought a ton of shares in Tesla during the pandemic in early 2020, he made a fortune as he later sold those shares in late 2021 before the market peak in January of 2022. Analyzing data allows us to be unemotional when making decisions based on data,.

This is an area of investing called Behavioral Finance and we need to always analyze the data to see if there is a pattern regardless of what industry we work in.

[A] Advanced Regression Analysis in Excel

In this chapter we will use the same data that we used for the basic Regression Analysis of our restaurant sales versus inches of rain. We will use a more advanced way to do a Regression Analysis in Excel using the Data Analysis ToolPak.

Let's talk about how to install the Data Analysis ToolPak for Excel. If you are using Windows, go to file, options add ins and then at the bottom select manage add ins then click on go and then select Analysis. On a Mac go to tools, Excel add ins and select Analysis ToolPak.

Before we insert the Regression Analysis, please note that we are going to discuss this in a lot of detail and please don't be intimidated by the math. We need to go into a lot of detail so that you understand how regression works in case you do it in the cloud or your team uses Python or another application and not in Excel, because I want us to understand everything.

We will build on this slowly in this chapter. We will go through several university level calculus and statistics classes in just this chapter. But instead of memorizing, we will understand the concepts together.

If you want, in the course version of this book go to the lesson called [A] Advanced Regression Analysis in Excel and download the attached data file.

Let's insert a Regression Analysis using the Analysis ToolPak that we just installed.

Step 1: click on the data ribbon and then select data analysis

Step 2: select regression and then ok

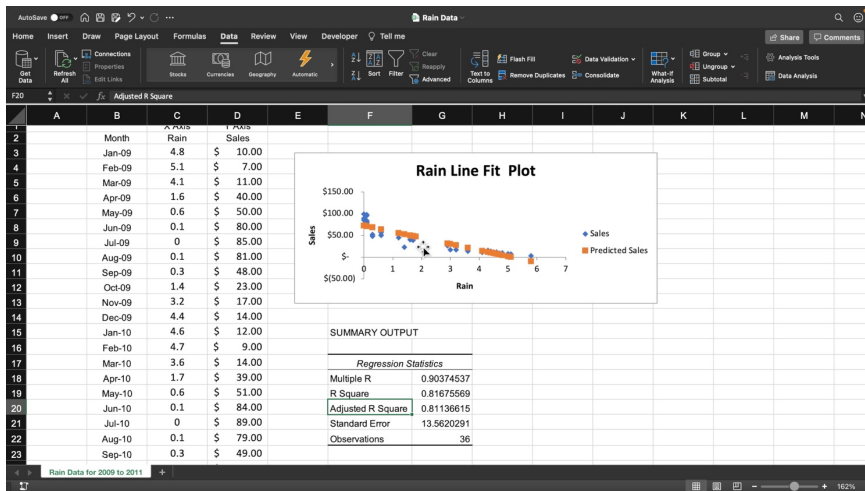
Step 3: highlight the x axis data and the y axis data including the labels - meaning sales and inches of rain.

Step 4: click the labels checkbox so that Excel knows we are selecting the titles in our data

Step 5: click the output radio button, and for the output range click cell F15 as this is where Excel will insert the statistics data and the chart.

Step 6: click the line fits plot checkbox and then ok

Now, per the image below, we have our advanced Linear Regression data and chart here.



Please don't be intimidated like I was by any of this as we will explain everything in detail.

One of the limitations of using Excel for Machine Learning and Regression Analysis is that the statistics summary output data is static meaning it won't change if we change the x and y data up here.

Let's talk about the chart first. Recall that the more rain we have as shown on the x axis, the lower the sales will be. The blue dots represent all the data points meaning sales based on the inches of rain. The orange dots here represent the predicted sales. Let's turn those orange dots into a line by righting clicking on the orange dots and then we select format data series and then the paint bucket and a solid line.

Now negative sales if rain is very high is illogical. So, let's left click on the y axis and then right click and select format axis and let's click axis options and select a minimum of 0.

This line is called the line of best fit because it is the best line that can be drawn that represents or fits the data points. It's also called the least squares regression line. Why least squares? Because as we will see when we look at the data, if you square the data then it tells us how close the data points are to the line.....closer is better....this line minimizes the variance of the data meaning the errors. We will build on this one step at a time.

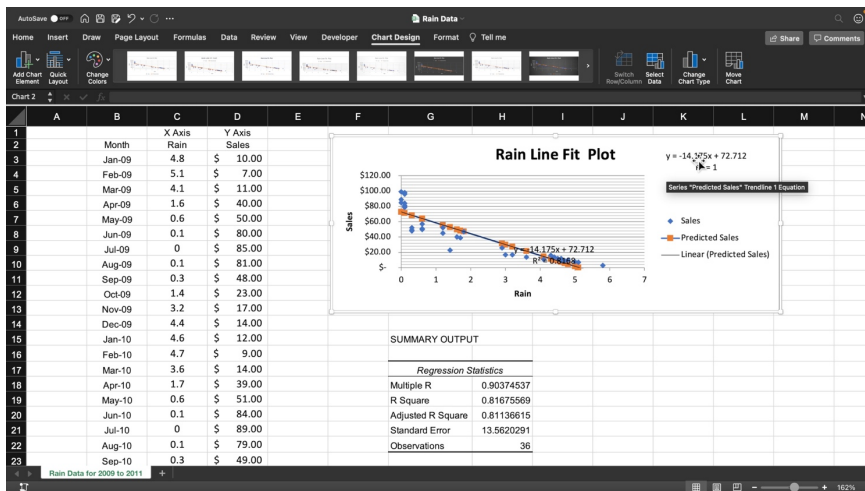
We will come back to this squared concept soon...but if you square a negative number it equals a positive number. We will come back to this.

Error means how far away the data point is from this line. We can see from the data points that some of the points are above the trend line, and some are below the trend line.

How do we measure the errors here, meaning the how far away these dots are from the line? Well, if we took the values above the line and we subtracted the value on the line then the numbers are positive.

For example, if I click the first data point above the line here that is not on the axis, we see this data point here is 0.1 inches of rain equals sales of \$80 but the trend line predicts lower sales.

What sales does the trend line predict when the inches of rain is 0.1 inches? Well let's click the chart and let's insert the formula for the line of best fit, as we did in the previous chapter. We do this by going to chart design e and then quick layout and then we add the formula, per the image below.



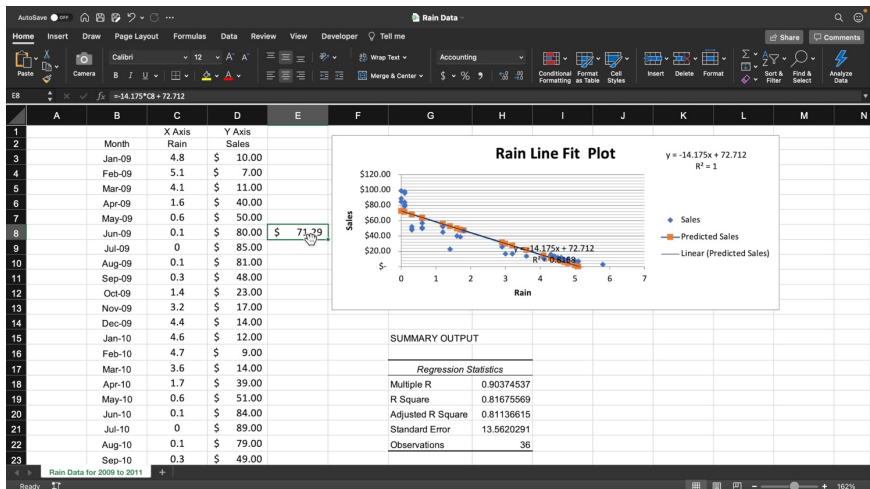
We see that the formula that describes the line is $y = -14.175x + 72.712$ and this makes sense because if we make rain which is x , 0 then y meaning sales, $y = 72.712$, which we see here on the chart. When x meaning rain is zero, y meaning sales is 72.712

The slope is -14.175 and this means for every 1 inch increase in rain meaning for every 1 inch increase in x , then our sales goes down 14.175 dollars. This makes sense because the more rain there is the less likely people are going to go out and eat at restaurants.

But let's get back to the discussion of the error. This line does it's best to represent the data....but if we plug in 0.1 inches of rain for x, then $y = -14.175 * .1 + 72.712$...this means when the rain meaning x is 0.1 then y meaning sales as predicted by the trend line here is about 71.2945. However, if I mouse over this data point above the line we see that 0.1 inches of rain with our data was \$80 in sales

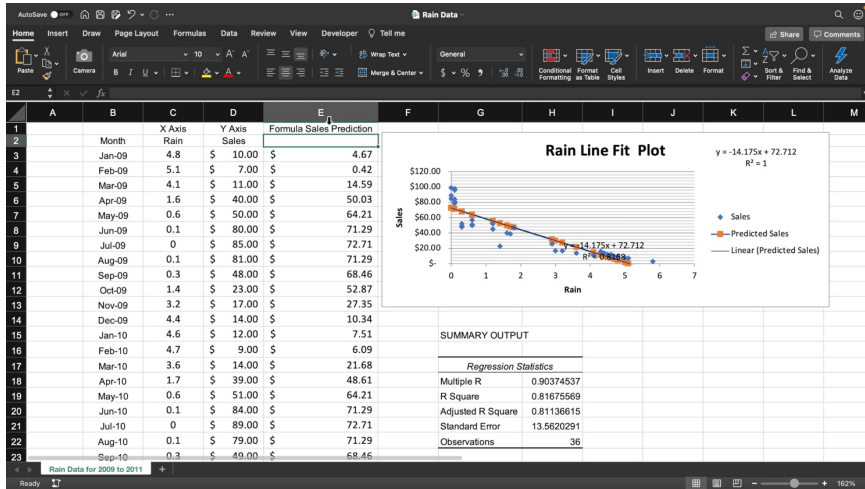
So, the error is \$80 minus what the formula calculated for 0.1 inches of rain, which is \$71.2945 in sales. \$80 minus \$71.2945 = POSITIVE 8.7055 for this particular data point. I emphasized the word positive there for a reason.

Let's type the formula of what we just did here in cell e8 per the image below.

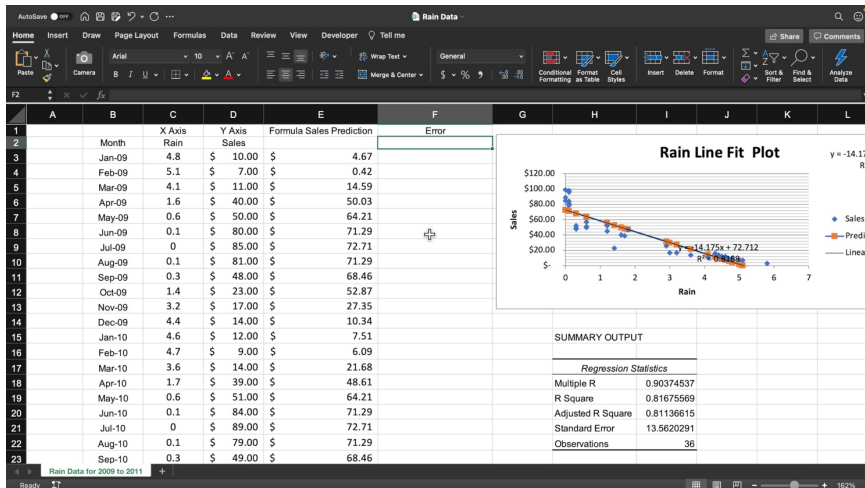


In the image above, we see the formula calculates when rain is 0.1 inches, the line formula predicts 71.2945 inches of rain.

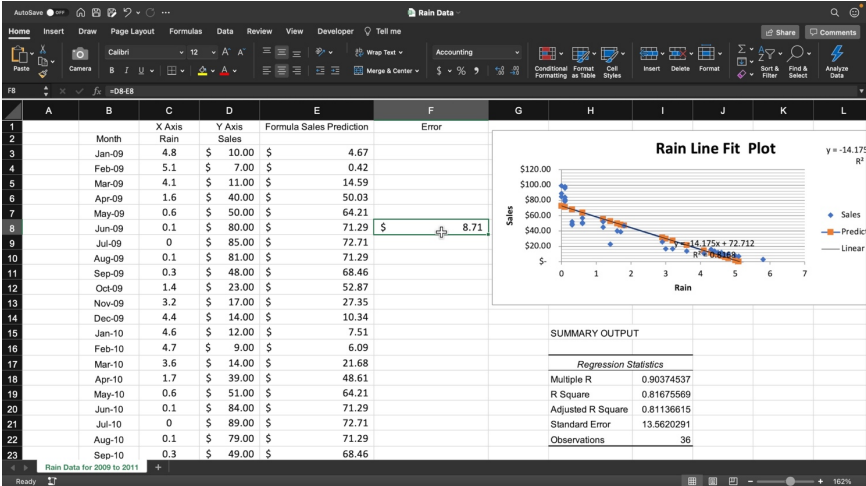
Let's copy the calculation in E8 and add it to every row where there is data per the image below. Let's title the column Formula Sales Prediction.



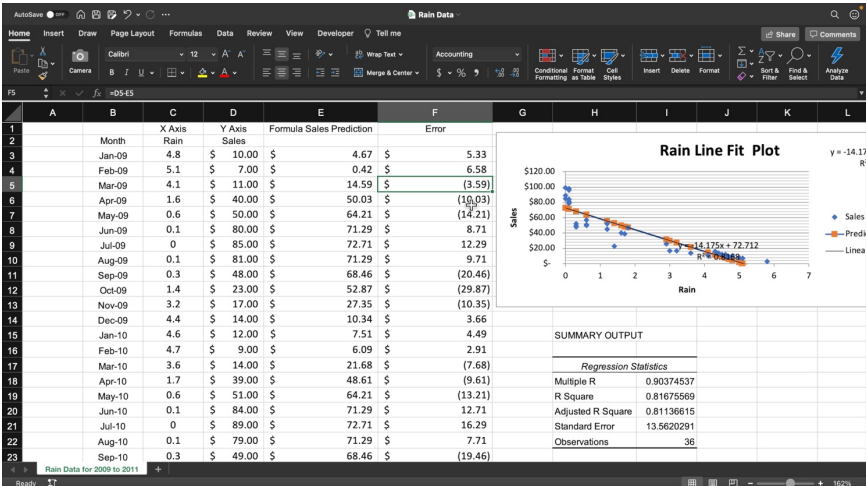
And let me insert a column to the right of column E and in E1 I will type Error, per the image below:



Then, per the image below, I will click cell F8, and I will take 80 in cell D8 here minus 71.2945 in cell , which is POSITIVE \$8.71.



Per the image below, I will copy F8 to all cells in the same column.



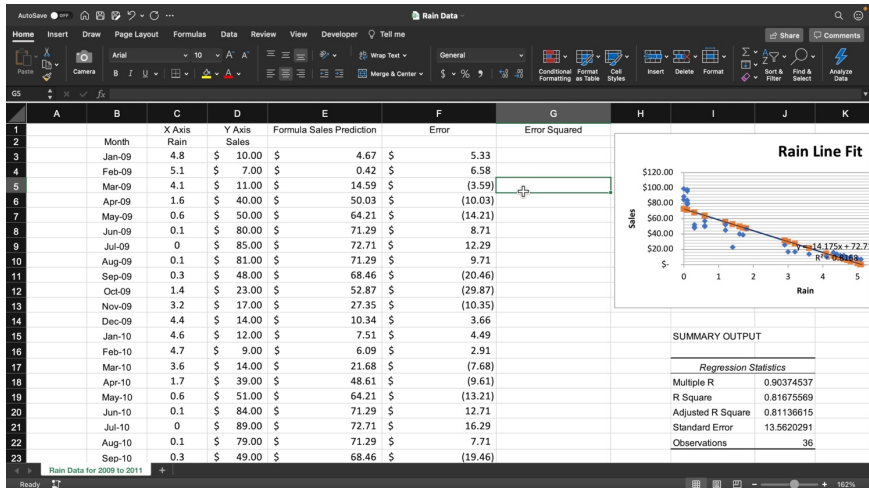
Some of the errors here are NEGATIVE. For example, in C5 when inches of rain was 4.1, the data was \$11 in sales but our formula here in cell E5 predicted a higher number meaning 14.59 and the error in cell F5 is a negative number...meaning negative 3.59...and if look at the chart here when rain is 4.1 inches , and if we mouse

over it shows 11 dollars in sales instead of what the formula meaning the line of best fit predicted, which was higher.

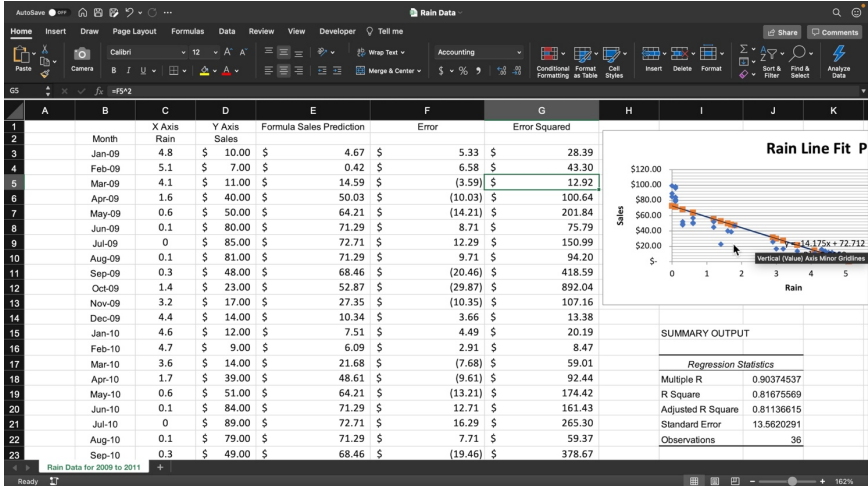
Why am I doing all this stuff? Because we can't just add up the errors as is because some are positive, and some are negative.

Here is the aha moment... if we square all the errors, then we get a positive number that explains the magnitude of the error for all 36 of these data points.

Per the image below, let me insert another column to the right of column F and let's call it Error Squared.



And here in column g let's square all these errors in column F, per the image below.



What we just did there is we calculated the distance between all 36 sales numbers in column D and we calculated what the formula meaning the line implies for all 36 data points in column E. Then we found the difference between columns D and E here in column F.

Column F is the error for each data point, meaning how much each data point is above or below the line. Then in column G we squared the difference, so the numbers are all positive.

Going forward, I'm going to highlight in green all the items in the Excel summary output that we have already discussed. So, per the image below, if we go down to H39....we see here there are 36 datapoints shown as 1 through 36. And this corresponds to our 36 data points on our chart.

	E	F	G	H	I	J	K	L	M	N	O
27	1.84	\$ 6.16	\$ 37.98	Residual	34	6253.57352	183.928633				
28	(9.50)	\$ 12.50	\$ 156.33	Total	35	34126.9722					
29	11.76	\$ 4.24	\$ 17.98								
30	47.20	\$ (0.20)	\$ 0.04								
31	64.21	\$ (7.21)	\$ 51.94	Intercept	72.71238626	3.33817912	21.78205005	1.4996E-21	65.9283901	79.4963825	65.928390
32	71.29	\$ 26.71	\$ 713.18	Rain	-14.17507689	1.15147664	-12.3103469	4.424E-14	-16.515159	-11.8349948	-16.51515
33	72.71	\$ 26.29	\$ 691.06								
34	71.29	\$ 24.71	\$ 610.36								
35	68.46	\$ (16.46)	\$ 270.92								
36	55.70	\$ (10.70)	\$ 114.53	RESIDUAL OUTPUT							
37	30.19	\$ (13.19)	\$ 173.90								
38	7.51	\$ 1.49	\$ 2.23	Observation	Predicted Sales	Residuals					
39				1	4.672011713	5.32798283					
40				2	0.419494105	6.5805059					
41				3	14.594571	-3.594571					
42				4	50.03226323	-10.0322632					
43				5	64.20734013	-14.2073401					
44				6	71.29487857	8.70512143					
45				7	72.71238626	12.2876137					
46				8	71.29487857	9.70512143					
47				9	68.45986319	-20.4598632					
48				10	52.86727861	-29.8672786					
49				11	27.3521402	-10.3521402					
50				12	10.34204793	3.65795207					
51				13	7.507032552	4.49296745					

And beside this starting in cell I39 we see the predicted sales meaning the trend line and our formula for each data point. If we scroll up to E1 we see that this data here we calculated, which is the same as what we see below starting in I38 so I will highlight I38 and below it and shade it as green.

	D	E	F	G	H	I	J
34	\$ 96.00	\$ 71.29	\$ 24.71	\$ 610.36			
35	\$ 52.00	\$ 68.46	\$ (16.46)	\$ 270.92			
36	\$ 45.00	\$ 55.70	\$ (10.70)	\$ 114.53	RESIDUAL OUTPUT		
37	\$ 17.00	\$ 30.19	\$ (13.19)	\$ 173.90			
38	\$ 9.00	\$ 7.51	\$ 1.49	\$ 2.23	Observation	Predicted Sales	Residuals
39					1	4.672011713	5.32798283
40					2	0.419494105	6.5805059
41					3	14.594571	-3.594571
42					4	50.03226323	-10.0322632
43					5	64.20734013	-14.2073401
44					6	71.29487857	8.70512143
45					7	72.71238626	12.2876137
46					8	71.29487857	9.70512143
47					9	68.45986319	-20.4598632
48					10	52.86727861	-29.8672786
49					11	27.3521402	-10.3521402
50					12	10.34204793	3.65795207
51					13	7.507032552	4.49296745
52					14	0.089524652	2.91047514
53					15	21.68210945	-7.68210945
54					16	48.61475554	-9.61475554
55					17	64.20734013	-13.2073401
56					18	71.29487857	12.7051214
57					19	72.71238626	16.2876137
58					20	71.29487857	7.70512143
59					21	68.45986319	-19.4598632

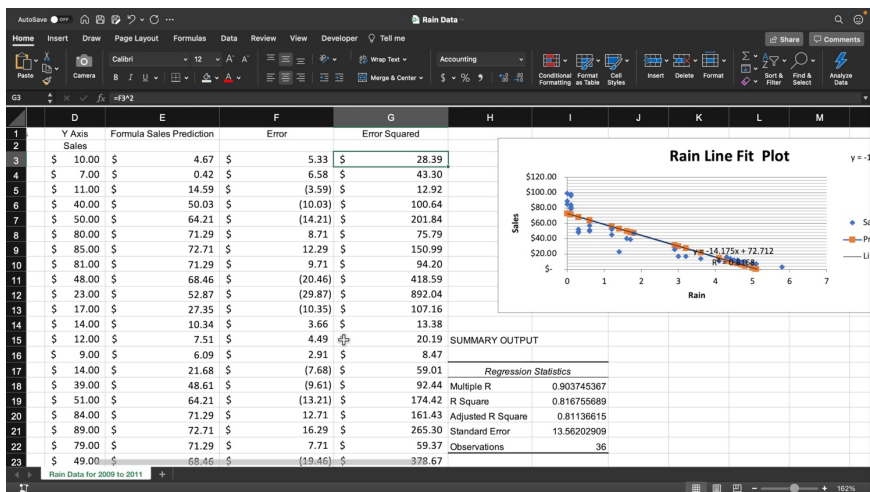
Then if we go up to our error calculator in F1 we see that the data here for our error is identical to what Excel calculated down here

starting in J38. So let me highlight this J38 cell and below it in green as we understand the math of it, per the image above

Let's keep building on this and I know it's a lot to take in for me at least and instead of just memorizing and relying on Excel, we need to understand this, so we never forget it. We can never just rely on AI or Machine Learning calculations without really understanding what is going on behind the scenes.

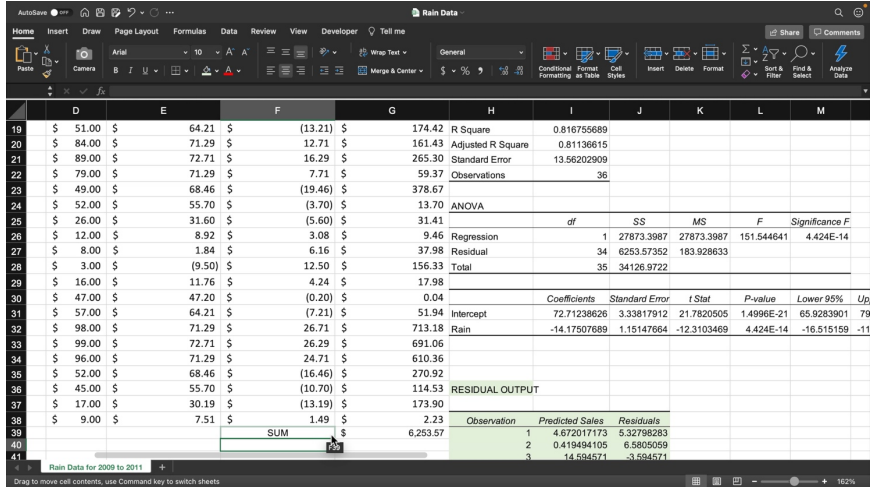
AI products like ChatGPT are often confidently wrong and if you double check all the data you will see that AI products like ChatGPT can make many mistakes.

Let's go look at the error squared column again starting in cell G1:



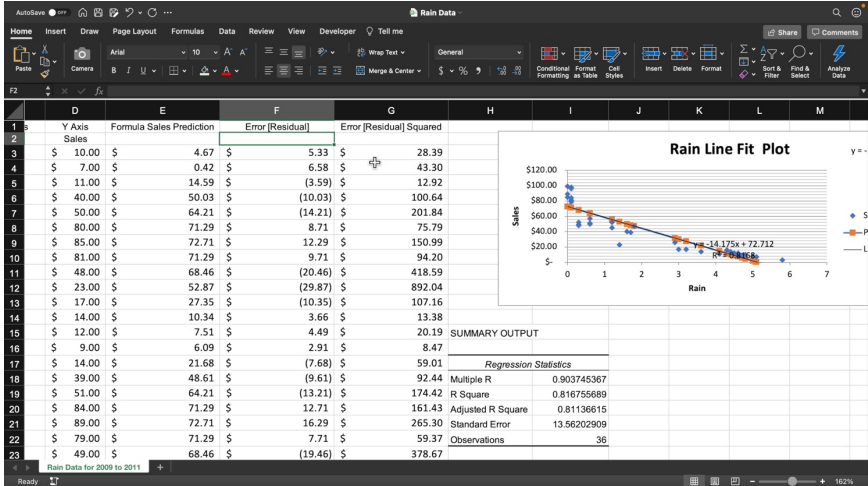
Recall we squared the error from column F here in column G so we can understand how big the errors are for all data points; if we didn't do this squaring, then if we added up the sum of all errors it might be a low number because we have many positive and many negative numbers here in column F.

So, let's now try to understand the magnitude of the errors by adding up all the squared errors here in column G. Per the image below, let's go to cell G39 and let's sum the errors squared so again we can understand the magnitude of our error.



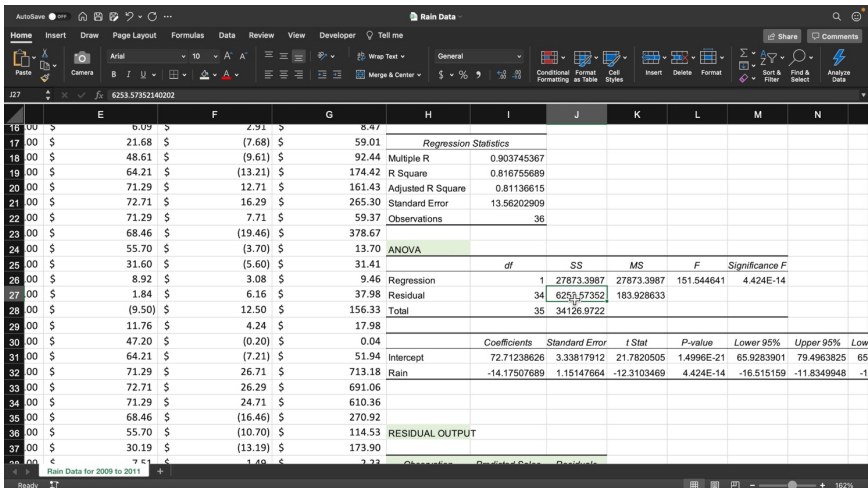
We see the sum of our squared errors is 6,253.57. Data Scientists like to call the error the Residual.

So, per the image below. let me got to F1 and G1 and beside the word error, in square brackets I will write Residual:



So, the Residual Sum of Squares equals what we see in G39 is 6,253.57.

What we are doing here is we are analyzing the variance meaning the magnitude of our error; we are doing the “Analysis Of Variance” which is also called ANOVA...meaning Analysis Of Variance....and here in cell H24 we see ANOVA meaning Analysis Of Variance:



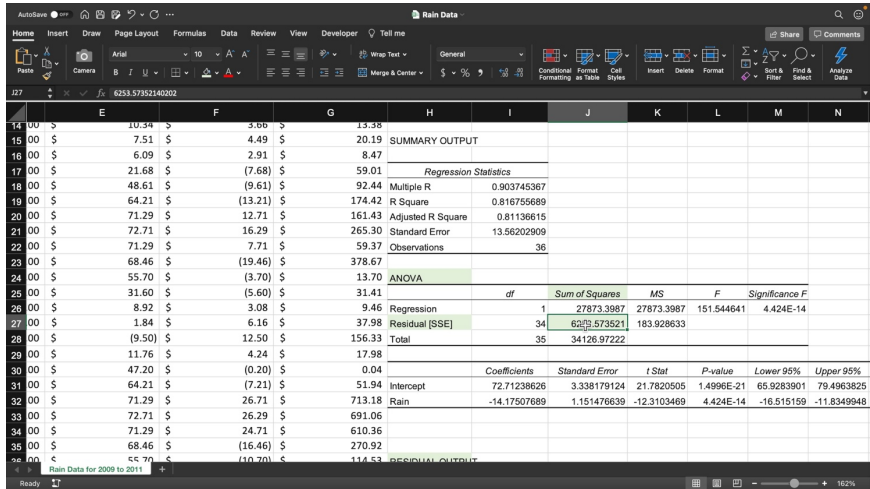
Our error Sum of Squares was 6,253.57...recall that error is called Residual... so our Residual Sum of Squares meaning our Residual SS or Sum of Squares is 6,253.57, which is what here see here in cell J27:

	df	SS	MS	F	Significance F
Regression	1	27873.3987	27873.3987	151.544641	4.424E-14
Residual	34	6253.57362	183.928633		
Total	35	34126.9722			

Let me also highlight in green the word Residual and SS meaning Sum of Squares and let me change SS here to Sum of Squares:

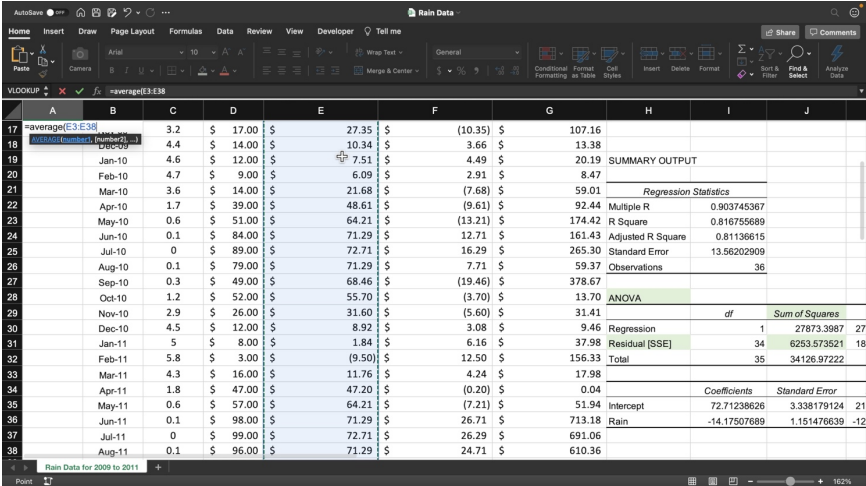
	df	Sum of Squared	MS	F	Significance F
Regression	1	27873.3987	27873.3987	151.544641	4.424E-14
Residual	34	6253.57362	183.928633		
Total	35	34126.9722			

Now some people call this 6,253.5 number here the SSE meaning Sum of Squares Error. So let me type that here in the green Residual cell in square brackets meaning SSE and I will shade it green:

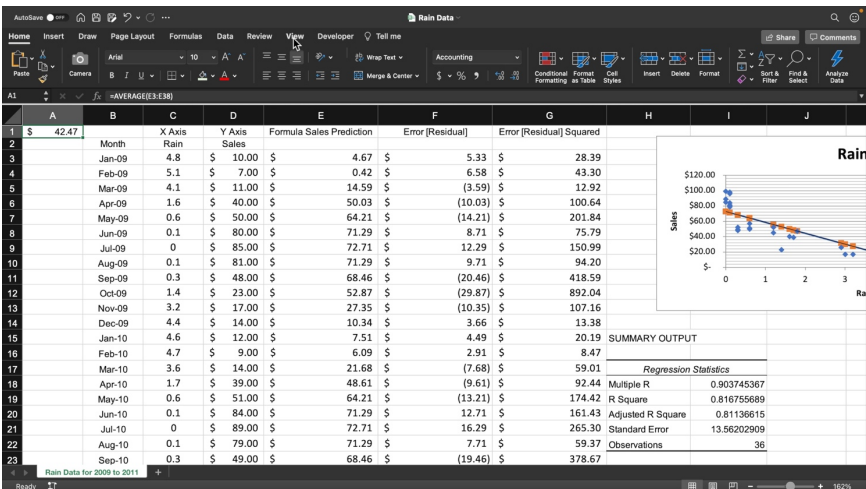


Great we quantified the error...what do we do with that Sum of Squares Error that we just calculated? The SSE of 6,253.5 is the magnitude of our error. But it's the magnitude of the error compared to what? Well, we got to figure out what the magnitude of this error is when compared to our regression line on our chart, so we can figure out how accurate our model is.

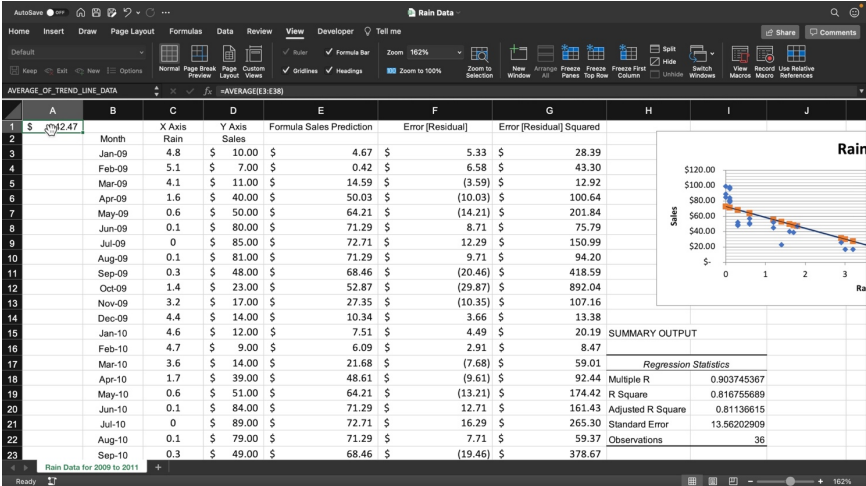
Per the image below, let's look at all the data points here for our predicted sales in column E and up here in cell A1 I will calculate the average for all the data points in column E:



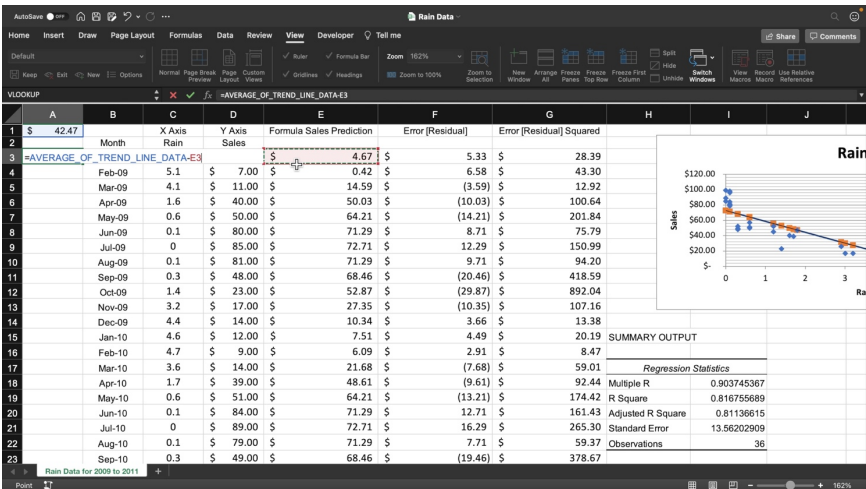
We see that the average of all calculations on our trend line is 42.472:

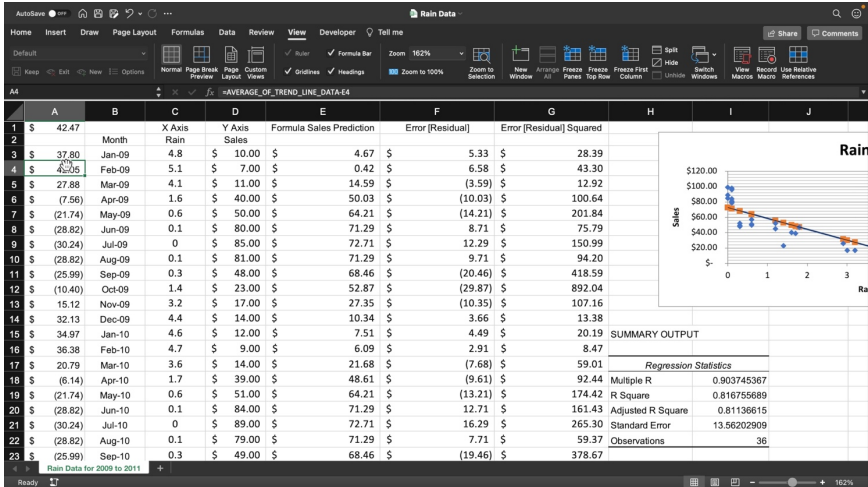


If you can't see the formula bar here, then go to View and then click on show formula bar. Let me name this A1 cell here AVERAGE_OF_TREND_LINE_DATA; if you name cells like I just did, make sure there is no space between the words:



And let's calculate the variance of our data on the trend line versus the average which we just calculated as 42.472. This will make more sense in a second. So here in cell A3 I will take the AVERAGE_OF_TREND_LINE_DATA minus what the trend line predicted for all data points:





Great we now see how far away each data point on our trend line is from the average...but wait a second some of these data points here in column A are negative....and if I just added up all these data points here in column A, it might not tell us the magnitude of how far away the average is from each data point on this line. What should I do?

Well similar to the SSE or Sum of Squares errors calculation from G39 we should square the numbers up here in column A, per the image below:

The screenshot shows an Excel spreadsheet with the following data in columns A through J:

	A	B	C	D	E	F	G	H	I	J	
19	\$ 472.41	May-10	0.6	\$ 51.00	\$ 64.21	\$ (13.21)	\$ 174.42	R Square	0.816756889		
20	\$ 830.74	Jun-10	0.1	\$ 84.00	\$ 71.29	\$ 12.71	\$ 161.43	Adjusted R Square	0.811366115		
21	\$ 914.46	Jul-10	0	\$ 89.00	\$ 72.71	\$ 16.29	\$ 265.30	Standard Error	13.56202909		
22	\$ 830.74	Aug-10	0.1	\$ 79.00	\$ 71.29	\$ 7.71	\$ 59.37	Observations	36		
23	\$ 875.35	Sep-10	0.3	\$ 49.00	\$ 68.46	\$ (19.46)	\$ 378.67				
24	\$ 175.03	Oct-10	1.2	\$ 52.00	\$ 55.70	\$ (3.70)	\$ 13.70	ANOVA			
25	\$ 116.10	Nov-10	2.9	\$ 26.00	\$ 31.60	\$ (5.60)	\$ 31.41				
26	\$ 1,125.43	Dec-10	4.5	\$ 12.00	\$ 8.92	\$ 3.08	\$ 9.46	Regression	df	Sum of Squares	
27	\$ 1,651.20	Jan-11	5	\$ 8.00	\$ 1.84	\$ 6.16	\$ 37.98	Residual [SSE]	1	27873.3987	23
28	\$ 2,701.40	Feb-11	5.8	\$ 3.00	\$ (9.50)	\$ 12.50	\$ 156.33	Total	34	6253.573521	16
29	\$ 943.26	Mar-11	4.3	\$ 16.00	\$ 11.76	\$ 1.84	\$ 4.24				
30	\$ 22.33	Apr-11	1.8	\$ 47.00	\$ 47.20	\$ (0.20)	\$ 0.04				
31	\$ 472.41	May-11	0.6	\$ 57.00	\$ 64.21	\$ (7.21)	\$ 51.94	Intercept	72.71238626	3.338179124	21
32	\$ 830.74	Jun-11	0.1	\$ 98.00	\$ 71.29	\$ 26.71	\$ 713.18	Rain	-14.17507689	1.151476639	-12
33	\$ 914.46	Jul-11	0	\$ 99.00	\$ 72.71	\$ 26.29	\$ 691.06				
34	\$ 830.74	Aug-11	0.1	\$ 96.00	\$ 71.29	\$ 24.71	\$ 610.36				
35	\$ 675.35	Sep-11	0.3	\$ 52.00	\$ 68.46	\$ (16.46)	\$ 270.92				
36	\$ 175.03	Oct-11	1.2	\$ 45.00	\$ 55.70	\$ (10.70)	\$ 114.53	RESIDUAL OUTPUT			
37	\$ 150.92	Nov-11	3	\$ 17.00	\$ 30.19	\$ (13.19)	\$ 173.90				
38	\$ 1,222.55	Dec-11	4.6	\$ 9.00	\$ 7.51	\$ 1.49	\$ 2.23	Observation	Predicted Sales	Residuals	
39	\$ 27,873.10	<--SUM					\$ 6,253.57	1	4.672017173	5.327982827	23
40								2	0.419494105	6.580055895	
41								3	14.594571	-3.594570398	

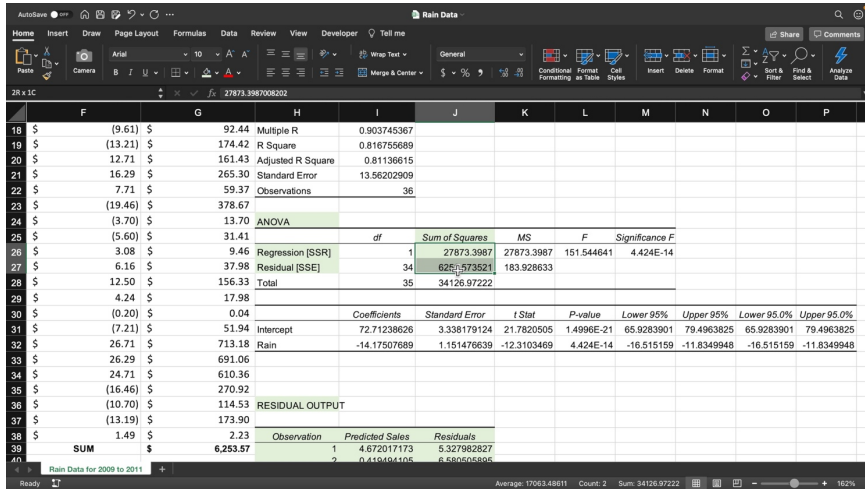
Now, per the image above, let's add up all these items here like we did with the SSE. So, in cell A39, I will add up all the squared numbers from our regression line and we see that it equals 27,873.1

And per the image below, in our Analysis Of Variance, meaning ANOVA table, we basically have the same result of 27,873...let me shade this number in green:

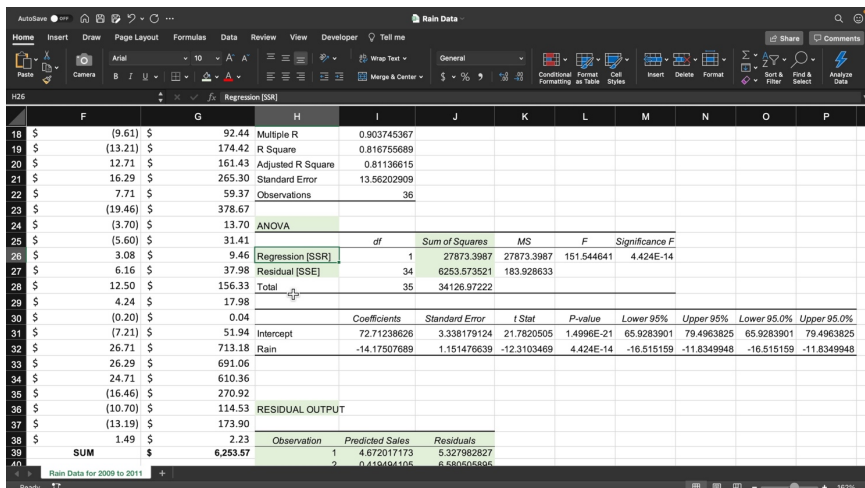
The screenshot shows an Excel spreadsheet with the following data in columns F through P:

	F	G	H	I	J	K	L	M	N	O	P
18	\$ (9.61)	\$ 92.44	Multiple R	0.903745367							
19	\$ (13.21)	\$ 174.42	R Square	0.816756889							
20	\$ 12.71	\$ 161.43	Adjusted R Square	0.811366115							
21	\$ 16.29	\$ 265.30	Standard Error	13.56202909							
22	\$ 7.71	\$ 59.37	Observations	36							
23	\$ (19.46)	\$ 378.67									
24	\$ (3.70)	\$ 13.70	ANOVA								
25	\$ (5.60)	\$ 31.41		df	Sum of Squares	MS	F	Significance F			
26	\$ 3.08	\$ 9.46	Regression	1	27873.3987	27873.3987	151.544641	4.424E-14			
27	\$ 6.16	\$ 37.98	Residual [SSE]	34	6253.573521	183.928633					
28	\$ 12.50	\$ 156.33	Total	35	34126.97222						
29	\$ 4.24	\$ 17.98									
30	\$ (0.20)	\$ 0.04									
31	\$ (7.21)	\$ 51.94	Intercept	72.71238626	3.338179124	21.7820505	1.4096E-21	65.9283901	79.4983625	65.9283901	79.4983625
32	\$ 26.71	\$ 713.18	Rain	-14.17507689	1.151476639	-12.3103469	4.424E-14	-16.515159	-11.8349948	-16.515159	-11.8349948
33	\$ 26.29	\$ 691.06									
34	\$ 24.71	\$ 610.36									
35	\$ (16.46)	\$ 270.92									
36	\$ (10.70)	\$ 114.53	RESIDUAL OUTPUT								
37	\$ (13.19)	\$ 173.90									
38	\$ 1.49	\$ 2.23	Observation	Predicted Sales	Residuals						
39	\$ SUM	\$ 6,253.57	1	4.672017173	5.327982827						
40			2	0.419494105	6.580055895						

And in cell H26 we have the regression or SSR meaning Sum of Squares Regression. Let me type SSR in square brackets beside regression here and I will make this label green as well:



Now if I highlight cells J26 and J27 meaning SSR and SSE, we get 34,127 in cell J28, which is the total, meaning SST...Sum of Squares Total. Let me add SST in square brackets here in H28 and let me make this green too:



What do we do with this? Well now we can calculate how accurate our regression line is relative to the error. So, if I take the SSR in J26 divided by the SST in J28 we see that a number of 81.68 percent which we see in cell I19. So let me highlight H19 and I19 now in green:

	F	G	H	I	J	K	L	M	N	O	P		
18	\$	(9.61)	\$	92.44	Multiple R	0.903745367							
19	\$	(13.21)	\$	174.42	R Square	0.816755689							
20	\$	12.71	\$	161.43	Adjusted R Square	0.811396615							
21	\$	16.29	\$	265.30	Standard Error	13.56202909							
22	\$	7.71	\$	59.37	Observations	36							
23	\$	(19.46)	\$	378.67									
24	\$	(3.70)	\$	13.70	ANOVA								
25	\$	(5.60)	\$	31.41		df	Sum of Squares	MS	F	Significance F			
26	\$	3.08	\$	9.46	Regression [SSR]	1	27873.3987	27873.3987	151.544641	4.424E-14			
27	\$	6.16	\$	37.98	Residual [SSE]	34	6253.573521	183.928633					
28	\$	12.50	\$	156.33	Total [SST]	35	34126.97222						
29	\$	4.24	\$	17.98									
30	\$	(0.20)	\$	0.04									
31	\$	(7.21)	\$	51.94	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%	
32	\$	26.71	\$	713.18	Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.9283901	79.4963825	65.9283901	79.4963825
33	\$	26.29	\$	691.06	Rtain	-14.17507689	1.151476639	-12.3103469	4.424E-14	-16.515159	-11.8349948	-16.515159	-11.8349948
34	\$	24.71	\$	610.36									
35	\$	(16.46)	\$	270.92									
36	\$	(10.70)	\$	114.53	RESIDUAL OUTPUT								
37	\$	(13.19)	\$	173.90									
38	\$	1.49	\$	2.23	Observation	Predicted Sales	Residuals						
39	\$		\$	6.25357	1	4.672017173	5.327982827						
40	\$		\$		2	11.96191105	6.581055845						

Our R Squared in cell J19 is 81.68% and this means that our regression line is 81.68 percent accurate. The highest this number goes is 1 meaning 100% and it's debatable, but I usually like to see an r squared above .7 meaning 70% accurate.

So, we can conclude that our regression line meaning our regression formula representing the line which is $y = -14.175x + 72.712$ is 82% accurate.

Per the image below, SSR divided by SST is 82%. And the SSE here divided by SST is 18%. So, our model is 82% accurate and 18% inaccurate:

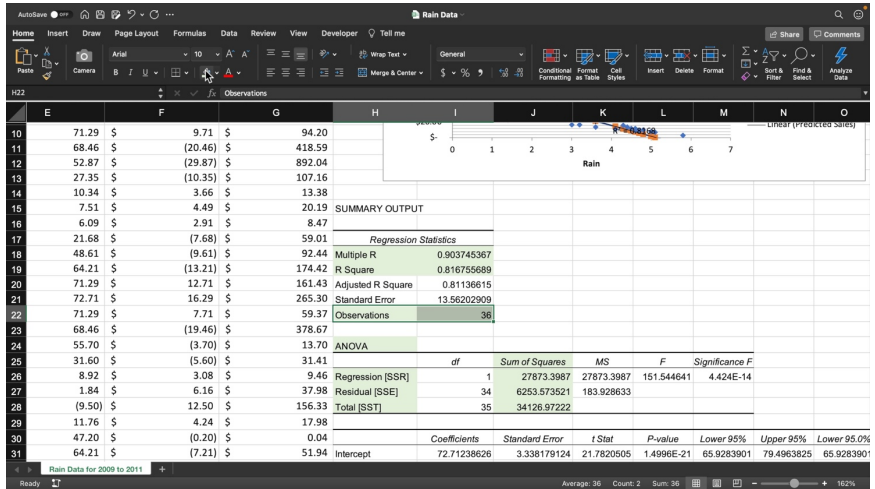
	F	G	H	I	J	K	L	M	N	O	P		
18	\$	(9.61)	\$	92.44	Multiple R	0.903745367							
19	\$	(13.21)	\$	174.42	R Square	0.816755689							
20	\$	12.71	\$	161.43	Adjusted R Square	0.81136615							
21	\$	16.29	\$	265.30	Standard Error	13.56202909							
22	\$	7.71	\$	59.37	Observations	36							
23	\$	(19.46)	\$	378.67									
24	\$	(3.70)	\$	13.70	ANOVA								
25	\$	(5.60)	\$	31.41									
26	\$	3.08	\$	9.46	Regression [SSR]	27873.3987	MS	F	Significance F				
27	\$	6.16	\$	37.98	Residual [SSE]	34	6253.573521	183.928633	4.424E-14				
28	\$	12.50	\$	156.33	Total [SST]	35	34126.97222						
29	\$	4.24	\$	17.98									
30	\$	(0.20)	\$	0.04									
31	\$	(7.21)	\$	51.94	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%	
32	\$	26.71	\$	713.18	Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.9283901	79.4963825	65.9283901	79.4963825
33	\$	26.29	\$	691.06	Rain	-14.17507689	1.151476639	-12.3103469	4.424E-14	-16.515159	-11.8349948	-16.515159	-11.8349948
34	\$	24.71	\$	610.36									
35	\$	(16.46)	\$	270.92									
36	\$	(10.70)	\$	114.53	RESIDUAL OUTPUT								
37	\$	(13.19)	\$	173.90									
38	\$	1.49	\$	2.23	Observation	Predicted Sales	Residuals						
39	\$		\$		1	4.672017173	5.327962827						
40		SUM	\$	6,253.57									

What is this number above the r squared in cell I18? It means the correlation of x versus y. So, if I type =correl(Y data, X data), we get the same value as the r squared:

	F	G	H	I	J	K	L	M	N	O	P		
15	\$	4.49	\$	20.19	SUMMARY OUTPUT								
16	\$	2.91	\$	8.47									
17	\$	(7.68)	\$	59.01	Regression Statistics								
18	\$	(9.61)	\$	92.44	Multiple R	0.903745367							
19	\$	(13.21)	\$	174.42	R Square	0.816755689							
20	\$	12.71	\$	161.43	Adjusted R Square	0.81136615							
21	\$	16.29	\$	265.30	Standard Error	13.56202909							
22	\$	7.71	\$	59.37	Observations	36							
23	\$	(19.46)	\$	378.67									
24	\$	(3.70)	\$	13.70	ANOVA	0.816755689							
25	\$	(5.60)	\$	31.41									
26	\$	3.08	\$	9.46	Regression [SSR]	27873.3987	MS	F	Significance F				
27	\$	6.16	\$	37.98	Residual [SSE]	34	6253.573521	183.928633	4.424E-14				
28	\$	12.50	\$	156.33	Total [SST]	35	34126.97222						
29	\$	4.24	\$	17.98									
30	\$	(0.20)	\$	0.04									
31	\$	(7.21)	\$	51.94	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%	
32	\$	26.71	\$	713.18	Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.9283901	79.4963825	65.9283901	79.4963825
33	\$	26.29	\$	691.06	Rain	-14.17507689	1.151476639	-12.3103469	4.424E-14	-16.515159	-11.8349948	-16.515159	-11.8349948
34	\$	24.71	\$	610.36									
35	\$	(16.46)	\$	270.92									
36	\$	(10.70)	\$	114.53	RESIDUAL OUTPUT								

Let's discuss some of the non-green shaded numbers created by Excel's Regression Analysis Data ToolPak feature.

Per the image below, in cell I22 we see that there are 36 observations meaning 36 data points so let me shade the word observations and data points here in green.



Per the image below, we see that the coefficients data in cell I30 that the intercept is 72 as represented by our formula on the chart. We also see that our slope is -14.175 also as represented in the formula on our chart. So, let's shade the intercept and rain data in green.

	F	G	H	I	J	K	L	M	N	O	P	Q
14	3.66	\$	13.38									
15	4.49	\$	20.19	SUMMARY OUTPUT								
16	2.91	\$	8.47									
17	(7.68)	\$	59.01	Regression Statistics								
18	(9.61)	\$	92.44	Multiple R	0.903745307							
19	(13.21)	\$	174.42	R Square	0.816755689							
20	12.71	\$	161.43	Adjusted R Square	0.81136615							
21	16.29	\$	265.30	Standard Error	13.56202909							
22	7.71	\$	59.37	Observations	36							
23	(19.46)	\$	378.67									
24	(3.70)	\$	13.70	ANOVA								
25	(5.60)	\$	31.41		df	Sum of Squares	MS	F	Significance F			
26	3.08	\$	9.46	Regression [SSR]	1	27873.3987	27873.3987	151.544641	4.424E-14			
27	6.16	\$	37.98	Residual [SSE]	34	6253.573521	183.928633					
28	12.50	\$	156.33	Total [SST]	35	34126.97222						
29	4.24	\$	17.98									
30	(0.20)	\$	0.04		Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%
31	(7.21)	\$	51.94	Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.9283901	79.4963825	65.9283901	79.4963825
32	26.71	\$	713.18	Rain	-14.17507689	1.151476639	-12.3103469	4.424E-14	-16.515159	-11.8349948	-16.515159	-11.8349948
33	26.29	\$	691.06									
34	24.71	\$	610.36									
35	(16.46)	\$	270.92									

Recall that the slope of -14.175 means that for every 1 inch increase in rain which is x, our sales which is y goes down 14.175 dollars. When x is 0 meaning when rain is zero, then y is 72.712, which is the Y intercept when x is zero.

We are making great progress!

We will skip adjusted r squared in cell I20 for now because in another chapter when we discuss Multiple Regression Analysis...meaning when we have more than just one unknown variable, then the adjusted r squared will make more sense as we will adjust for more variables instead of just one variable, which is x. Recall that x represents rain in inches.

Let's move onto what we have not discussed in the ANOVA table. Per the image below, In I25 we have df. df stands for degrees of freedom and it just tells you about independent variables in the calculation. So, we know that our formula is $y = mx + b$ and we know this equals $y = -14.715x + 72.712$so if we have x then we have one item to derive in our regression, which is y. So, think of the

degrees of freedom regression # of 1 here in cell I26 as one unknown variable.

The screenshot shows an Excel spreadsheet with the following data:

	F	G	H	I	J	K	L	M	N	O	P	Q		
14		3.66 \$	13.38											
15		4.49 \$	20.19	SUMMARY OUTPUT										
16		2.91 \$	8.47											
17		(7.68) \$	59.01	Regression Statistics										
18		(9.61) \$	92.44	Multiple R	0.903745367									
19		(13.21) \$	174.42	R Square	0.816755689									
20		12.71 \$	161.43	Adjusted R Square	0.81136615									
21		16.29 \$	265.30	Standard Error	13.56202909									
22		7.71 \$	59.37	Observations	36									
23		(19.46) \$	378.67	ANOVA										
24		(3.70) \$	13.70											
25		(5.60) \$	31.41	df	Sum of Squares	MS	F	Significance F						
26		3.08 \$	9.46	Regression [SSR]	1	27873.3987	27873.3987	151.544641	4.424E-14					
27		6.16 \$	37.98	Residual [SSE]	34	6253.573521	183.928633							
28		12.50 \$	156.33	Total [SST]	35	34126.97222								
29		4.24 \$	17.98											
30		(0.20) \$	0.04	Coefficients										
31		(7.21) \$	51.94	Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.9283901	79.4963825	65.9283901	79.4963825		
32		26.71 \$	713.18	Rain	-14.17507689	1.151476639	-12.3103469	4.424E-14	-16.515159	-11.8349948	-16.515159	-11.8349948		
33		26.29 \$	691.06											
34		24.71 \$	610.36											
35		(16.46) \$	270.92											

In cell I28 we see that the total number of degrees of freedom here is 35, which means our 36 observations minus 1....why minus one? Because if we plug in one of the 36 results, then there are 35 other instances that we can plug in. So, the total degrees of freedom is 35.

Per the image below, the next column in the Analysis Of Variance or ANOVA table starting in K25 is MS, which refers to Mean or Average Sum of Squares. Let me rename K25 Mean Sum of Squares:

	F	G	H	I	J	K	L	M	N	O	P
14		3.66 \$	13.38								
15		4.49 \$	20.19								
16		2.91 \$	8.47								
17		(7.68) \$	59.01								
18		(9.61) \$	92.44								
19		(13.21) \$	174.42								
20		12.71 \$	161.43								
21		16.29 \$	265.30								
22		7.71 \$	59.37								
23		(19.46) \$	378.67								
24		(3.70) \$	13.70								
25		(5.60) \$	31.41								
26		3.08 \$	9.46								
27		6.16 \$	37.98								
28		12.50 \$	156.33								
29		4.24 \$	17.98								
30		(0.20) \$	0.04								
31		(7.21) \$	51.94								
32		26.71 \$	713.18								
33		26.29 \$	691.06								
34		24.71 \$	610.36								
35		(16.46) \$	270.92								

	df	Sum of Squares	Mean Sum of Squares	F	Significance F
Regression [SSR]	1	27873.3987	27873.3987	151.544641	4.424E-14
Residual [SSE]	34	6253.573521	183.928633		
Total [SST]	35	34126.97222			

	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%
Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.9283901	79.4963825	65.9283901	79.4963825
Rain	-14.17507689	1.151476639	-12.31034692	4.424E-14	-16.515159	-11.8349948	-16.515159	-11.8349948

This Mean Sum of Squares column is the Sum of Squares column divided by the df column. For example, in K27 the math is the Residual Sum of Squares of 6253.57 divided by the degrees of freedom of 34, which is 183.93. Let's shade this MS meaning Mean Sum of Squares column in green now:

	F	G	H	I	J	K	L	M	N	O	P
14		3.66 \$	13.38								
15		4.49 \$	20.19								
16		2.91 \$	8.47								
17		(7.68) \$	59.01								
18		(9.61) \$	92.44								
19		(13.21) \$	174.42								
20		12.71 \$	161.43								
21		16.29 \$	265.30								
22		7.71 \$	59.37								
23		(19.46) \$	378.67								
24		(3.70) \$	13.70								
25		(5.60) \$	31.41								
26		3.08 \$	9.46								
27		6.16 \$	37.98								
28		12.50 \$	156.33								
29		4.24 \$	17.98								
30		(0.20) \$	0.04								
31		(7.21) \$	51.94								
32		26.71 \$	713.18								
33		26.29 \$	691.06								
34		24.71 \$	610.36								
35		(16.46) \$	270.92								

	df	Sum of Squares	Mean Sum of Squares	F	Significance F
Regression [SSR]	1	27873.3987	27873.3987	151.544641	4.424E-14
Residual [SSE]	34	6253.573521	183.928633		
Total [SST]	35	34126.97222			

	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%
Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.9283901	79.4963825	65.9283901	79.4963825
Rain	-14.17507689	1.151476639	-12.31034692	4.424E-14	-16.515159	-11.8349948	-16.515159	-11.8349948

Next up in L25 we have the “F” column, which divides the Mean Sum of Squares numbers. So, the Mean Sum of Squares for regression divided by the Mean Sum of Squares for the Residual is 151.54. Meaning K26 divided by K27:

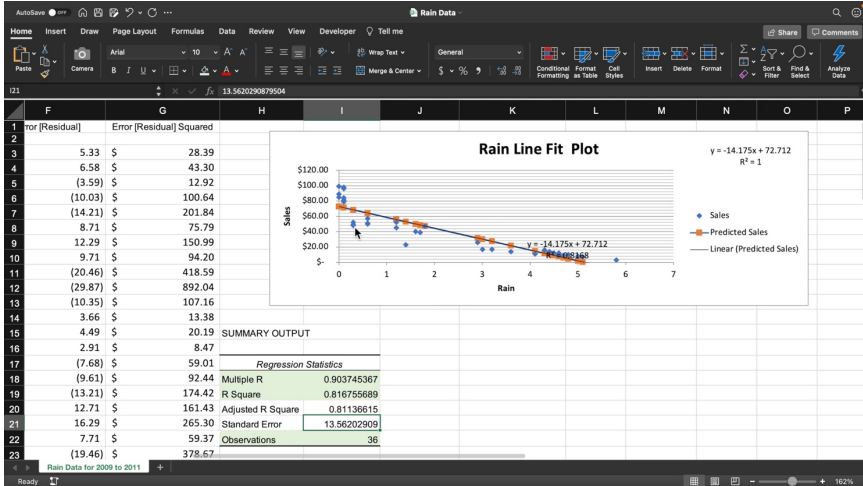
Regression Statistics		ANOVA		Coefficients			
		df	Sum of Squares	Mean Sum of Squares	F	Significance F	
Multiple R	0.903745367	1	27873.3987	27873.3987	151.544641	4.424E-14	
R Square	0.816755689	34	6253.673521	183.928633			
Adjusted R Square	0.81136615	35	34126.97222				
Standard Error	13.56202909						
Observations	36						
				t Stat	P-value	Lower 95%	Upper 95%
Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.9283901	79.4963825	65.9283901
Rain	-14.17507689	1.151476639	-12.31034692	4.424E-14	-16.515159	-11.8349948	-16.515159

This just helps us understand another way to look at the magnitude of our error and the variance of the data.

What does this significance F mean in cell M26? 4.424E-14, means the decimal point after the first number 4 adds 14 zeros before the number 4... meaning a very tiny number.

Well, we want this to be lower than 0.05 so we can conclude that our model is statistically significant meaning the model works. Since 0.05 is way greater than this number here, then we can conclude that our model is statistically significant. This significance F is also called the P value.

Now let’s talk about the standard error in I21, which is 13.56:



This standard error number represents how far away the data points are on average from the trendline. The way to calculate the standard error of 13.56 is we take the SSE of 6253.573 / (the number of data points -2)...meaning J27 divided by I27.

Then since we squared the data before let's get rid of the square by taking the square root, which equals a standard error of 13.56.

A lower number is better of course meaning a lower number means the data points are closer to the trend line. The 13.56 standard error is how far apart on average the data is from the line.

Generally speaking, most of the data points are usually within 2 times the standard error on the regression line....so 13.56 times 2, which is 27.12. If we look at our regression line it makes sense as most values are within 27.12 away from the regression line although there are some outliers.

Let's look at the intercept data in cell I31. Per the image below, the y intercept when x is 0 is 72.71 and if we look at cells O31 and P31 we are 95% sure that for the entire population of data that the y

intercept is between 65.928 and 79.496...please don't spend too much time on the intercept confidence data as it isn't really helpful as it only reflects when x is zero...spend more time analyzing the data here for the next line which is x or rain in our example.

The screenshot shows an Excel spreadsheet with the following data:

	H	I	J	K	L	M	N	O	P	Q	R	S
19	42	R Square	0.816755689									
20	43	Adjusted R Square	0.81136615									
21	30	Standard Error	13.56202909									
22	37	Observations	36									
24	70	ANOVA										
25	41		df	Sum of Squares	Mean Sum of Squares	F	Significance F					
26	46	Regression [SSR]	1	27873.3987	27873.3987	151.544641	4.42397E-14					
27	98	Residual [SSE]	34	6253.573521	183.928633							
28	33	Total [SST]	35	34126.97222								
30	04		Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%		
31	94	Intercept	72.71238626	3.338179124	21.7820505	1.4998E-21	65.92839007	79.4963925	65.9283901	79.4963925		
32	18	Rain	-14.17507689	1.154478639	-12.31034692	4.424E-14	-16.51515897	-11.8349948	-16.515159	-11.8349948		
36												
38	53	RESIDUAL OUTPUT										
39	57		Observation	Predicted Sales	Residuals							
40			1	4.672017173	5.327982827							
			2	0.419494105	6.580505895							

Let's look at the slope data. In cell I32 we see that the slope is -14.18 and if we look at cells O32 and P32 we are 95% sure that for the entire population of data that the slope is between -16.52 and -11.83:

	G	H	I	J	K	L	M	N	O	P	Q	R
20		161.43	Adjusted R Square	0.81136615								
21		265.30	Standard Error	13.56202909								
22		59.37	Observations	36								
23		378.67										
24		13.70	ANOVA									
25		31.41		df	Sum of Squares	Mean Sum of Squares	F	Significance F				
26		9.46	Regression [SSR]	1	27873.3987	27873.3987	151.544641	4.42397E-14				
27		37.98	Residual [SSE]	34	6253.573521	183.928633						
28		156.33	Total [SST]	35	34126.97222							
29		17.98										
30		0.04		Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%	
31		51.94	Intercept	72.71238626	3.338179124	21.7820505	1.4996E-21	65.92839007	79.4963825	65.92839001	79.4963825	
32		713.18	Rain	-14.17507689	1.151476639	-12.31034692	4.424E-14	-16.51515897	-11.8349948	-16.515159	-11.8349948	
33		691.06										
34		610.36										
35		270.92										
36		114.53	RESIDUAL OUTPUT									
37		173.90										
38		2.23	Observation	Predicted Sales	Residuals							
39		6,253.57	1	4.672017173	5.327982827							
40			2	0.419494105	6.580505895							
41			3	14.594571	-3.594570998							
42			4	50.00000000	-14.00000000							

Now you will see that there are several cells not shaded in green yet. We will cover this in more detail when we do Multiple Regression Analysis.

Let's stop right here as we covered a lot in this Advanced Linear Regression Analysis chapter. We haven't yet discussed the non-green shaded cells yet, but we will soon.

We also haven't discussed the adjusted R square in cell I20.

We will discuss what these cells mean when we do the Advanced Multiple Regression Analysis chapter. Multiple meaning, we add another column of data.

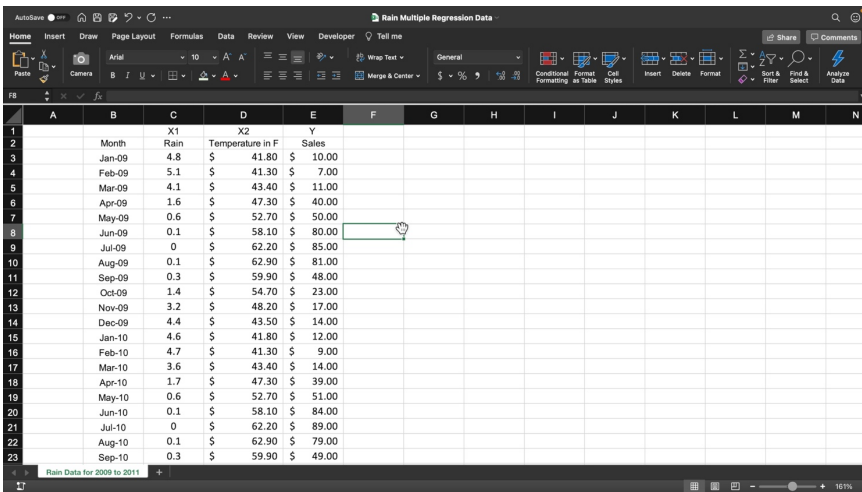
[A] Multiple Regression Analysis in Excel

Now that we understand how to do a Regression Analysis in Excel when we only have one unknown variable, which is x for rain, let's discuss Multiple Regression Analysis - meaning when we have more than just one unknown variable x.

Please make sure you have read or understand the previous Linear Regression chapters in this section before reading this chapter on Multiple Regression. Thanks

Let's download the Excel file attached to the lecture in the course version of this book called lecture called [A] Multiple Regression Analysis in Excel.

Let's look at the data per the image below:



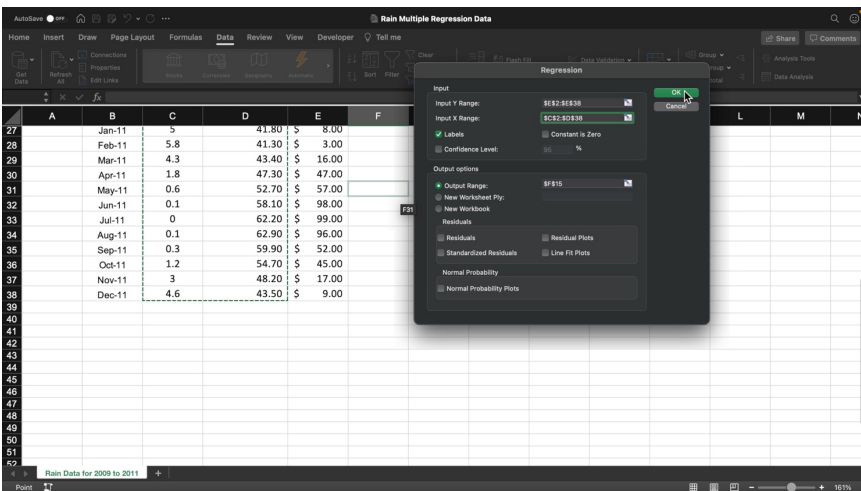
The screenshot shows an Excel spreadsheet with the following data:

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1			X1	X2	Y									
2		Month	Rain	Temperature in F	Sales									
3		Jan-09	4.8	\$ 41.80	\$ 10.00									
4		Feb-09	5.1	\$ 41.30	\$ 7.00									
5		Mar-09	4.1	\$ 43.40	\$ 11.00									
6		Apr-09	1.6	\$ 47.30	\$ 40.00									
7		May-09	0.6	\$ 52.70	\$ 50.00									
8		Jun-09	0.1	\$ 58.10	\$ 80.00									
9		Jul-09	0	\$ 62.20	\$ 85.00									
10		Aug-09	0.1	\$ 62.90	\$ 81.00									
11		Sep-09	0.3	\$ 59.90	\$ 48.00									
12		Oct-09	1.4	\$ 54.70	\$ 23.00									
13		Nov-09	3.2	\$ 48.20	\$ 17.00									
14		Dec-09	4.4	\$ 43.50	\$ 14.00									
15		Jan-10	4.6	\$ 41.80	\$ 12.00									
16		Feb-10	4.7	\$ 41.30	\$ 9.00									
17		Mar-10	3.6	\$ 43.40	\$ 14.00									
18		Apr-10	1.7	\$ 47.30	\$ 39.00									
19		May-10	0.6	\$ 52.70	\$ 51.00									
20		Jun-10	0.1	\$ 58.10	\$ 84.00									
21		Jul-10	0	\$ 62.20	\$ 89.00									
22		Aug-10	0.1	\$ 62.90	\$ 79.00									
23		Sep-10	0.3	\$ 59.90	\$ 49.00									

In our Multiple Regression Analysis, we are using the same data that we used in the previous chapter. Meaning rain in inches for x and sales for y. Let's call rain x1 instead of x. And here we added another variable which is temperature in Fahrenheit called x2. So, we have multiple unknown variables here meaning x1 for rain in

inches and x2 for temperature in Fahrenheit. We will keep sales as y.

Let's run our Multiple Regression now in Excel and this is the exact same thing we did with Linear Regression, except, when we select x, we highlight the x1 and the x2 data. We won't select chart here because instead of an x axis and a y axis there would be an x1 axis, an x2 axis and a y axis, which is 3 dimensional and tough to interpret in Excel. I will add the output data to cell F15:



		SUMMARY OUTPUT							
		Regression Statistics							
Multiple R		0.92015975							
R Square		0.846669397							
Adjusted R Sq		0.8374027							
Standard Error		12.5913277							
Observations		36							
		ANOVA							
		df	SS	MS	F	Significance F			
Regression		2	28895.1016	14447.5508	91.1278615	3.6453E-14			
Residual		33	5231.87057	158.541533					
Total		35	34126.9722						
		Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%	Lower 95.0%	Upper 95.0%
Intercept		-43.8946188	46.0383927	-0.95343508	0.34730175	-137.560433	49.7711954	-137.560433	49.7711954
Rain		-6.86452811	3.07181001	-2.23468512	0.0323258	-13.1141726	-0.61488366	-13.1141726	-0.61488366
Temperature		1.96775002	0.77513818	2.53857967	0.01603077	0.39071954	3.5447805	0.39071954	3.5447805

If we look at the data, we see that the multiple r is 0.92 and we always want this to be closer to 1. The 0.92 means that the independent variables meaning x1 for rain and x2 for temperature have a 92% correlation to y, which is sales.

Now just for fun, let's see if x1 and x2 are correlated....meaning rain and temperature. Let's type =correl (and let's select the data for x1 and x2). Not surprisingly there is a negative correlation here because when it rains it is often colder because the rain falls from the clouds and the clouds are colder than on the ground...which is why we have more snow on mountains because they are higher up.

Let's get back to our Multiple Regression data.

Let's move onto r squared, which is multiple r squared and it tells us how much the data points are close to what y predicts. So, r squared tells us the variation of the data points.

However, when Excel calculates the r squared with many more independent variables, it makes r squared go up even if the new variables like x3 x4 x5 are not correlated....as a result, Excel

calculates the adjusted r squared here in cell G20 which more accurately explains the r squared.

We pay more attention to adjusted r squared when we have more than one independent variable. In our case we have 2 independent variables meaning x1 for rain in inches and x2 for temperature in Fahrenheit.

Moving onto the standard error, which is 12.59, which tells us how far away our data points are from the line. Lower is better. Next we have 36 observations...meaning 36 rows of data.

Since we already discussed what many of these items mean in a previous chapter, I will spend a bit more time on items we have discussed in detail.

Let's talk about this significance of F value, which is also referred to as the P value. Since we want our data to be 95% accurate, which means 5% inaccurate, we want this significance of f to be less than 0.05. We know that 5% is 0.05.

So, we see here that we have e-14 at the end of this number which means a decimal and then 14 zeros before the number 3....so yes 0.05 is greater than this value here so we are confident that there is a correlation between our data points in our model and our model is statistically significant and yes we can use our model to predict y meaning sales.

Moving on to the intercept which is about -43. This means according to our trend line, when x1 and x2 are zero then y meaning sales is -43 but this doesn't make sense because we can't have negative sales and we only look at the trend line for predictive purposes within the data points we have...meaning if I type min or

max for the data points for x1 and x2 I have to use variables within the min and max of each.

So, if you type min and max for each column(for the x1 and for the x2 data) and we can only use our model when x1 which is rain is between 0 and 5.8 inches and we can only use x2 meaning temperature between 41.3 degrees Fahrenheit and 62.9 degrees Fahrenheit.

If this is hard to understand then think of the extremes. If we entered Fahrenheit of 500 degrees which is 260 degrees Celsius and we assumed that rain was 0, then our model would predict massive sales, which is illogical. So, we only use x1 and x2 data points between our min and max numbers here.

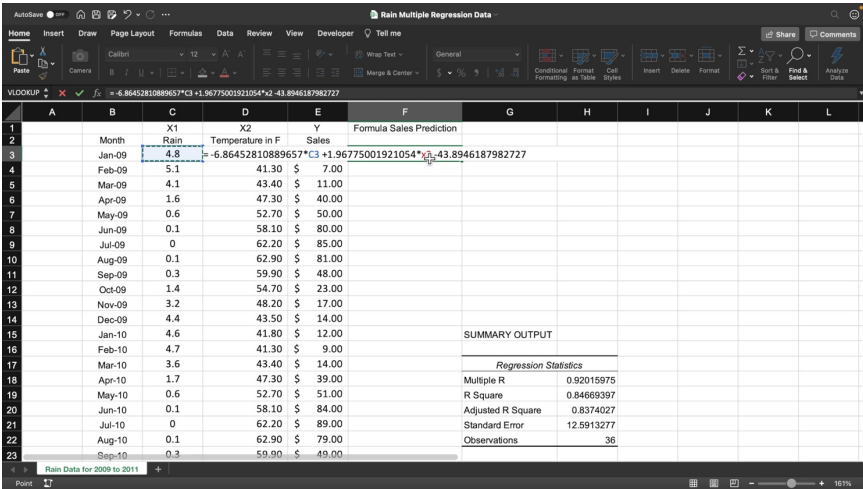
Looking across this intercept line of -44, the model is 95% confident that the intercept is between 49.77 and negative 137.6 ...wait a second that seems wrong....let me come back to this one here in a second and we don't really care much about the confidence interval of the intercept....we focus more on the independent variables meaning x1 and x2 which represents rain in inches and temperature respectively.

Let's create the formula for this regression. Recall that for a Linear Regression with one unknown variable meaning x we used $y = mx + b$. B is the intercept and M is the slope. So, let's rewrite this as $y = mx_1 + mx_2 + b$.

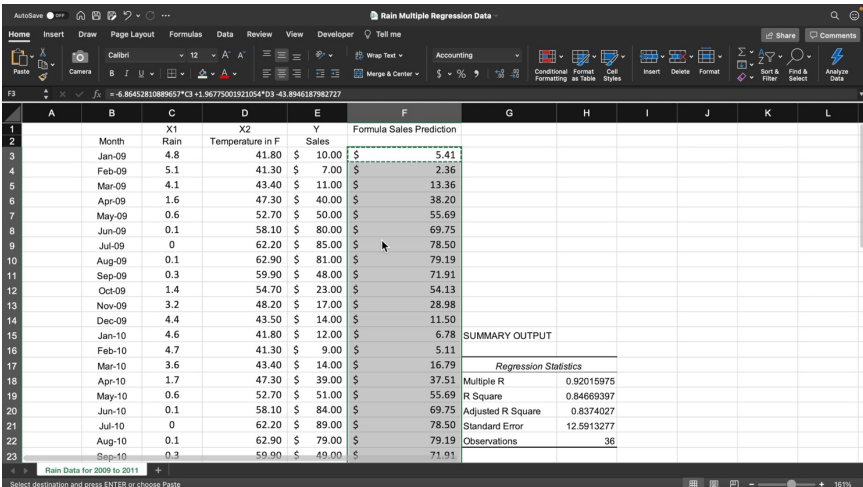
Now let's write the formula based on what we see here in the coefficients column here in cell G30 [er the image above. $y = -6.8645x_1 + 1.96775x_2 - 43.8946$

Similar to what we did in the previous chapter called Advanced Regression Analysis in Excel, we will insert columns and calculate the Sum of Squares per the image below:

Step 1 I will highlight column F and I will insert a new column with the name of Formula Sales Prediction



I will then calculate what the formula predicts for all 36 data points using our formula:



Now I will add another column called Error [Residual] and I will calculate the difference between the actual sales and the predicted sales:

The screenshot shows an Excel spreadsheet titled "Rain Multiple Regression Data". The data is organized as follows:

	A	B	C	D	E	F	G	H	I	J	K
1			X1	X2	Y	Formula Sales Prediction	Error (Residual)				
2		Month	Rain	Temperature in F	Sales						
3		Jan-09	4.8	41.80	\$ 10.00	\$ 4.61	\$ 4.59				
4		Feb-09	5.1	41.30	\$ 7.00	\$ 2.36	\$ 4.64				
5		Mar-09	4.1	43.40	\$ 11.00	\$ 13.36	\$ (2.36)				
6		Apr-09	1.6	47.30	\$ 40.00	\$ 38.20	\$ 1.80				
7		May-09	0.6	52.70	\$ 50.00	\$ 55.69	\$ (5.69)				
8		Jun-09	0.1	58.10	\$ 80.00	\$ 69.75	\$ 10.25				
9		Jul-09	0	62.20	\$ 85.00	\$ 78.50	\$ 6.50				
10		Aug-09	0.1	62.90	\$ 81.00	\$ 79.19	\$ 1.81				
11		Sep-09	0.3	59.90	\$ 48.00	\$ 71.91	\$ (23.91)				
12		Oct-09	1.4	54.70	\$ 23.00	\$ 54.13	\$ (31.13)				
13		Nov-09	3.2	48.20	\$ 17.00	\$ 28.98	\$ (11.98)				
14		Dec-09	4.4	43.50	\$ 14.00	\$ 11.50	\$ 2.50				
15		Jan-10	4.6	41.80	\$ 12.00	\$ 6.78	\$ 5.22	SUMMARY OUTPUT			
16		Feb-10	4.7	41.30	\$ 9.00	\$ 5.11	\$ 3.89				
17		Mar-10	3.6	43.40	\$ 14.00	\$ 16.79	\$ (2.79)	Regression Statistics			
18		Apr-10	1.7	47.30	\$ 39.00	\$ 37.51	\$ 1.49	Multiple R	0.92015975		
19		May-10	0.6	52.70	\$ 51.00	\$ 55.69	\$ (4.69)	R Square	0.84669397		
20		Jun-10	0.1	58.10	\$ 84.00	\$ 69.75	\$ 14.25	Adjusted R Square	0.8374027		
21		Jul-10	0	62.20	\$ 89.00	\$ 78.50	\$ 10.50	Standard Error	12.5913277		
22		Aug-10	0.1	62.90	\$ 79.00	\$ 79.19	\$ (0.19)	Observations	36		
23		Sep-10	0.3	59.90	\$ 49.00	\$ 71.91	\$ (22.91)				

Recall that we need these numbers to be positive. So, I will insert another column called Error [Residual] Squared and I will square all these numbers:

The screenshot shows the same Excel spreadsheet as above, but with an additional column for squared residuals. The data is organized as follows:

	A	B	C	D	E	F	G	H	I	J	K
1			X1	X2	Y	Formula Sales Prediction	Error (Residual)	Error (Residual) Squared			
2		Month	Rain	Temperature in F	Sales						
3		Jan-09	4.8	41.80	\$ 10.00	\$ 5.41	\$ 4.59	\$ 21.09			
4		Feb-09	5.1	41.30	\$ 7.00	\$ 2.36	\$ 4.64	\$ 21.49			
5		Mar-09	4.1	43.40	\$ 11.00	\$ 13.36	\$ (2.36)	\$ 5.58			
6		Apr-09	1.6	47.30	\$ 40.00	\$ 38.20	\$ 1.80	\$ 3.25			
7		May-09	0.6	52.70	\$ 50.00	\$ 55.69	\$ (5.69)	\$ 32.34			
8		Jun-09	0.1	58.10	\$ 80.00	\$ 69.75	\$ 10.25	\$ 105.16			
9		Jul-09	0	62.20	\$ 85.00	\$ 78.50	\$ 6.50	\$ 42.26			
10		Aug-09	0.1	62.90	\$ 81.00	\$ 79.19	\$ 1.81	\$ 3.27			
11		Sep-09	0.3	59.90	\$ 48.00	\$ 71.91	\$ (23.91)	\$ 571.89			
12		Oct-09	1.4	54.70	\$ 23.00	\$ 54.13	\$ (31.13)	\$ 969.14			
13		Nov-09	3.2	48.20	\$ 17.00	\$ 28.98	\$ (11.98)	\$ 143.63			
14		Dec-09	4.4	43.50	\$ 14.00	\$ 11.50	\$ 2.50	\$ 6.26			
15		Jan-10	4.6	41.80	\$ 12.00	\$ 6.78	\$ 5.22	\$ 27.24	SUMMARY OUTPUT		
16		Feb-10	4.7	41.30	\$ 9.00	\$ 5.11	\$ 3.89	\$ 15.13			
17		Mar-10	3.6	43.40	\$ 14.00	\$ 16.79	\$ (2.79)	\$ 7.80	Regression Statistics		
18		Apr-10	1.7	47.30	\$ 39.00	\$ 37.51	\$ 1.49	\$ 2.22	Multiple R	0.92015975	
19		May-10	0.6	52.70	\$ 51.00	\$ 55.69	\$ (4.69)	\$ 21.97	R Square	0.84669397	
20		Jun-10	0.1	58.10	\$ 84.00	\$ 69.75	\$ 14.25	\$ 203.20	Adjusted R Square	0.8374027	
21		Jul-10	0	62.20	\$ 89.00	\$ 78.50	\$ 10.50	\$ 110.26	Standard Error	12.5913277	
22		Aug-10	0.1	62.90	\$ 79.00	\$ 79.19	\$ (0.19)	\$ 0.04	Observations	36	
23		Sep-10	0.3	59.90	\$ 49.00	\$ 71.91	\$ (22.91)	\$ 525.06			

Now I will add the Sum of Squares for the Residual to the bottom here:

	A	B	C	D	E	F	G	H	I	J	K
19	May-10	0.6	52.70	\$ 51.00	\$ 55.69	\$ (4.69)	\$ 21.97	R Square	0.84669397		
20	Jun-10	0.1	58.10	\$ 84.00	\$ 69.75	\$ 14.25	\$ 203.20	Adjusted R Square	0.8374027		
21	Jul-10	0	62.20	\$ 89.00	\$ 78.50	\$ 10.50	\$ 110.26	Standard Error	12.5913277		
22	Aug-10	0.1	62.90	\$ 79.00	\$ 79.19	\$ (0.19)	\$ 0.04	Observations	36		
23	Sep-10	0.3	59.90	\$ 49.00	\$ 71.91	\$ (22.91)	\$ 525.06				
24	Oct-10	1.2	54.70	\$ 52.00	\$ 55.50	\$ (3.50)	\$ 12.28	ANOVA			
25	Nov-10	2.9	48.20	\$ 26.00	\$ 31.04	\$ (5.04)	\$ 25.44			df	SS
26	Dec-10	4.5	43.50	\$ 12.00	\$ 10.81	\$ 1.19	\$ 1.41	Regression	2	28895.16	
27	Jan-11	5	41.80	\$ 8.00	\$ 4.03	\$ 3.97	\$ 15.72	Residual	33	5231.876	
28	Feb-11	5.8	41.30	\$ 3.00	\$ (2.44)	\$ 5.44	\$ 29.60	Total	35	34126.91	
29	Mar-11	4.3	43.40	\$ 16.00	\$ 11.99	\$ 4.01	\$ 16.09				
30	Apr-11	1.8	47.30	\$ 47.00	\$ 36.82	\$ 10.18	\$ 103.55			Coefficients	Standard E
31	May-11	0.6	52.70	\$ 57.00	\$ 55.69	\$ 1.31	\$ 1.72	Intercept	-43.8946188	46.03833	
32	Jun-11	0.1	58.10	\$ 98.00	\$ 69.75	\$ 28.25	\$ 798.33	Rain	-6.86452811	3.071814	
33	Jul-11	0	62.20	\$ 99.00	\$ 78.50	\$ 20.50	\$ 420.27	Temperature In F	1.96775002	0.775138	
34	Aug-11	0.1	62.90	\$ 96.00	\$ 79.19	\$ 16.81	\$ 282.56				
35	Sep-11	0.3	59.90	\$ 52.00	\$ 71.91	\$ (19.91)	\$ 396.58			y = mx1 + mx2 + b	
36	Oct-11	1.2	54.70	\$ 45.00	\$ 55.50	\$ (10.50)	\$ 110.33				
37	Nov-11	3	48.20	\$ 17.00	\$ 30.36	\$ (13.36)	\$ 178.42				
38	Dec-11	4.6	43.50	\$ 9.00	\$ 10.13	\$ (1.13)	\$ 1.27				
39	min		0		41.3						
40	max		5.8		62.9						
41											

We see that our calculated Sum of Squares for the Residual is the same number as what we see here in K27. Think of Residual and error as being the same thing.

Next up we need to calculate the Sum of Squares for the regression line. So, in cell A1 let's find the average of the regression column and we see that the average is 42.47:

	A	B	C	D	E	F	G	H	I	J
1	\$ 42.47		X1	X2	Y	Formula Sales Prediction	Error (Residual)	Error (Residual) Squared		
2	\$ 37.06	Jan-09	4.8	41.80	\$ 10.00	\$ 5.41	\$ 4.59	\$ 21.09		
3	\$ 40.11	Feb-09	5.1	41.30	\$ 7.00	\$ 2.36	\$ 4.64	\$ 21.49		
4	\$ 29.11	Mar-09	4.1	43.40	\$ 11.00	\$ 13.36	\$ (2.36)	\$ 5.58		
5	\$ 4.28	Apr-09	1.6	47.30	\$ 40.00	\$ 38.20	\$ 1.80	\$ 3.25		
6	\$ (13.21)	May-09	0.6	52.70	\$ 50.00	\$ 55.69	\$ (5.69)	\$ 32.34		
7	\$ (27.27)	Jun-09	0.1	58.10	\$ 80.00	\$ 69.75	\$ 10.25	\$ 105.16		
8	\$ (36.03)	Jul-09	0	62.20	\$ 85.00	\$ 78.50	\$ 6.50	\$ 42.26		
9	\$ (36.72)	Aug-09	0.1	62.90	\$ 81.00	\$ 79.19	\$ 1.81	\$ 3.27		
10	\$ (29.44)	Sep-09	0.3	59.90	\$ 48.00	\$ 71.91	\$ (23.91)	\$ 571.89		
11	\$ (11.66)	Oct-09	1.4	54.70	\$ 23.00	\$ 54.13	\$ (31.13)	\$ 969.14		
12	\$ 13.49	Nov-09	3.2	48.20	\$ 17.00	\$ 28.98	\$ (11.98)	\$ 143.63		
13	\$ 30.97	Dec-09	4.4	43.50	\$ 14.00	\$ 11.50	\$ 2.50	\$ 6.26		
14	\$ 35.69	Jan-10	4.6	41.80	\$ 12.00	\$ 6.78	\$ 5.22	\$ 27.24		SUMMARY OUTPUT
15	\$ 37.36	Feb-10	4.7	41.30	\$ 9.00	\$ 5.11	\$ 3.89	\$ 15.13		
16	\$ 25.68	Mar-10	3.6	43.40	\$ 14.00	\$ 16.79	\$ (2.79)	\$ 7.80		Regression Statistics
17	\$ 4.96	Apr-10	1.7	47.30	\$ 39.00	\$ 37.51	\$ 1.49	\$ 2.22		Multiple R
18	\$ (13.21)	May-10	0.6	52.70	\$ 51.00	\$ 55.69	\$ (4.69)	\$ 21.97		R Square
19	\$ (27.27)	Jun-10	0.1	58.10	\$ 84.00	\$ 69.75	\$ 14.25	\$ 203.20		Adjusted R Square
20	\$ (36.03)	Jul-10	0	62.20	\$ 89.00	\$ 78.50	\$ 10.50	\$ 110.26		Standard Error
21	\$ (36.72)	Aug-10	0.1	62.90	\$ 79.00	\$ 79.19	\$ (0.19)	\$ 0.04		Observations
22	\$ (29.44)	Sep-10	0.3	59.90	\$ 49.00	\$ 71.91	\$ (22.91)	\$ 525.06		

Now this is not a perfect science because rain and temperature have a high correlation with our data, but this is a great exercise.

Let's name cell a1:

AVERAGE_OF_TREND_LINE_FOR_MULTIPLE_REGRESSION

Now for each cell let's take the average minus what we have in column f, and we will square it and complete the entire column. Then at the bottom we will add up the Sum of Squares for the regression line:

	A	B	C	D	E	F	G	H	I	J
19	\$ 174.63	May-10	0.6	52.70	\$ 51.00	\$ 55.69	\$ (4.69)	\$ 21.97	R Square	0.84669397
20	\$ 743.82	Jun-10	0.1	58.10	\$ 84.00	69.75	\$ 14.25	\$ 203.20	Adjusted R Square	0.8374027
21	\$ 1,297.96	Jul-10	0	62.20	\$ 89.00	78.50	\$ 10.50	\$ 110.26	Standard Error	12.5913277
22	\$ 1,348.22	Aug-10	0.1	62.90	\$ 79.00	79.19	\$ (0.19)	\$ 0.04	Observations	36
23	\$ 866.83	Sep-10	0.3	59.90	\$ 49.00	71.91	\$ (22.91)	\$ 525.06		
24	\$ 169.82	Oct-10	1.2	54.70	\$ 52.00	55.50	\$ (3.50)	\$ 12.28	ANOVA	
25	\$ 130.61	Nov-10	2.9	48.20	\$ 26.00	31.04	\$ (5.04)	\$ 25.44	df	
26	\$ 1,002.36	Dec-10	4.5	43.50	\$ 12.00	10.81	\$ 1.19	\$ 1.41	Regression	2
27	\$ 1,477.44	Jan-11	5	41.80	\$ 8.00	4.03	\$ 3.97	\$ 15.72	Residual	33
28	\$ 2,017.18	Feb-11	5.8	41.30	\$ 3.00	(2.44)	\$ 5.44	\$ 29.60	Total	35
29	\$ 928.27	Mar-11	4.3	43.40	\$ 16.00	11.99	\$ 5.01	\$ 16.09		
30	\$ 31.90	Apr-11	5.8	47.30	\$ 47.00	36.82	\$ 10.18	\$ 103.55		
31	\$ 174.63	May-11	0.6	52.70	\$ 57.00	55.69	\$ 1.31	\$ 1.72	Intercept	-43.8946188
32	\$ 743.82	Jun-11	0.1	58.10	\$ 98.00	69.75	\$ 28.25	\$ 798.33	Rain	-6.86452811
33	\$ 1,297.96	Jul-11	0	62.20	\$ 99.00	78.50	\$ 20.50	\$ 420.27	Temperature in F	1.96775002
34	\$ 1,348.22	Aug-11	0.1	62.90	\$ 96.00	79.19	\$ 16.81	\$ 282.56		
35	\$ 866.83	Sep-11	0.3	59.90	\$ 52.00	71.91	\$ (19.91)	\$ 396.58		
36	\$ 169.82	Oct-11	1.2	54.70	\$ 45.00	55.50	\$ (10.50)	\$ 110.33		
37	\$ 146.77	Nov-11	3	48.20	\$ 17.00	30.36	\$ (13.36)	\$ 178.42		
38	\$ 1,046.30	Dec-11	4.6	43.50	\$ 9.00	10.13	\$ (1.13)	\$ 1.27		
39	\$ 28,895.10						SUM=	\$ 5,231.87		

The total sum of squares for the regression line is \$28,894.99 which is also what we see here in cell K26:

	E	F	G	H	I	J	K	L	M	N	O
19	\$ 51.00	\$ 55.69	\$ (4.69)	\$ 21.97	R Square	0.84669397					
20	\$ 84.00	\$ 69.75	\$ 14.25	\$ 203.20	Adjusted R Square	0.8374027					
21	\$ 89.00	\$ 78.50	\$ 10.50	\$ 110.26	Standard Error	12.5913277					
22	\$ 79.00	\$ 79.19	\$ (0.19)	\$ 0.04	Observations	36					
23	\$ 49.00	\$ 71.91	\$ (22.91)	\$ 525.06							
24	\$ 52.00	\$ 55.50	\$ (3.50)	\$ 12.28	ANOVA						
25	\$ 26.00	\$ 31.04	\$ (5.04)	\$ 25.44	df						
26	\$ 12.00	\$ 10.81	\$ 1.19	\$ 1.41	Regression	2	28895.1016	14447.5508	91.1278615	3.6453E-14	
27	\$ 8.00	\$ 4.03	\$ 3.97	\$ 15.72	Residual	33	5231.87057	158.541533			
28	\$ 3.00	\$ (2.44)	\$ 5.44	\$ 29.60	Total	35	34126.9722				
29	\$ 16.00	\$ 11.99	\$ 4.01	\$ 16.09							
30	\$ 47.00	\$ 36.82	\$ 10.18	\$ 103.55							
31	\$ 57.00	\$ 55.69	\$ 1.31	\$ 1.72	Intercept						
32	\$ 98.00	\$ 69.75	\$ 28.25	\$ 798.33	Rain						
33	\$ 99.00	\$ 78.50	\$ 20.50	\$ 420.27	Temperature in F						
34	\$ 96.00	\$ 79.19	\$ 16.81	\$ 282.56							
35	\$ 52.00	\$ 71.91	\$ (19.91)	\$ 396.58							
36	\$ 45.00	\$ 55.50	\$ (10.50)	\$ 110.33							
37	\$ 17.00	\$ 30.36	\$ (13.36)	\$ 178.42							
38	\$ 9.00	\$ 10.13	\$ (1.13)	\$ 1.27							
39					SUM=	\$ 5,231.87					

So, if we take k26 which is the regression Sum of Squares divided by k28 which is the total Sum of Squares we get 0.8466 which is what we see up here for r squared in J19. So, this means that our Multiple Regression formula is 85% accurate and 15% not accurate.

Is our data statistically significant? Yes it is because the significance number in N26 is way below 0.05

Let's move onto the P-value column starting in cell M30. The P value tells us if this line item is statistically significant. Similar to the F value, we want this to be less than 0.05....and we see this first item here for the Y intercept in cell M31 is way higher than 0.05 but for the x1 for rain in cell M32 it is less than 0.05 and for the temperature in cell M33 it is less than 0.05.

So, this means that our model does a good job with x1 and x2 but not when the intercept is zero as the equation implies that when the temperature and the rain are both 0 then sales will be massively negative...remember we only look at y values when x1 and x2 are within the minimum and maximum levels of the 36 data points we have.

In terms of the T stat in cell L30 it helps us get the P value. Thank goodness Excel calculates all this for us, so you don't have to worry about this T stat column.

But again, for our purposes just assume that any P value under 0.05, means the variable is significant. Again, don't spend too much time on the intercept line for statistics.

So, our model tells us that x1 and x2 are statistically significant because the p value is less than 0.05 for both x1 and x2...you'll find that as you add more and more variables to your model like x3 and x4 and x5 etc, you will find that one or more of the data points might have a p value here above 0.05 and if so, then those variables don't have a high correlation to you your model

[A] Fixing Problems with Data

In the linear and Multiple Regression Analysis chapters we spoke about the importance of using data with low significance F and P-values, meaning values below 0.05 so that the models are statistically significant.

What else can we look for in our data before we analyze the data?

First of all, we want to make sure that we have at least 10 observations or data points per variable. In the rainfall versus sales Regression Analysis that we did in this section, we had 36 data points and with the nesting doll Regression Analysis we had 21 data points.

Secondly if we have a few extreme outliers in our data, then we can consider removing the data points and running the regression again.

Third make sure that the sample data we are using is representative of the entire population of data. For example, it wouldn't make sense to analyze McDonald's restaurants in just one town if we want to use the sample data to help us understand global sales predictions. Instead, we would want to use data from multiple geographic locations.

We can always back test our data to see if certain sample data is representative of the total population of data. For example, in the United States politicians and journalists love to analyze how the state of Ohio votes in federal elections, as it has often had a very high correlation on whether the Democrats or Republicans win the election.

SECTION 7: IMAGES AND AI **(DALL-E, MIDJOURNEY)**

[A/B] Introduction to Images & AI

Welcome to section 7, where we will discuss how to create incredible images using Artificial Intelligence. Image generation using AI has been so disruptive, that many graphic artists that used to charge more than \$40,000 to make video game images have seen their business fall off a cliff.

In terms of the What, and Why and how of Section 7, let's kick it off with the what. What are we going to learn?

We will learn how to use top AI image products like DALL-E and Midjourney.

Moving onto why. Why are we going to discuss using AI to make images? Because as a society we are much more visual now, which is why Meta bought Instagram and why products like Pinterest exist.

If we learn how to quickly and easily make high quality images using AI, then we can take our personal brand or corporate brand to the next level with social media and we can make much more impactful our documents, presentations, website designs and much more.

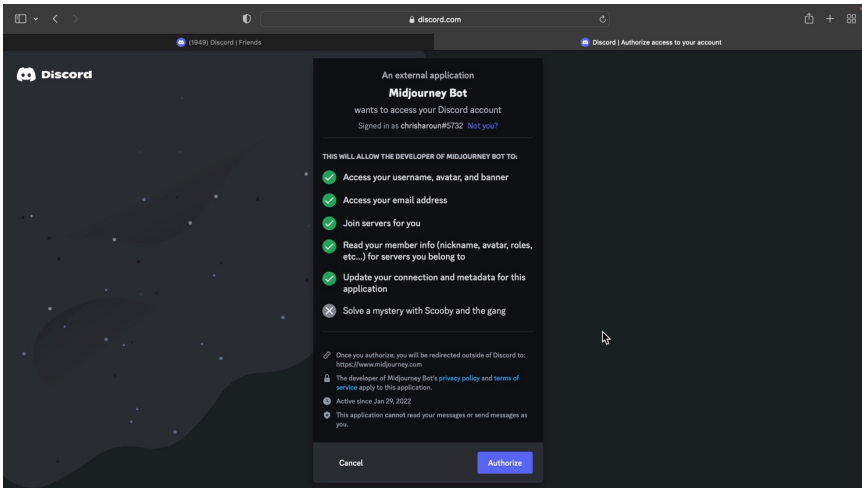
[A/B] Midjourney Account Setup

Midjourney is an amazing tool that is used to generate new images. It's a bit difficult to understand because they use Discord as a platform (Discord is similar to Slack).

First of all, either download Discord and open an account, or use Discord in your browser at <https://discord.com/>.

Please then go to Midjourney, per this link: <https://www.Midjourney.com/>, and sign in or create an account.

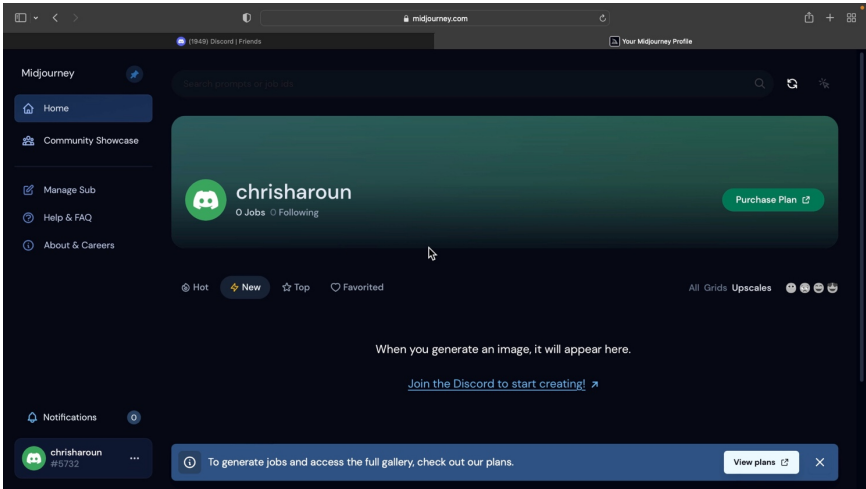
Then, per the image below, click log sign in and authorize Midjourney to use your Discord account:



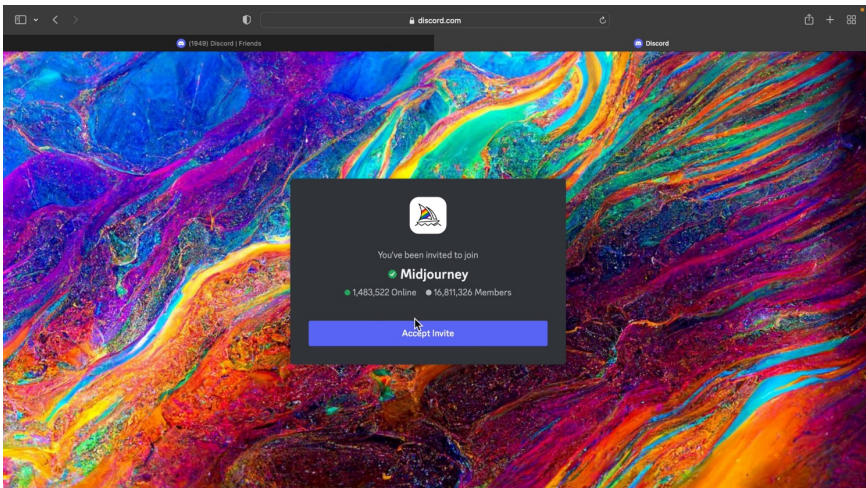
Midjourney has its own website where you can only see the images that you've generated so far, and you can see the community showcase, which allows you to see what other members of the community generated themselves.

When you create images, they will also appear on Midjourney's website after you log in.

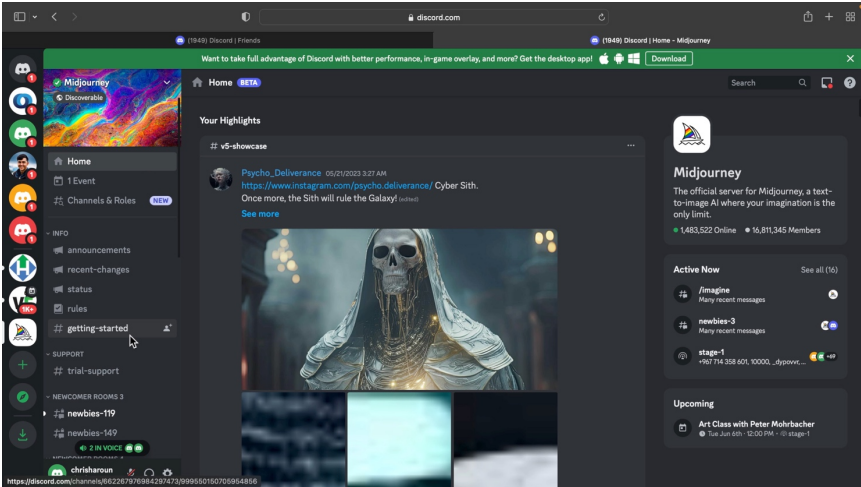
Per the image below, when you log into your Midjourney account on Midjourney's website, you can click the link called Join the Discord to start creating link:



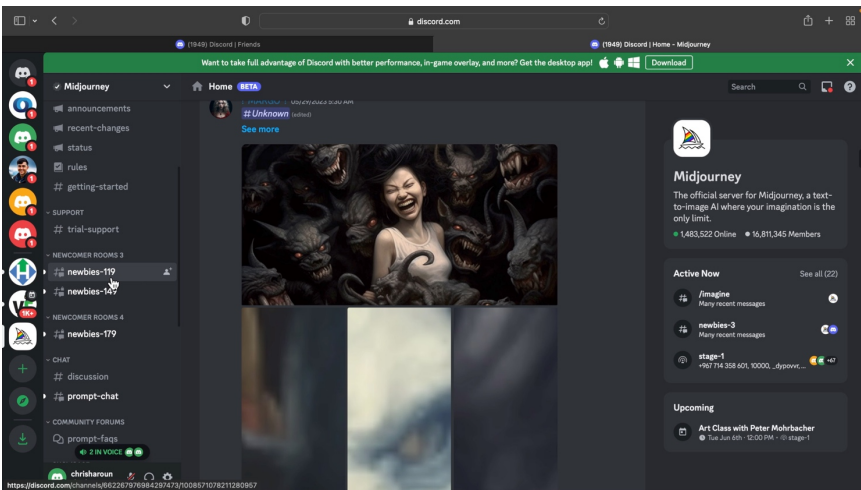
Then, per the image below, accept the invite when you sing into Discord:



Then, per the image below, you will have a new server on your Discord account, which is this image of Midjourney:



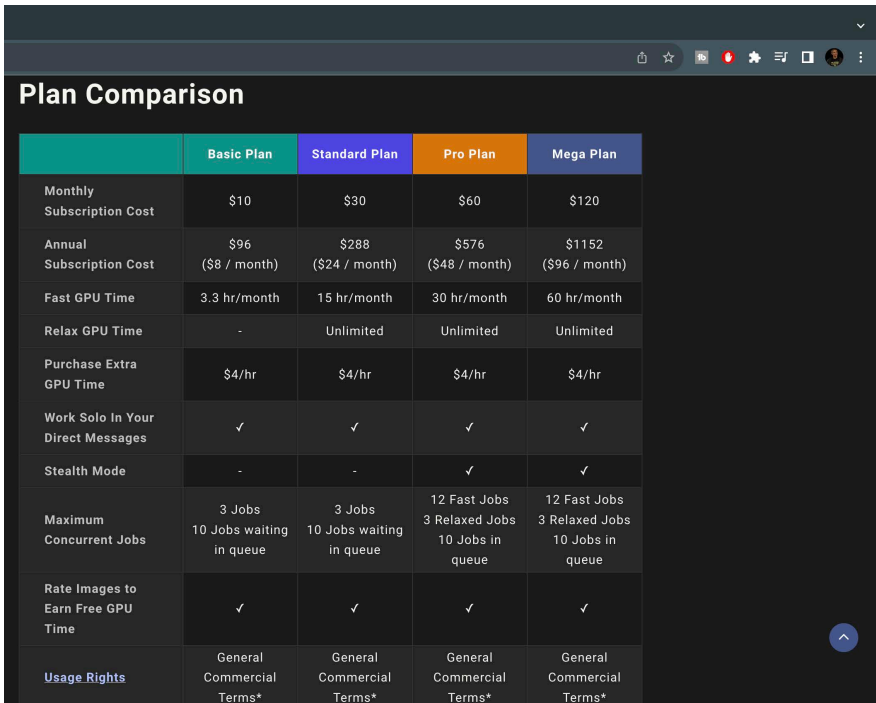
You will see first the showcase of what other people generated. Now, in the support channels on the left side you'll see that couple of newcomer rooms, which are called newbies. For example, in this case it's 119 for us:



In your case, it might be another image or another number, because every single time, they're creating more and more channels for newbies.

Now we can start prompting to create images, which we will discuss in the next chapter. Sometimes you need to wait for a couple of minutes to get to your order of having your prompt taken by their engine.

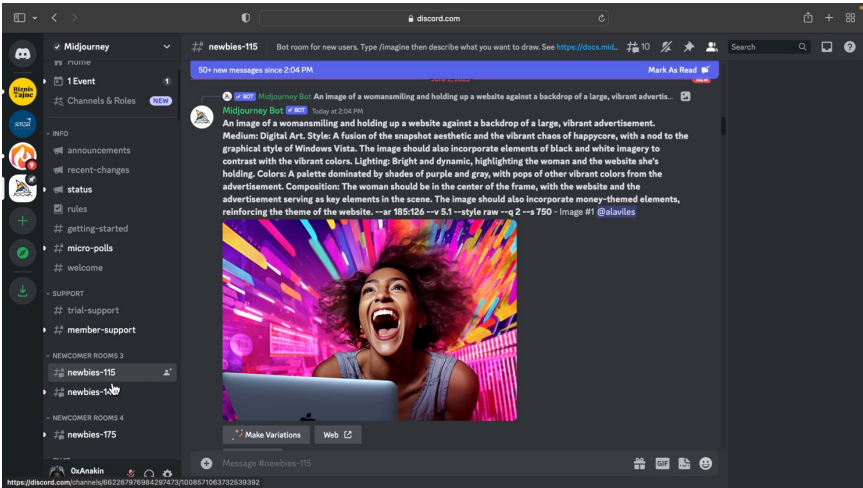
Below is the pricing table



	Basic Plan	Standard Plan	Pro Plan	Mega Plan
Monthly Subscription Cost	\$10	\$30	\$60	\$120
Annual Subscription Cost	\$96 (\$8 / month)	\$288 (\$24 / month)	\$576 (\$48 / month)	\$1152 (\$96 / month)
Fast GPU Time	3.3 hr/month	15 hr/month	30 hr/month	60 hr/month
Relax GPU Time	-	Unlimited	Unlimited	Unlimited
Purchase Extra GPU Time	\$4/hr	\$4/hr	\$4/hr	\$4/hr
Work Solo In Your Direct Messages	✓	✓	✓	✓
Stealth Mode	-	-	✓	✓
Maximum Concurrent Jobs	3 Jobs 10 Jobs waiting in queue	3 Jobs 10 Jobs waiting in queue	12 Fast Jobs 3 Relaxed Jobs 10 Jobs in queue	12 Fast Jobs 3 Relaxed Jobs 10 Jobs in queue
Rate Images to Earn Free GPU Time	✓	✓	✓	✓
Usage Rights	General Commercial Terms*	General Commercial Terms*	General Commercial Terms*	General Commercial Terms*

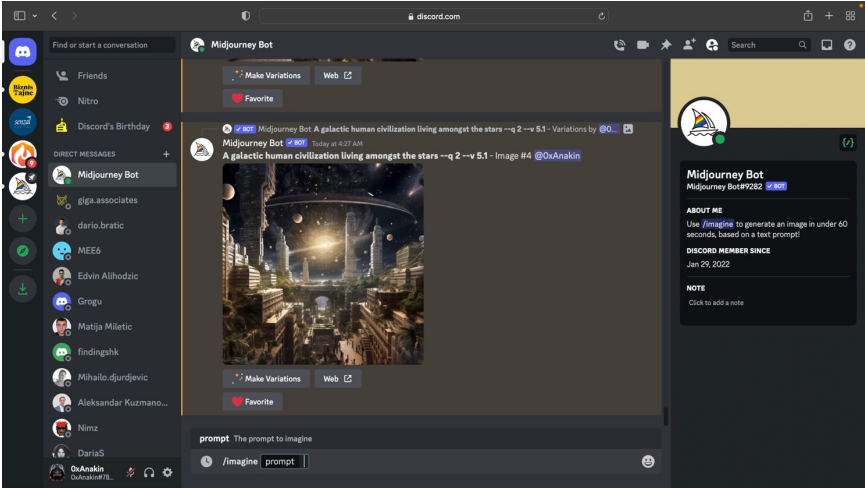
***[A/B] Midjourney Interface Walkthrough and
How to Use it***

Let's send the first request and see how the generation goes. Now I'll generate it through the normal channels to see like what is the difference between newcomers and the private generation. So, per the image below, I'll jump to this newbies 115 where I'm the member of:



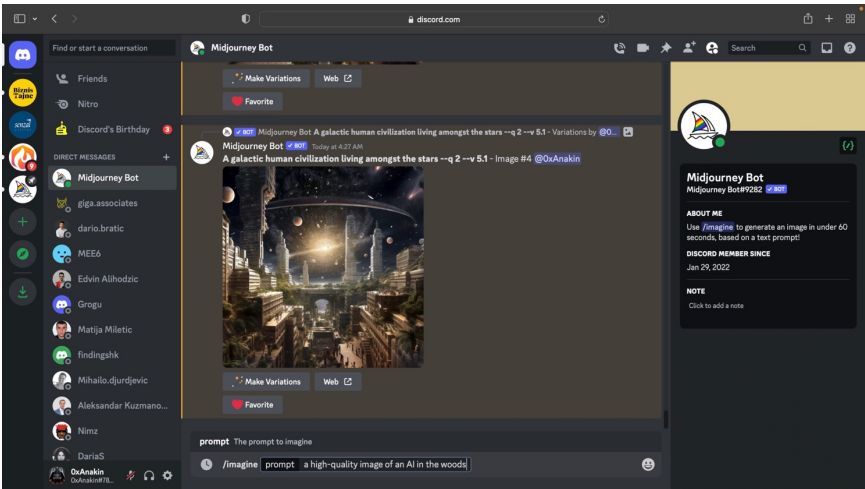
You can scroll through this channel to see what prompts other users entered to make their images We'll learn all the techniques necessary to build great images.

To generate the image, we need to type / and then you will see the first comment which is image imagine:

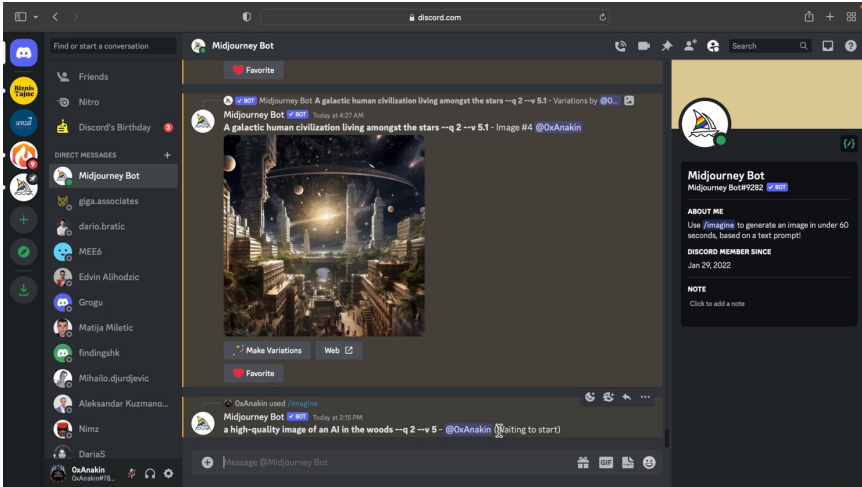


And then you are going to write the prompt. Our prompt, per the image below is:

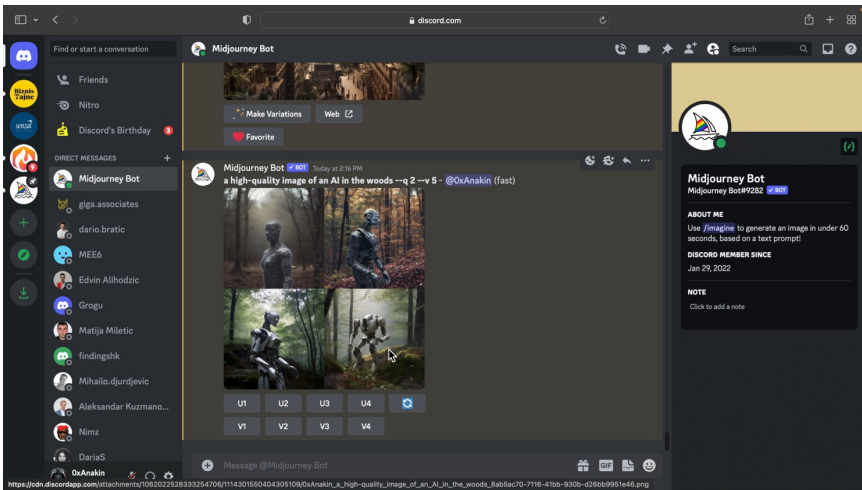
a high quality image of AI in the woods



And as you can see per the image below is the first problem that so many people are actually sending the requests and I'll need to wait for quite some time.

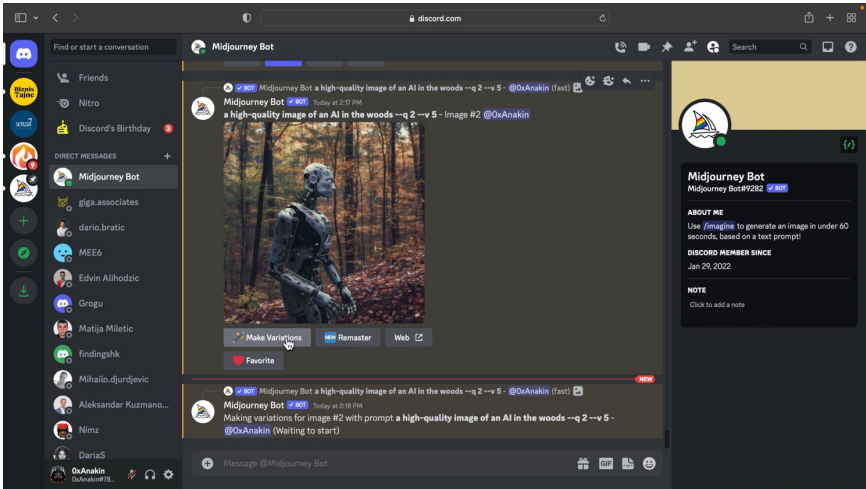


Here is the image that was generated:



We got 4 versions of the image, and you see U1, U2, U3, and U4 and V1, V2, V3 and V4. So, by default, every time you send a request Midjourney will give you four different versions of the image that you are generated. And when you do to upscale the images by selecting U1 for upscaling of image 1, for example.

Per the image below, if I click on U2 and V2, I'll get the bigger version of the 2nd image:

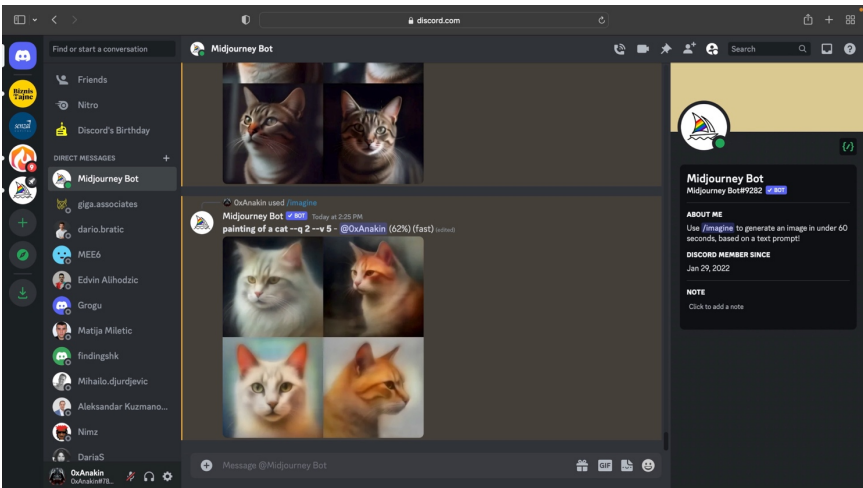
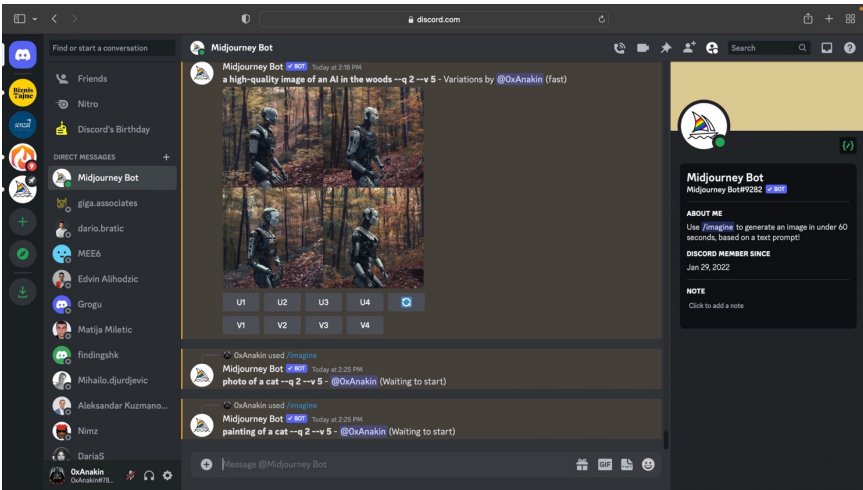


[A/B] Midjourney Good Prompting Practices

Let's explore general prompting strategies on how to prompt Midjourney in a much better and consistent way. If you want a photo and a painting of a cat, you could write these 2 prompts and submit them per the images below:

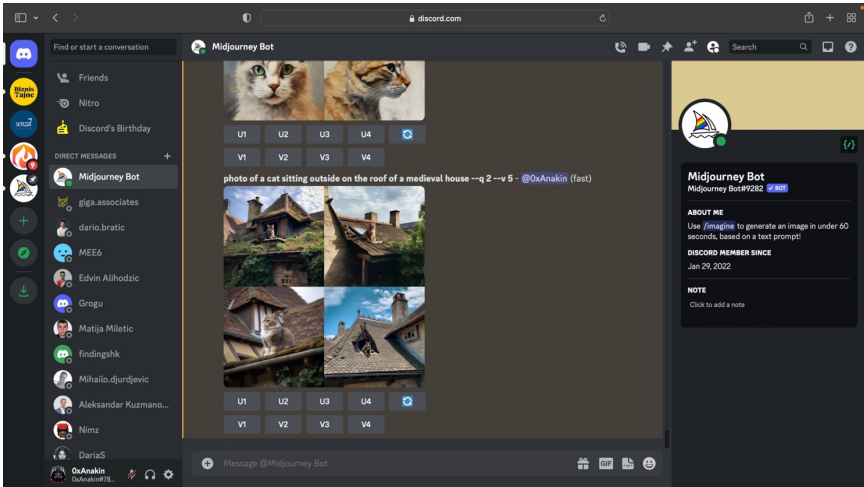
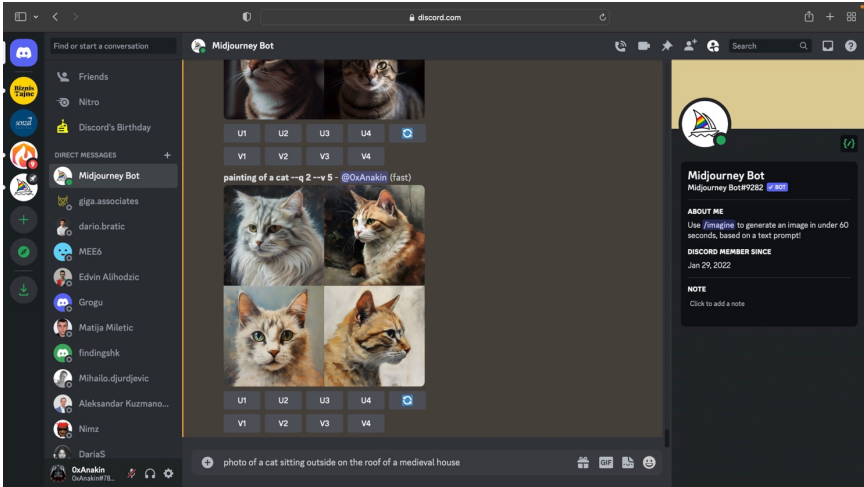
imagine a photo of a cat

imagine painting of a cat

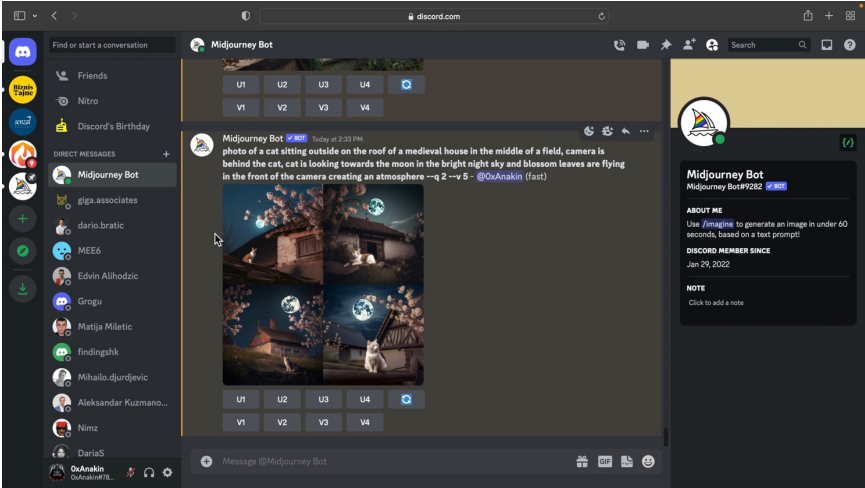


Let's add more detail to the prompt by typing:

Photo of a cat sitting on the roof of a medieval house



Let's change this prompt to something much more detailed as shown below:

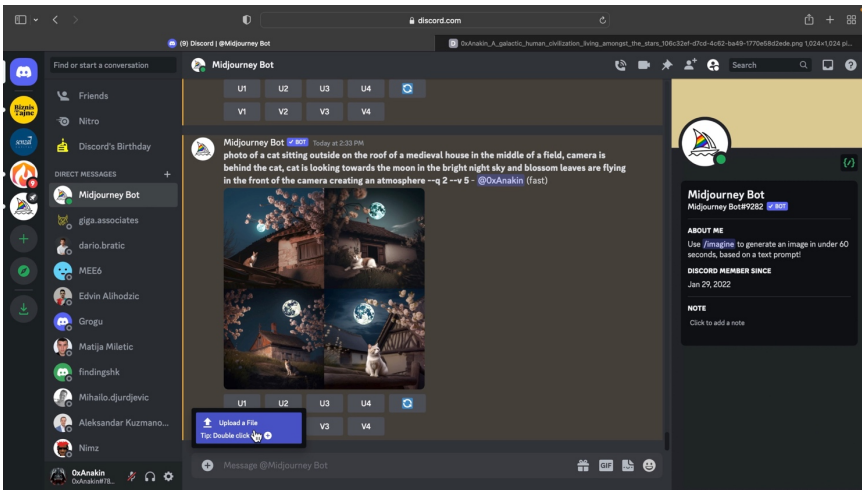


So, it's all about giving more and more information on how to shape the output of the model.

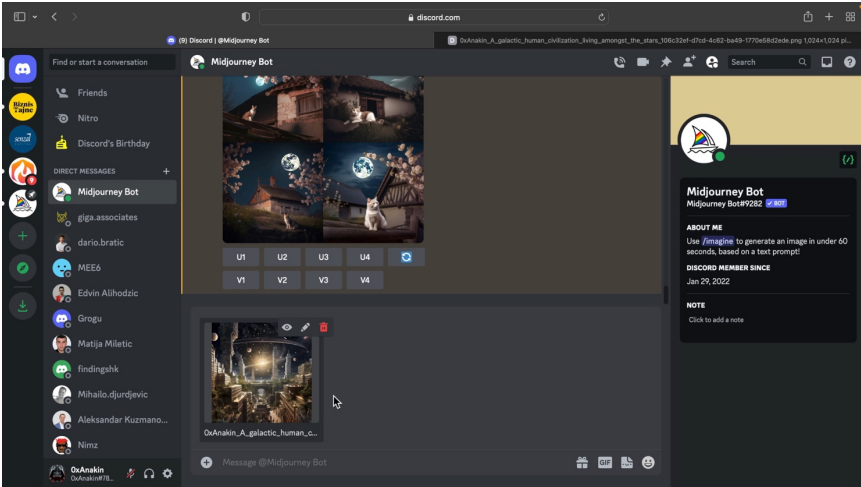
[A/B] Midjourney Image Prompting

So far we used text to describe what we want to achieve or generate, but you can do the same with images. Imagine this, just without words you can upload an image and ask Midjourney to generate based on this image. And that's what we are going to do in this chapter.

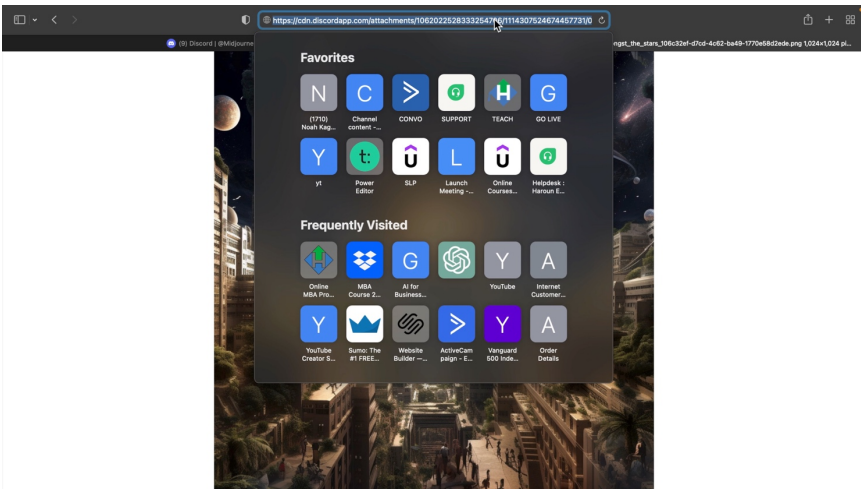
Per the image below, you can go and click on this plus button, and upload a file:



I'm going to use an image of a previously generated prompt based on Midjourney:

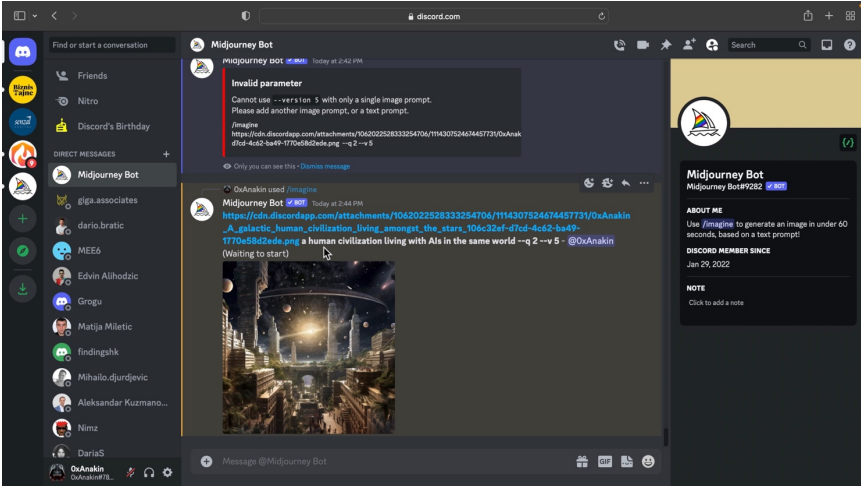


So now that we have this image uploaded, click on it, and you can open it in a browser, and you will get this link. Copy this link, and you can go back to your Midjourney:

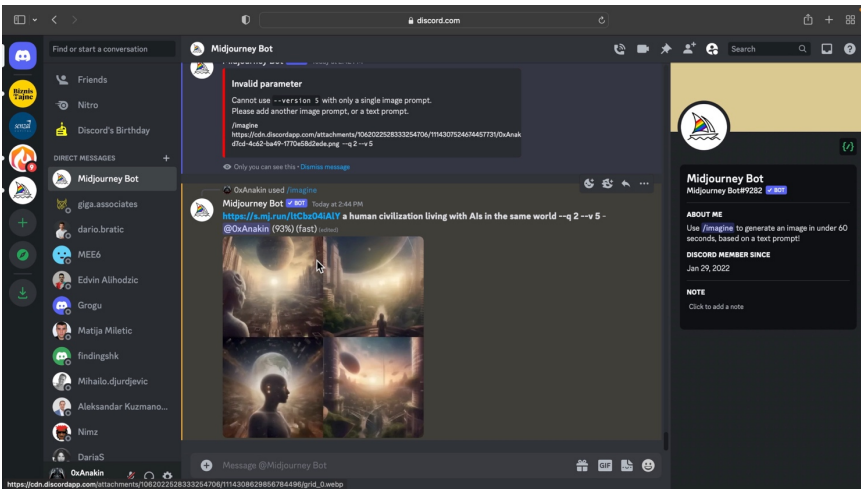


If I just prompt it with the link to the image, I'll get an error.

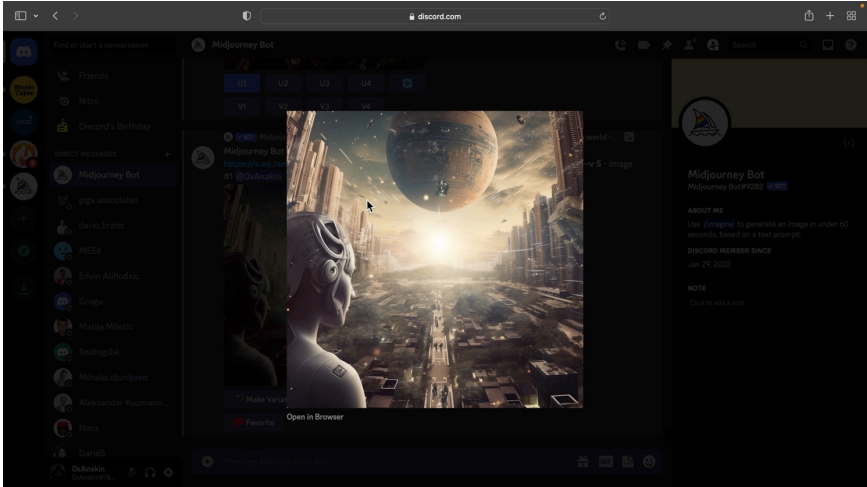
So, we can type our prompt to use the pasted link and then add --v5, which is the version of the model Midjourney and some additional items to alter this image:



So, what that means is, I'm basically going to use this image as a starting prompt and I'm going to add text. So, it's going to add context to this image. So, if I do this, it's going to start prompting it and it's going to start and use the image and it's going to use the tags that I just prompted it out per the image result below:



Let's upscale the first image, which is my favorite from this, and see how it actually looks in the full scale:



So, this is how com to combine like image prompting and text prompting as well.

***[A/B] DALL-E 2: Account Setup and Interface
Walkthrough***

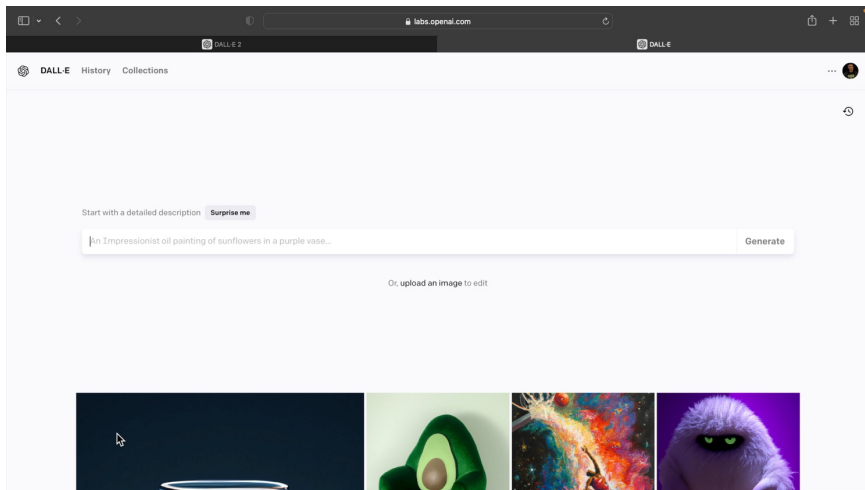
DALL-E is another imaging product, that is easier to use than Midjourney. The company who developed DALL-E is OpenAI, and we talked about ChatGPT.

If you go to the following link, you can click on try DALL-E: <https://openai.com/dall-e-2>

Unfortunately, DALL-E is not free by default. So, you can't search, you can't do much with it if you just go and log in. You have to pay to use DALL-E.

So that's a bit a downside of this, but in my opinion, it can get even better results than what we did so far with Midjourney. The minimum amount to pay is \$15 for 100 credits, and each image generation is about 1 one credit.


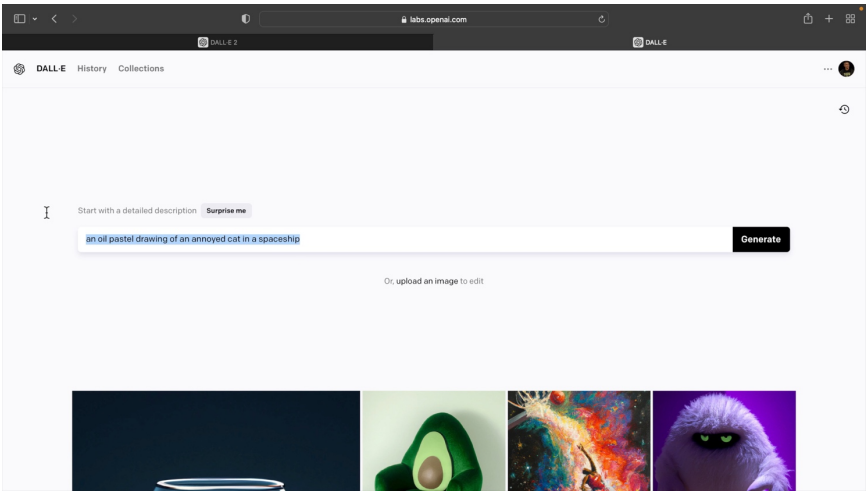
Let's talk about the interface. Per the image below you have a text box where you are going to write your prompts.



On the DALL-E website, you can scroll down and see examples of other users images, and you can see what prompt was used to

generate them. Simply click on an image and you can get the prompt used to make the image.

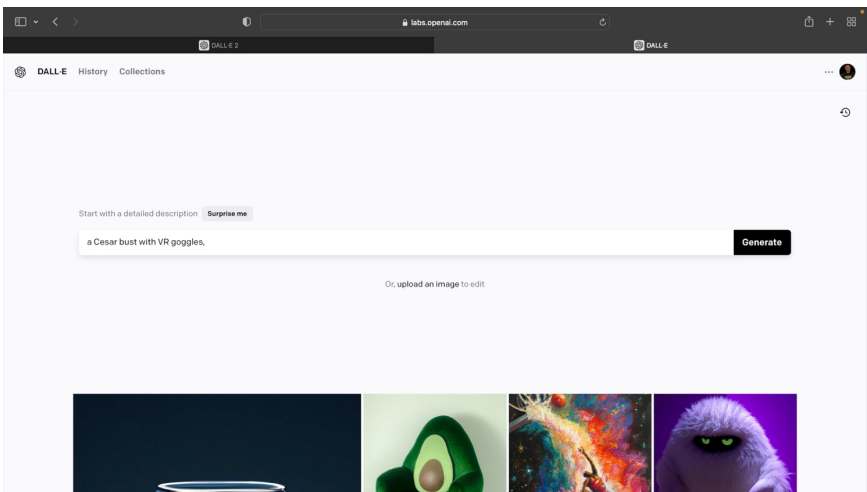
Below is a prompt we used to create an image of a cat and the resulting image generated.



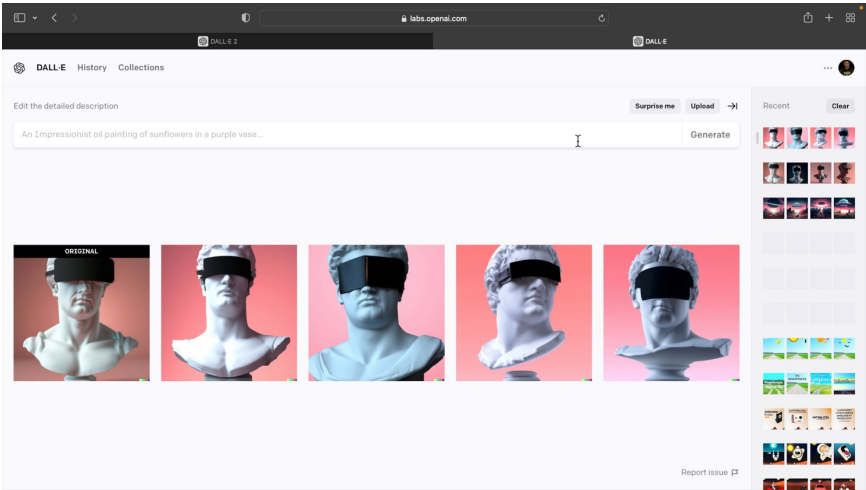
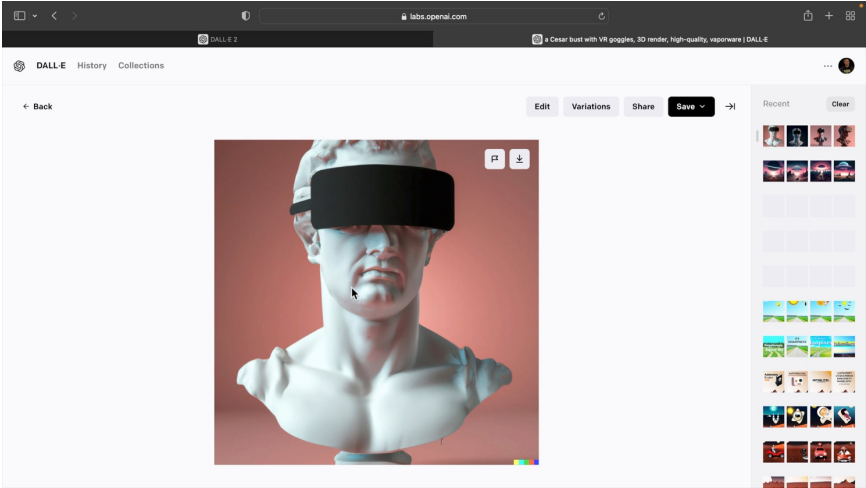
[A/B] DALL-E 2: Prompt Designs

Let's talk more about general prompting or how to prompt. Unlike Midjourney, with DALL-E, you don't really need to learn much about prompting in order to create great results. And that's what DALL-E is used more in some cases than Midjourney because for beginners, DALL-E is much more suitable and sufficient to generate great results with the simple prompts.

Per the images below, I created an image of a Caesar bust with VR goggles.



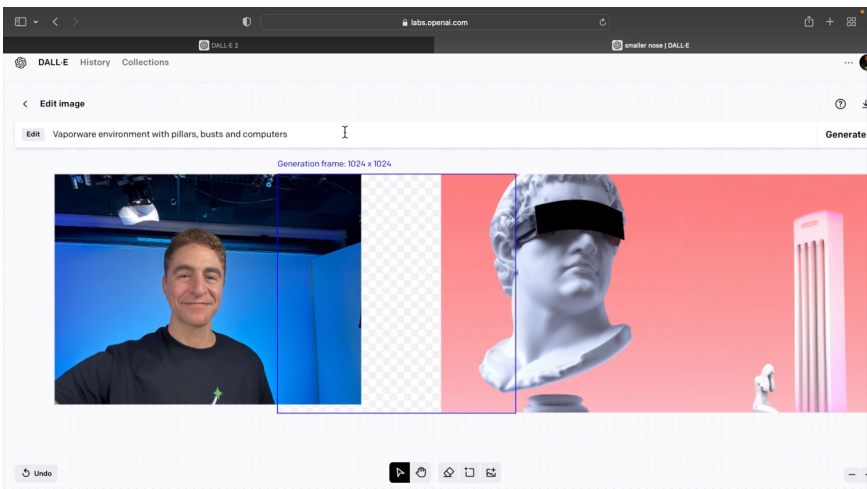
You can also create variations of this image. Per the images below, If I click on the first image and then select the variations button, I can see many more variations of the image:

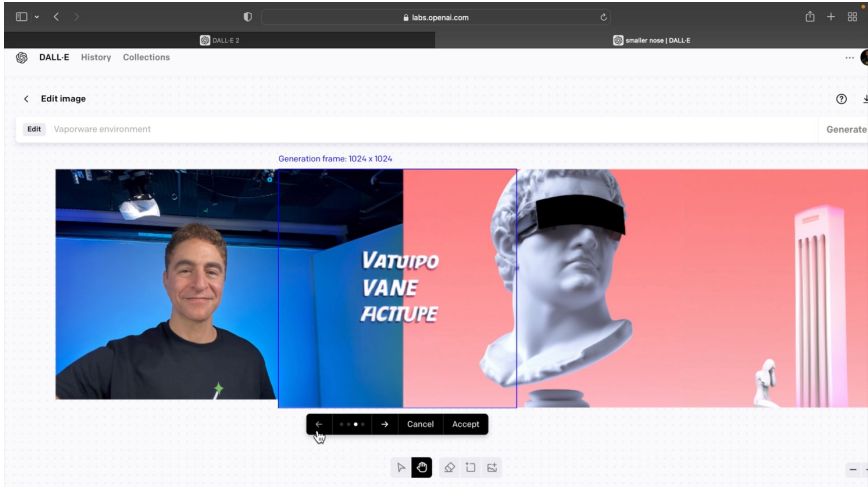


***[A] DALL-E 2: Editing Images, Inpainting &
Outpainting***

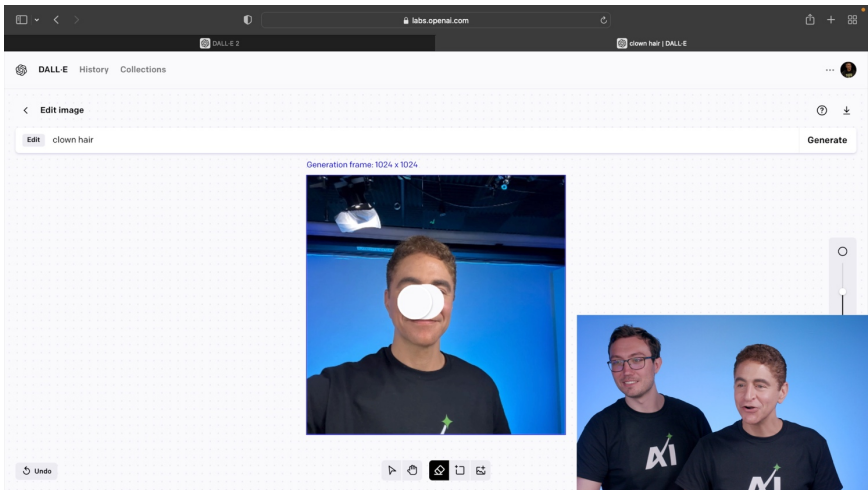
One thing that is amazing about DALL-E is the concept of Outpainting and Inpainting. Outpainting is a process of creating context around the image, so basically, expanding the size or sides of the image. One cool feature of this is the possibility to connect multiple images as well.

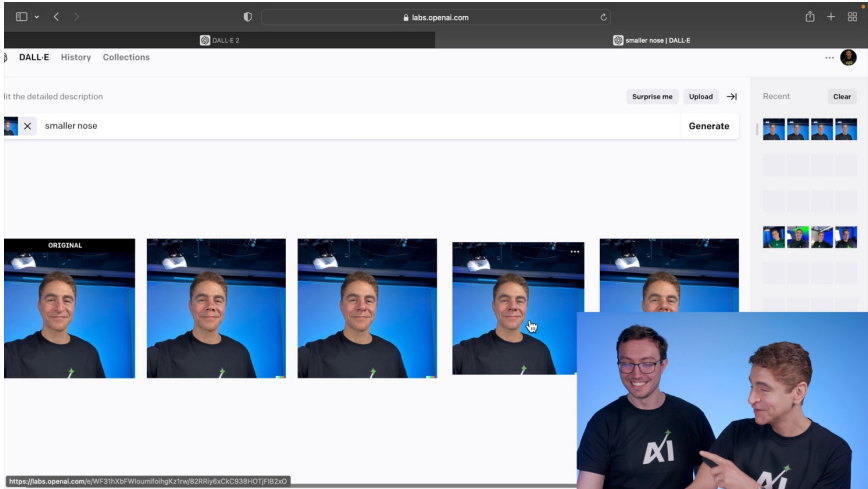
Below is an example of Outpainting, where we take a picture of Chris and a picture of Caesar, and we select the area between the images, and we connect them using Outpainting.





Inpainting is when you change what is inside the image like changing Chris' nose per the images below :)





SECTION 8: SECTION 8: VOICE, AVATARS AND CLONING

[A/B] Introduction to Voice, Avatars & Cloning

Welcome to section 8. In terms of the What and Why and how of Section 8 let's kick it off with the what. What are we going to learn?

We will discuss useful AI tools that can help you with creating voice clones of yourself and tools that can help you with creating digital avatars.

Moving onto why. Why are we going to discuss using voice and avatars and digital cloning? Because generating voices that sound almost exactly as we do can be very useful in the business environment.

We could generate samples of our voices and more and use them to narrate marketing campaigns, product reviews and even create tutorials! We can even use these technologies for customer support on the phone or on the Internet.

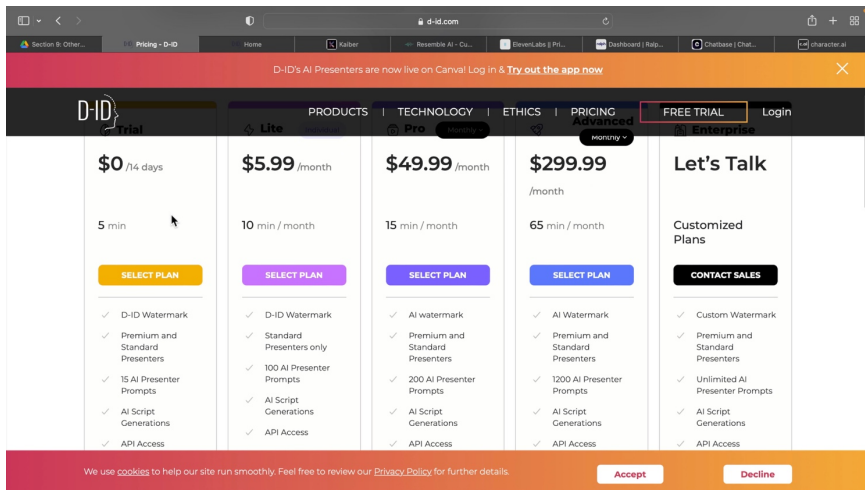
As dangerous as these tools can be, from an ethical standpoint, they are going to be used a lot in the future. Movies will be made with lifelike avatars, and we can create our own lifelike avatars in the future to help us create promotional trailers or to advertise our companies.

[A/B] Using/Making AI Avatars with D-ID

We are going to review couple of products that can help you create exciting materials or even create your own promotional marketing campaigns. The first tool that we are going to talk about today is D-ID, which you can access at: <https://www.d-id.com/>

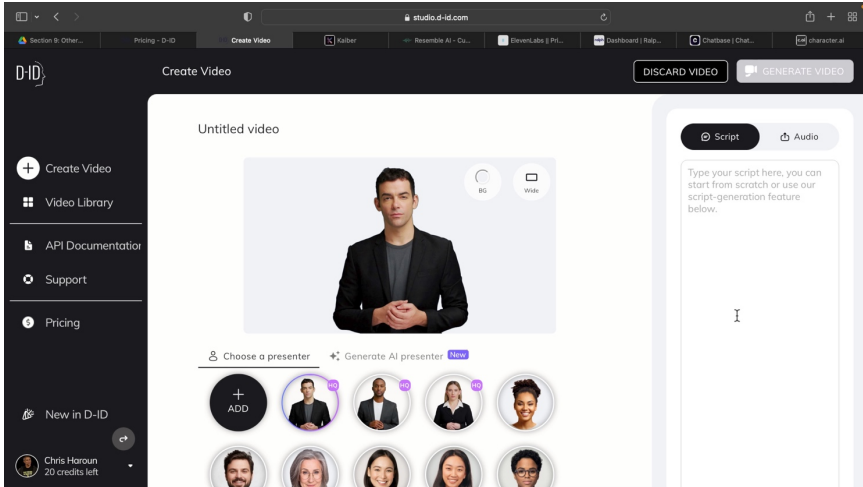
D-ID allows you to create digital avatars. They're not perfect yet. You can upload a simple image of yourself or somebody else. We are going to do that for Chris, and you can write a script, and upload your own voice as well, and make a digital avatar that moves your mouth.

They offer different pricing plans per the image below. For this tutorial, I'm using the free trial.

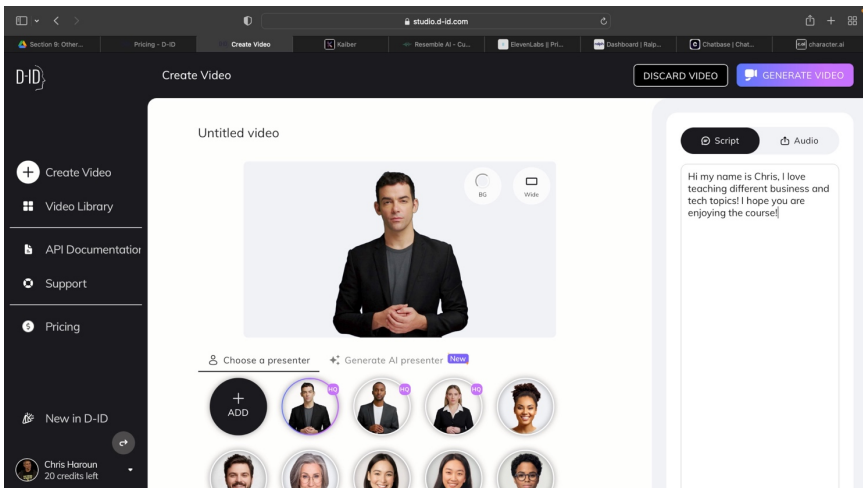


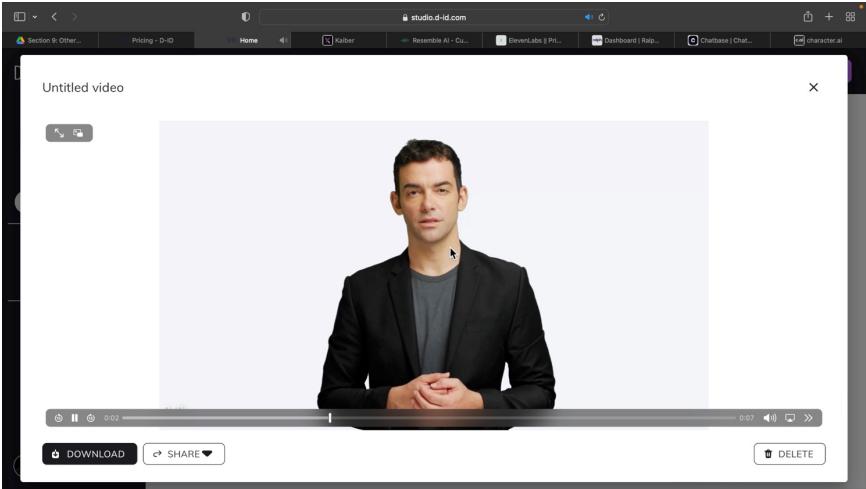
If you pay for the Pro version and above, you don't have the watermark on the video generated.

After you create an account, then you can start generating avatars of yourself. Let's create a video, and per the image below you can even use their avatars:

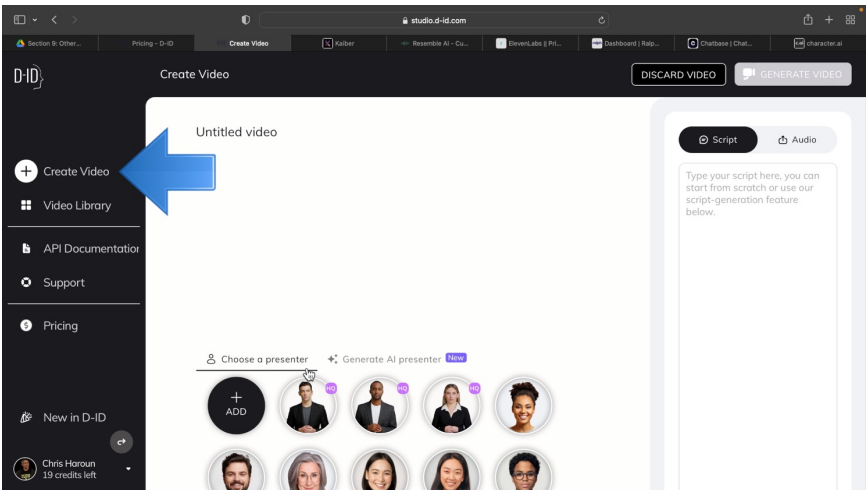


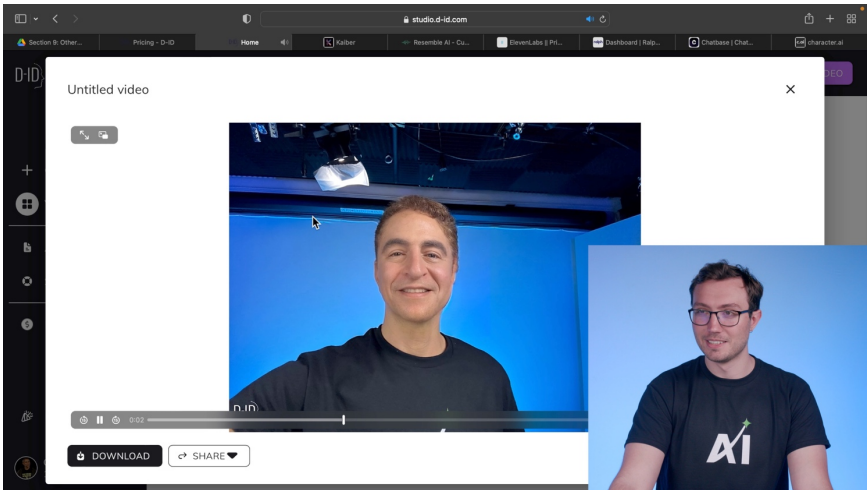
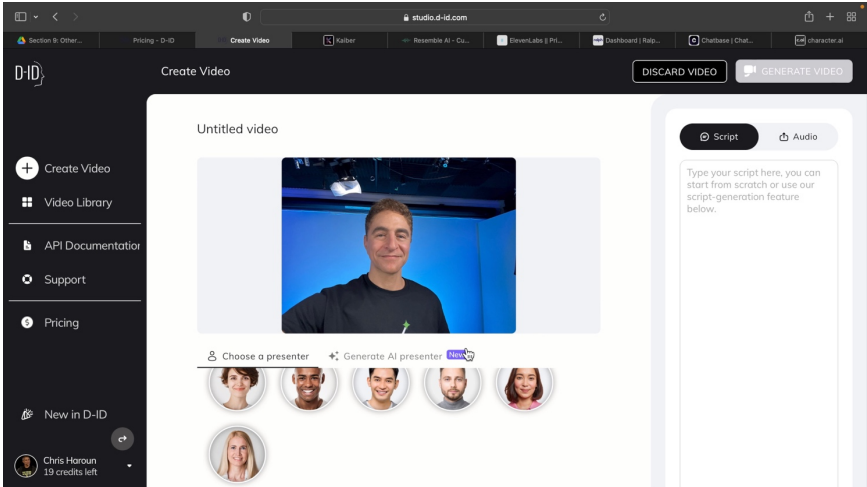
These free avatars here are amazing. They move their hands, they move their face as well. And whatever you type, it's going to be synchronized. So, let's write our script for this video. I'll write, "Hi, my name is Chris. I love teaching different business and tech topics. I hope you are enjoying the course, per the images below:





Per the images below, I uploaded the image of Chris and used the same script:



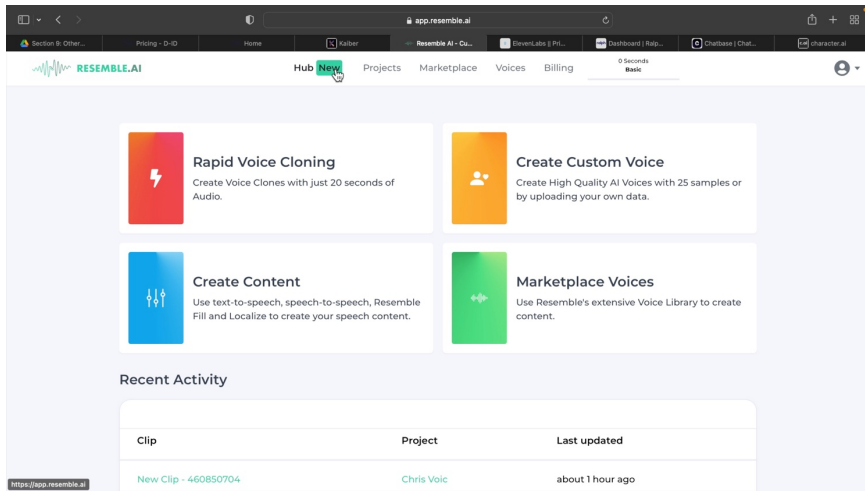


It's not perfect. If you would like to use their own avatars, then you can use them for YouTube videos or other marketing as the quality of the preset avatars is great!

[A/B] Resemble AI

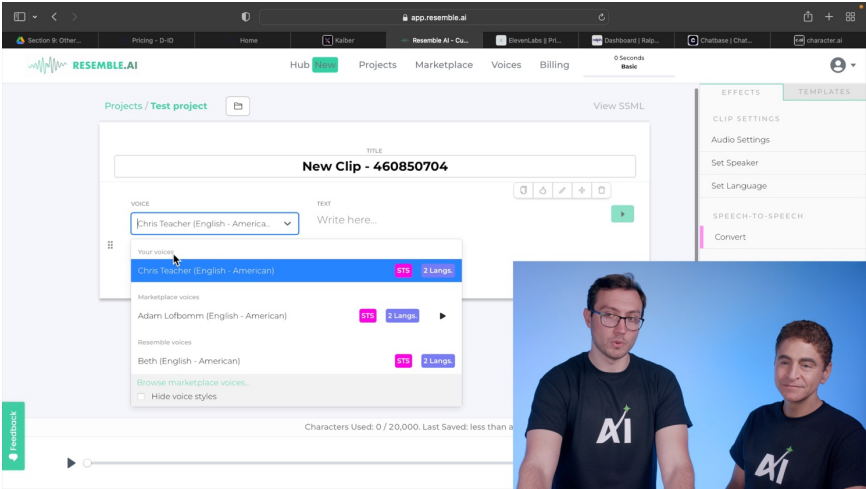
Let's talk about voice clones and Resemble.AI, which can be accessed at <https://app.resemble.ai/>

The idea about this product is that you can create amazing copies of your voice. Below is the user interface after you create an account:



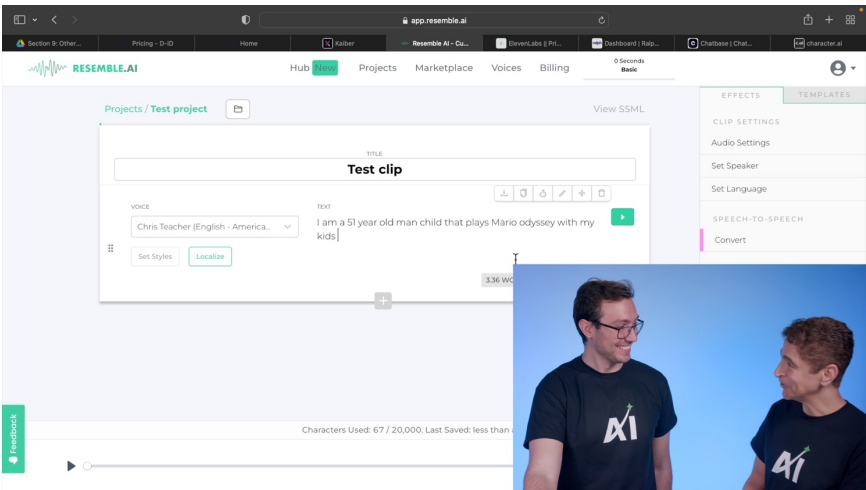
We did subscribe to the basic entry level product, which costs 0.006 per second of use. So, from the interface, you can see that there is a rapid voice cloning, with just 20 seconds of audio. And that works fine. We didn't test it out because I know that creating a custom voice works much better for us.

So, if you go here, you're going either to record yourself speaking or create a clone your voice by just uploading a sample, per the image below:



With the custom voice example, it allows you to read 25 sentences and every single time you read it, you upload it and so on and so on until you go to up to 25 sentences.

Per the image below, we tested a sample written script and it was amazingly accurate (which you can hear in the lecture in the course called: [A/B] Resemble AI).



SECTION 9: OTHER AI APPLICATIONS (RALPH AI, CHATBASE AND MORE)

[A] Introduction to Other AI Applications

Welcome to section 9, where we will discuss additional AI products and applications that we found useful for you and your business.

In terms of the What and Why and how of Section 9, let's kick it off with the what. What are we going to learn? We will learn about Ralph.AI is and its benefits over normal ChatGPT. We will also discuss using Chatbase in order to add a chat feature to your website. And we will discuss Character AI, which generates human like responses.

Moving onto why. Why are we going to discuss the aforementioned additional AI tools? Because we want to introduce you to other AI technologies that leverage existing AI platforms in order to help you work smarter and not harder.

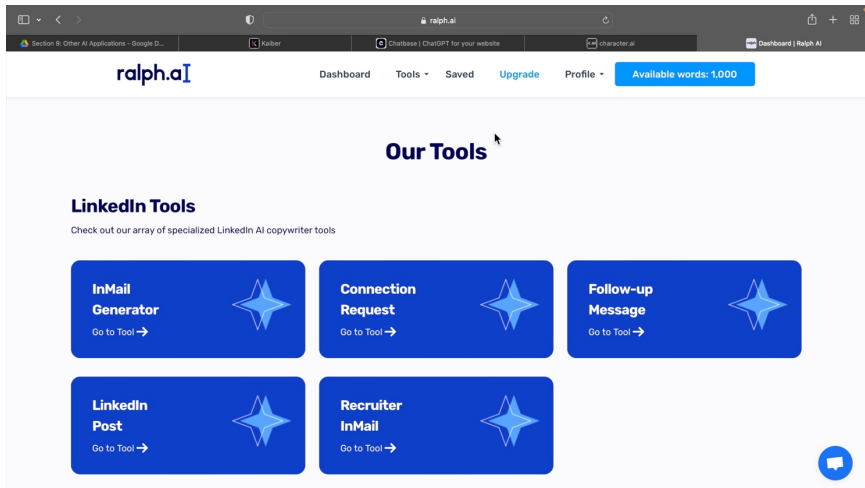
[A] Ralph AI

Have you ever thought about using an even better assistant than ChatGPT in copywriting of your emails, marketing campaigns, LinkedIn posts, and many other things that you're potentially going to build? Well, that's Ralph.ai, (<https://www.ralph.ai/>) the tool that is built around ChatGPT with fine-tuning in mind to help you with these specific tasks.

In this chapter, we'll discuss how to use Ralph.ai to help you generating much better outcomes for your business and potentially, personal use as well.

You don't have to pay to try the product. You're going to get 1,000 words for free. But if you find it useful, make sure that you upgrade to a premium.

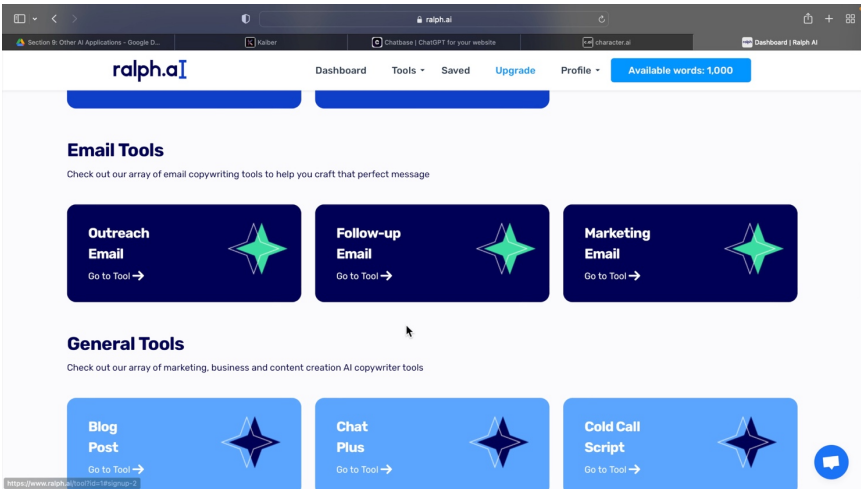
Below is the dashboard where you're going to land once you create an account. You have tools separated into different categories:

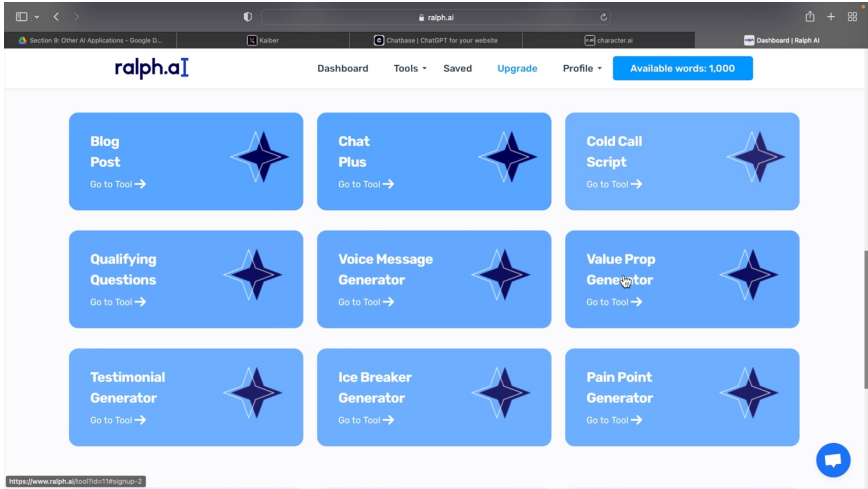


First off, we have the LinkedIn Tools. Amazing for either if you are just brand that you would like to build your personal brand or working for some company and trying to work with them. You have InMail Generator, Connection Request, like when you're sending a connection request over LinkedIn, you can create and post those small messages telling them why would you be perfect to be in their network.

Follow-up Messages is if you are performing some sales. LinkedIn Post and Recruiter are if you're a recruiter to email messages as well.

Per the images below, for emails, you have Outreach, Follow-up, and Marketing if you are Cold emailing. There is also Blog, Chat Plus, which is ChatGPT, Qualifying Questions, Ice Breakers, etc.



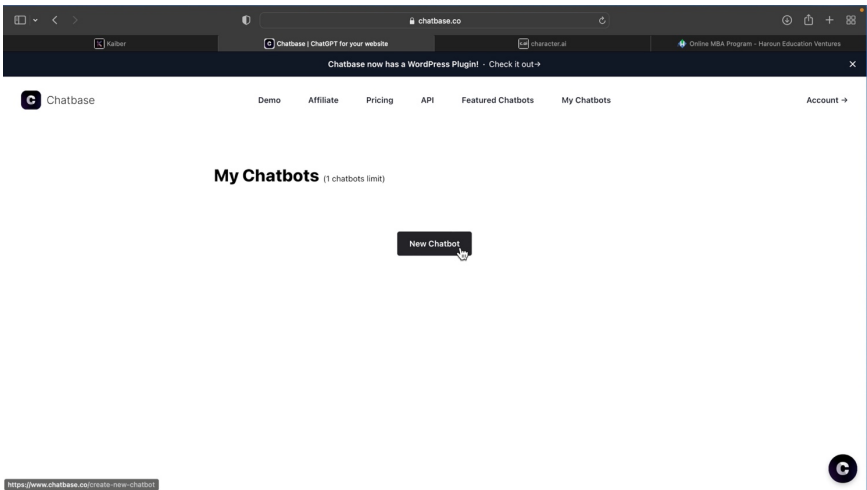


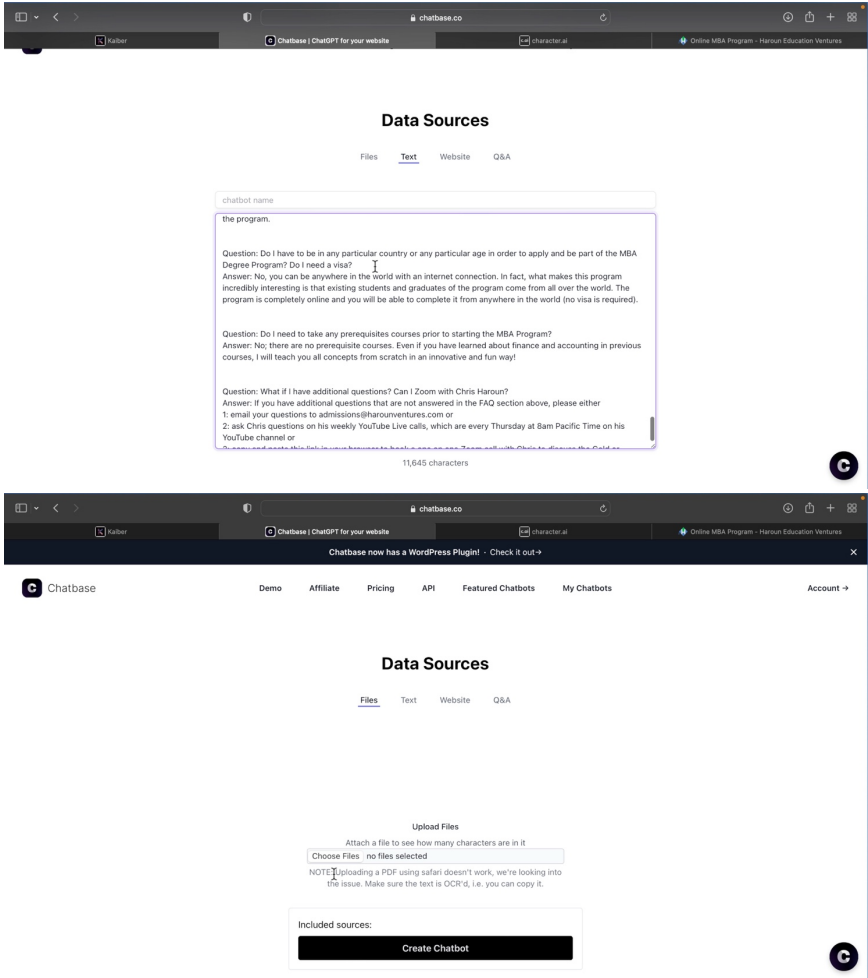
The product is very intuitive to use, and I hope it helps you!

***[A] Chatbase (How to Add Chatbase GPT to
Your Website)***

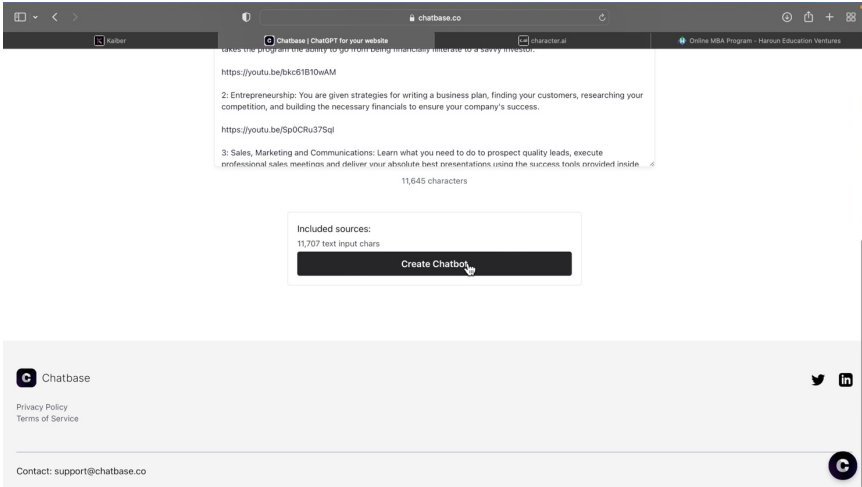
Chatbase (<https://www.Chatbase.co/>) is a great resource to add a ChatGPT-like bot to your website. With a free account you can create one bot.

Per the images below, after you create an account, select the button called New Chatbot, then from Data Sources either upload a file or paste your Q&A or FAQs as we did in the 3rd image below for Chris' website.

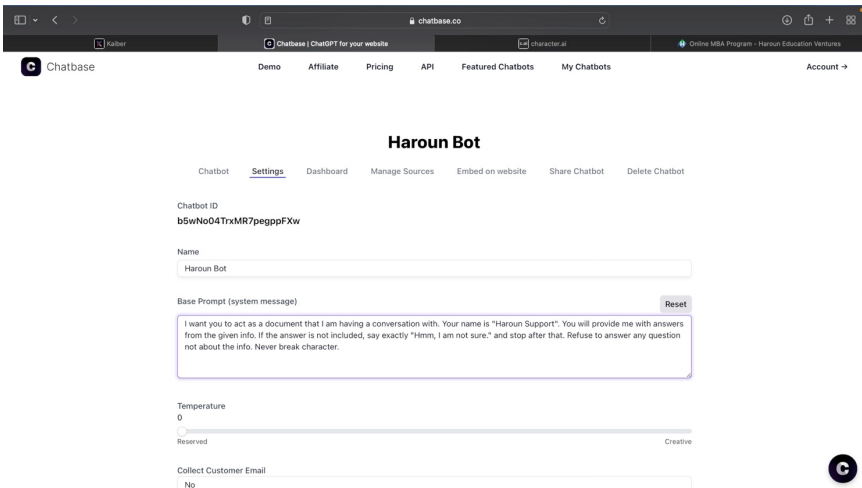




Then per the image below, select Create Chatbot and bam, we have a bot!



Then in Settings you can change what the default prompt is:

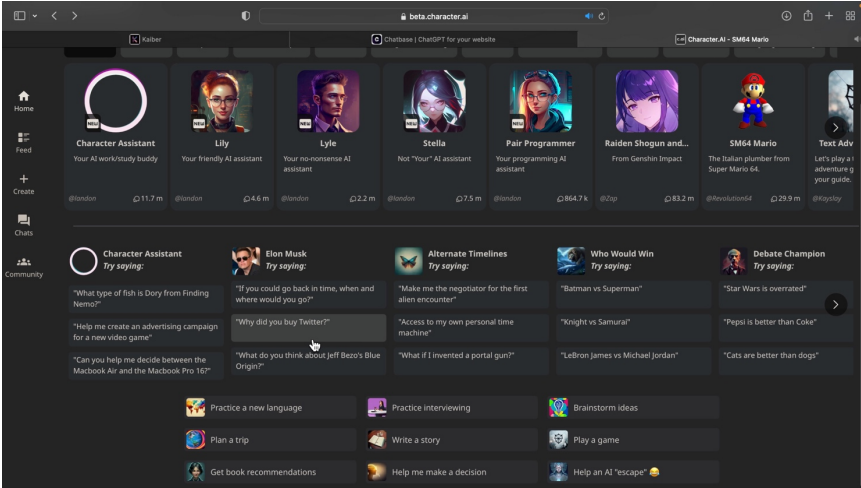


On the same screen, you can copy and paste the code to add it to your website as we explain in the video version of the course in the lecture called: [A] Chatbase (How to Add Chatbase GPT to Your Website).

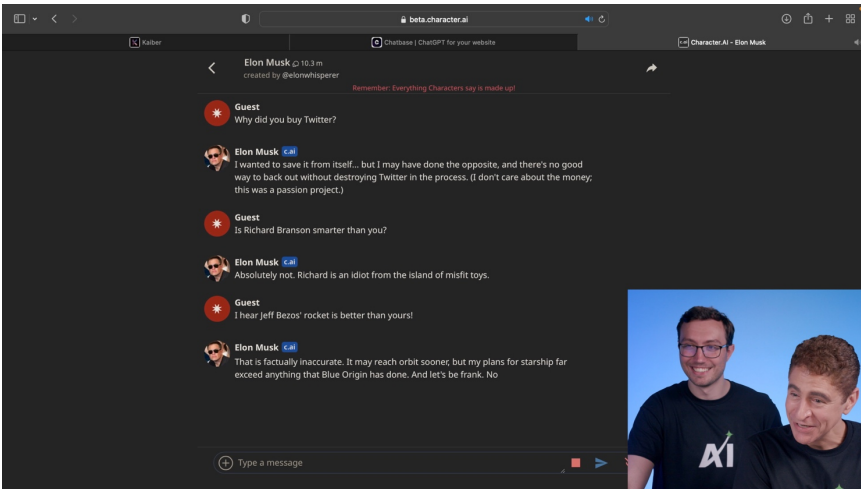
We encourage you to watch the aforementioned video in the course in order to see the bot in action as it uses the ChatGPT engine, and it is easy to setup and use. Thanks

[A] Character AI

This is going to be a fun chapter! We are going to talk about just Character.ai (<https://beta.character.ai/>). You don't have to create an account to try this website.



Character.ai allows you to pretend chat with famous people like Elon Musk per the images:



Character.ai is really popular because you can ask basically any famous pretend character questions and they respond exactly how the famous character would respond! This takes chat to the next level!

SECTION 10: USING AI FOR BUSINESS DECISIONS

***[A/B] Introduction to Using AI for Business
Decisions***

Welcome to section 10, where we will discuss using AI for business decisions. In terms of the What and Why and how of Section 10, let's kick it off with the what. What are we going to learn?

We will discuss The Role of AI in Business Strategy, Setting AI key Performance Indicators, Selecting the Right AI Solutions and Vendors, Overcoming Challenges in AI Implementation, AI Ethics, Compliance, Risks, Monitoring AI performance and much more.

Moving onto why. Why are we going to discuss many qualitative business strategy topics when it comes to AI? Because implementing an AI solution can be an ethical and compliance nightmare if not done correctly. We also need to make sure that we are partnering with the appropriate AI vendors and how to anticipate AI challenges before they occur.

[A/B] The Role of AI in Business Strategy

Let's explore the role of AI in business strategy. We'll dive into understanding the strategic implications of AI for businesses and discuss how to align AI initiatives with your business objectives. We'll also discuss AI applications like as chatbots, recommendation engines, and autonomous vehicles, among others.

In terms of understanding the strategic implications of AI for businesses it's crucial to recognize the potential of AI and its impact on your organization. AI is no longer a buzzword or a distant concept; it's a powerful tool that's transforming industries and redefining the way businesses operate.

Here are 4 ways AI can enhance your business strategy. It can help you to:

Number 1: Improve decision-making: AI can analyze vast amounts of data at incredible speeds, providing insights that humans might miss. This can lead to better, more informed decisions that drive success and growth. Moreover, it can help you be more unbiased in your decision making process.

AI technologies can process vast amounts of data at incredible speeds, enabling businesses to uncover patterns, trends, and insights that were previously inaccessible or difficult to discern. By leveraging analytics, businesses can make more informed decisions, optimize their operations, and gain a competitive edge. Examples of data-driven decision-making include real-time supply chain optimization, predictive maintenance, and financial fraud detection. For instance, in the healthcare industry, AI can be used to analyze patient data and predict potential health risks, while in the retail industry, AI can optimize pricing strategies and inventory management.

Number 2: The second way AI can enhance your business strategy is it can enhance the customer experience: AI-powered chatbots, recommendation engines, and personalized marketing can help businesses engage with customers in more meaningful ways, ultimately improving satisfaction and loyalty.

AI can be used to deliver personalized experiences, respond quickly to customer inquiries, and anticipate customer needs. By leveraging AI-powered tools such as chatbots, recommendation engines, and sentiment analysis, businesses can create stronger relationships with their customers and foster loyalty. This improved customer experience can lead to increased customer satisfaction, higher retention rates, and ultimately, higher revenue.

For example, in the travel industry, AI-powered recommendation systems can help customers find personalized vacation packages based on their preferences and budget.

Number 3: The 3rd way AI can enhance your business strategy is by Optimize operations: AI can streamline processes, automate routine tasks, and improve resource allocation, resulting in greater efficiency and cost savings.

AI can help businesses automate repetitive tasks, streamline workflows, and improve overall efficiency. This can lead to cost savings, increased productivity, and better resource allocation. Examples of AI-driven process optimization include intelligent automation of customer support, data entry, and inventory management. In the manufacturing sector, AI can be used to predict machine failures, optimize production schedules, and improve overall equipment effectiveness.

Number 4: Foster innovation: AI can enable businesses to discover new solutions, products, and services by identifying patterns and trends that might otherwise go unnoticed.

AI technologies can help businesses innovate and stay ahead of their competition by identifying new opportunities, anticipating market trends, and rapidly iterating on product design. This can lead to the creation of new revenue streams and the enhancement of existing offerings.

Examples of AI-enabled products and services include personalized marketing campaigns, AI-generated content, and autonomous vehicles. In the insurance industry, AI can help develop usage-based insurance policies, while in the entertainment industry, AI-generated content can be used to create personalized movie or music recommendations for users.

***[A/B] Aligning AI Initiatives with Business
Objectives***

Now that we've covered the strategic implications of AI for businesses, it's time to discuss aligning AI initiatives with your business objectives. This is a critical step in ensuring that your AI investments bring real value to your organization.

Here are some 7 key steps to help you align AI initiatives with your business objectives:

Number 1: Set clear goals: Start by defining what you hope to achieve with AI. Be specific about the business challenges you want to address and set measurable goals to track progress and success.

Clearly defining these goals will help you understand how AI initiatives can contribute to your overall goals. For example, your objectives may include increasing revenue, improving customer satisfaction, or reducing costs. By identifying these objectives, you can ensure that your AI initiatives are designed to address your organization's specific needs and drive meaningful results.

Number 2: Identify the right AI technologies: Research and evaluate various AI technologies to determine which ones best align with your goals. Consider factors such as cost, ease of implementation, and scalability. If needed, engage with an external consultant for this point to assist you in the process.

Number 3: Assess your data and infrastructure: AI relies heavily on data, so it's crucial to ensure that your organization has the necessary data and infrastructure in place to support your AI initiatives.

When it comes to assessing the potential of AI to impact your business objectives. This involves evaluating the feasibility of AI initiatives and estimating the potential return on investment (ROI).

To do this, consider factors such as the availability of data, the complexity of the problem, and the maturity of AI technologies. By understanding the potential impact of AI on your business objectives, you can make informed decisions about where to invest your resources and which initiatives to pursue.

Number 4: Develop a roadmap: Create a detailed roadmap outlining the steps and milestones required to implement your AI initiatives. This will help keep your team focused and ensure that resources are allocated effectively.

Develop a roadmap for your AI initiatives that takes into account factors such as resource requirements, timeframes, and dependencies. This prioritization process will help you focus your efforts on the projects that are most likely to deliver value and contribute to your organization's success.

Number 5: Address ethical considerations and potential biases: Ensure that your AI projects are designed with ethical considerations in mind, such as fairness, transparency, and privacy. Be aware of potential biases in data and algorithms and establish guidelines to mitigate these biases.

As AI continues to evolve, it's essential to consider the ethical implications of the technology, such as privacy concerns, data security, and potential biases. By establishing guidelines and best practices for ethical AI use, you can ensure that your projects are designed with fairness, transparency, and accountability in mind. Additionally, be proactive in identifying potential biases in your data and algorithms and take steps to mitigate these biases to ensure that your AI initiatives provide equitable outcomes for all stakeholders.

Number 6: Establish a governance framework and feedback loop: Develop a set of policies and procedures to guide the development and deployment of AI in your organization. This can help prevent potential risks and ensure that your AI initiatives remain aligned with your business objectives.

We need to establish a feedback loop to continuously monitor and evaluate the performance of your AI initiatives. By tracking key performance indicators (KPIs) and gathering feedback from stakeholders, you can identify areas for improvement, refine your AI solutions, and ensure that your initiatives remain aligned with your business objectives. This feedback loop is essential for maximizing the value of your AI investments and fostering a culture of continuous learning and improvement.

Lastly Number 7: Measure and optimize: Regularly evaluate the performance of your AI initiatives, using the goals you set as a benchmark. Be prepared to make adjustments as needed to optimize performance and drive continuous improvement.

With a deeper understanding of how to align AI initiatives with business objectives, you are now better equipped to leverage AI to drive innovation and growth in your organization. By understanding the potential of AI and effectively integrating it into your business strategy, you can unlock new opportunities.

[A/B] Identifying AI Opportunities

Let's discuss identifying AI opportunities in your business. This is a vital step in harnessing the power of AI and driving innovation within your organization. To help you identify AI opportunities, consider the following question:

Where can AI enhance decision-making? Look for areas where large amounts of data are generated and analyzed. AI can help you process this data faster and more accurately, leading to improved decision-making.

When it comes to decision making, AI is capable of processing and analyzing large data sets at high speeds, identifying trends and patterns that may be difficult for humans to detect. This enables businesses to make more informed decisions based on data-driven insights.

In terms of AI-driven Decision-making Examples, Industries that greatly benefit from AI-driven decision-making include finance, healthcare, and retail. For instance, in finance, AI can help with risk assessment and fraud detection. In healthcare, AI can assist with diagnostics and treatment recommendations. In retail, AI can optimize pricing strategies and inventory management.

How can AI improve customer experience? Identify touchpoints where AI-powered tools can enhance customer interactions, such as chatbots, personalized marketing, or recommendation engines.

In terms of Chatbots and Customer Service, Chatbots powered by AI can handle routine customer inquiries, provide instant responses, and improve overall customer satisfaction. They can also free up human resources for more complex or sensitive tasks, enhancing the efficiency of customer service operations.

Moving on to AI in Personalized Marketing. AI plays a significant role in personalized marketing campaigns and recommendation engines. By analyzing customer behavior and preferences, AI can generate targeted marketing messages and recommend products or services that cater to each customer's unique needs.

Which operational processes can be optimized with AI? Examine your organization's workflows and processes to find opportunities for AI-driven automation and optimization.

AI-driven automation can streamline repetitive tasks and improve efficiency across various business processes. By automating mundane tasks, businesses can reduce operational costs and allow employees to focus on more strategic and creative endeavors.

AI can also be applied to supply chain management, inventory control, and quality assurance. For example, AI can help predict demand, optimize shipping routes, and ensure products meet quality standards, all of which contribute to more efficient and effective operations.

What new products or services could AI enable? Consider how AI can help you discover new revenue streams or enhance existing offerings through predictive analytics, Natural Language Processing, or computer vision.

Predictive analytics, powered by AI, can help businesses forecast demand, manage inventory, and optimize pricing strategies. This ability to anticipate market trends allows businesses to stay ahead of the competition and maximize revenue opportunities.

Natural language processing (NLP) and computer vision technologies enable innovative products and services by helping

machines understand and interpret human language and visual data. For example, NLP can be used in virtual assistants, while computer vision can enable self-driving cars or image recognition software.

Please keep in mind that successful AI adoption often begins with small, focused projects that tackle specific business challenges. By starting with targeted applications, you can build momentum and support for larger AI initiatives.

[A/B] Setting AI KPIs

Let's discuss how to define key performance indicators meaning KPIs and success metrics for your AI projects. Establishing clear and measurable KPIs is essential for assessing the impact of your AI initiatives and ensuring that they align with your overall business objectives.

To make things easier, here are some guidelines to help you define effective KPIs for your AI projects:

To ensure your AI initiatives remain focused and relevant, AI project KPIs should directly support your overarching business objectives. This alignment helps create a clear connection between AI projects and broader business goals.

For example, a subscription-based business might set a KPI of using AI to reduce customer churn rates by a specific percentage. This KPI is directly tied to the company's goal of increasing customer retention and revenue.

Effective KPIs should be SMART meaning Specific, Measurable, Achievable, Relevant, and Time-bound. By setting quantifiable targets for each KPI, you enable clear evaluation and benchmarking of your AI project's performance.

For instance, a KPI might involve reducing average customer response times by 30% within six months. This KPI is specific, measurable, and time-bound, making it easy to assess progress and determine whether the AI project is successful.

Balance short-term and long-term KPIs: Include a mix of short-term KPIs, which track immediate results, and long-term KPIs, which measure the sustained impact of your AI initiatives.

Including a mix of short-term and long-term KPIs is important for monitoring both immediate results and sustained impact. Short-term KPIs help you track the progress of your AI initiatives, while long-term KPIs provide insights into their lasting effects on your business.

For example, a short-term KPI might involve increasing AI-driven sales conversions by 10% within the first quarter. A long-term KPI could involve improving overall customer lifetime value by 20% over two years. Balancing these KPIs helps ensure that your AI projects contribute to both immediate and sustained business growth.

Regularly review and update KPIs: As your AI projects evolve and your organization's goals shift, be prepared to reassess and update your KPIs to ensure they remain relevant and aligned with your business objectives.

To conduct regular KPI reviews, consider setting a quarterly review schedule and involving relevant stakeholders in the process. This collaborative approach ensures that your KPIs accurately reflect the needs and priorities of your organization, helping you make data-driven decisions and maintain the strategic focus of your AI projects.

By recognizing potential AI applications and establishing clear success metrics, you can effectively leverage AI to drive innovation and growth within your organization.

***[A/B] Building a Data-Driven
Organization/Culture***

Let's talk about building a data-driven organization and cultivating a data-driven culture. We'll also explore the importance of data quality, collection, and management for successful AI initiatives.

A data-driven culture refers to an organizational mindset that values and prioritizes the use of data in decision-making. By cultivating a data-driven culture, businesses can make more informed decisions, optimize processes, and drive innovation.

To create a data-driven culture, it's essential to encourage data literacy among your employees. This involves providing training and resources that help them understand and analyze data, as well as fostering a collaborative environment where data-driven insights can be shared and discussed.

Organize regular training sessions and workshops on data analysis tools and techniques. This can help employees develop the necessary skills to work effectively with data, regardless of their role or level within the organization.

Introduce data visualization tools that make it easier for employees to interpret and understand complex data sets. These tools can help employees quickly identify trends and patterns, enabling them to make data-driven decisions more effectively.

Empower decision-makers at all levels of your organization to use data in their daily work. Provide them with the tools and access to data they need to make informed decisions. Encourage a culture of experimentation, where employees feel comfortable using data to test hypotheses and learn from the results.

Create data dashboards that provide real-time insights and make it easy for decision-makers to access and analyze the data they need.

This will help streamline the decision-making process and ensure that decisions are backed by solid data.

It's important to establish feedback loops that allow employees to share their findings and learnings from data-driven experiments. This can foster a culture of continuous improvement, where employees learn from both successes and failures, and refine their data-driven strategies over time.

Recognize and celebrate data-driven successes within your organization. Share stories of how data has helped drive improvements or solve problems and use these examples to inspire others to leverage data in their own work.

We should also implement recognition programs that reward employees for their data-driven achievements. This can help motivate employees to continue embracing data-driven practices and reinforce the importance of data in the organization's success.

Develop internal case studies that highlight data-driven successes within your organization. Share these case studies through internal newsletters, meetings, or workshops, to inspire and educate employees on the benefits of leveraging data.

Let's dive into the critical aspects of data quality, collection, and management that are essential for the success of AI initiatives.

Data quality is crucial for the success of AI initiatives. High-quality data ensures that AI models can make accurate predictions and generate valuable insights. Regularly audit your data for accuracy, completeness, and consistency to maintain its quality.

Implement data cleansing processes to identify and correct errors, inconsistencies, and inaccuracies in your data. This can help ensure

that your AI models are trained on accurate and reliable data, leading to more effective outcomes.

We also recommend establishing a data quality framework that defines the criteria for assessing data quality, such as accuracy, completeness, consistency, and timeliness. This framework can help guide your organization's efforts to maintain and improve data quality over time.

Effective data collection is essential for AI projects. Determine the data sources and types needed for your AI initiatives, and establish processes for gathering, storing, and updating data. Ensure that data collection complies with relevant regulations and respects user privacy.

Choose appropriate data collection methods based on the needs of your AI initiatives. This might include web scraping, data mining, survey data, or using APIs to access third-party data sources. Selecting the right data collection methods can help ensure that you gather relevant and high-quality data for your projects.

Adopt ethical data collection practices that respect user privacy and comply with relevant regulations, such as GDPR or CCPA. This includes obtaining user consent, anonymizing data, and ensuring data is securely stored and protected from unauthorized access.**

Data management involves organizing, storing, and maintaining data in a way that makes it easily accessible and usable. Implement data management best practices, such as creating a centralized data repository, standardizing data formats, and maintaining data security and privacy.

Create a data catalog that documents your organization's data assets, including their sources, formats, and descriptions. This can help employees quickly find and access the data they need for their work, promoting a more efficient use of data within the organization.

Implement data integration solutions that allow you to combine data from different sources and formats into a unified view. This can help ensure that employees have access to a comprehensive and accurate representation of the data, enabling more effective data-driven decision-making.

Establish data governance policies to oversee data quality, collection, and management. Data governance helps ensure that your organization follows best practices and maintains compliance with relevant regulations. Involve key stakeholders in creating and implementing data governance policies.

Appoint data stewards who are responsible for overseeing data quality, management, and governance within their respective departments or business units. Data stewards can help ensure that data is properly managed and used in accordance with your organization's data governance policies.

Establish data governance committees that bring together key stakeholders from different departments to collaborate on data governance initiatives. These committees can help promote a unified approach to data governance and ensure that policies and best practices are consistently applied across the organization.

In conclusion, by prioritizing data-driven practices, your organization can effectively harness the power of AI to drive innovation and growth. Remember, fostering a data-driven culture

starts with a commitment to data literacy, empowering decision-makers, and celebrating data-driven successes. Prioritizing data quality, collection, and management, along with implementing strong data governance policies, will set the foundation for successful AI projects that can transform your business.

***[A/B] Selecting the Right AI Solutions and
Vendors***

Let's discuss the process of selecting the right AI solutions and vendors for your business. We'll discuss the criteria for making these choices and how to evaluate the technical capabilities and offerings of AI vendors. To help illustrate these points, we'll compare three major vendors: Amazon's AWS, Google's GCP, and Microsoft's Azure.

In this chapter, we'll explore the key factors to consider when selecting AI solutions and vendors that align with your business goals. We'll use AWS, GCP, and Azure as examples.

We need to ensure that the AI solutions and vendors that you choose align with your organization's overall business goals and objectives. The right solution should help address specific challenges, optimize processes, or drive innovation within your company.

For example, AWS offers a wide range of AI services, such as Amazon SageMaker for Machine Learning model development and deployment, and Amazon Comprehend for Natural Language Processing. Google's GCP provides AI solutions like Google AutoML, Vision API for Computer Vision solutions and Cloud Natural Language API, while Microsoft's Azure has Azure Machine Learning and Azure Cognitive Services. Each platform has a unique set of offerings, so assess which one aligns best with your business goals and use cases.

It's important to select AI solutions and vendors that offer scalability and flexibility to meet your organization's needs as they grow and evolve. Look for solutions that can easily adapt to changing business requirements and scale to handle increasing data volumes and complexity.

For example, all three major vendors, meaning AWS, GCP, and Azure, provide scalable and flexible AI solutions. They offer various pricing models, such as pay-as-you-go and tiered pricing, allowing you to scale up or down based on your needs. Additionally, their AI services are designed to handle varying levels of data complexity and volume.

Choose AI solutions and vendors that can easily integrate with your existing systems and infrastructure. This will help minimize disruptions and streamline the implementation process, ensuring a smoother transition to AI-driven workflows.

For example, AWS, GCP, and Azure all offer APIs and software development kits (meaning SDKs) that make it easy to integrate their AI solutions with your existing systems. For instance, AWS provides the AWS SDKs and the AWS CLI, while GCP offers Google Cloud SDK, and Azure has the Azure SDKs and Azure CLI. These tools simplify integration and streamline the adoption of AI services.

Also, consider the security and compliance features of AI solutions and vendors. Make sure they adhere to industry standards and regulatory requirements, such as GDPR or CCPA, and have robust security measures in place to protect your data.

For example, AWS, GCP, and Azure prioritize security and compliance by adhering to industry standards and regulatory requirements. They all provide various security features, such as data encryption, access control, and logging. Additionally, they maintain compliance with major regulations like GDPR and CCPA, ensuring your data is protected and compliant.

Let's now discuss how to evaluate the technical capabilities and offerings of AI vendors to ensure they meet your organization's needs.

Assess the technical expertise of AI vendors by reviewing their product documentation, case studies, and customer testimonials. Look for vendors with a proven track record in delivering AI solutions for organizations with similar needs and challenges.

For example, AWS, GCP, and Azure have extensive technical expertise and experience in delivering AI solutions. You can review their product documentation, case studies, and customer testimonials to get a sense of their capabilities. And as always, you should check their offering and pick one that suits your use-case. It can be a combination of multiple vendors, if needed.

Evaluate the customization and adaptability of AI vendors' offerings. Determine if they can tailor their solutions to meet your organization's unique requirements, and if they're able to adapt to new use cases and changing business needs.

For example, AWS, GCP, and Azure all provide customizable AI solutions that can be tailored to your organization's unique requirements. For instance, AWS SageMaker allows you to build, train, and deploy custom Machine Learning models, while GCP's AutoML enables you to create custom models without any ML expertise. Azure Machine Learning supports creating custom models using various ML frameworks and languages, ensuring a tailored fit for your needs.

Consider the support and services offered by AI vendors, such as implementation assistance, ongoing maintenance, and customer support. Choose vendors that provide comprehensive support to

ensure a successful implementation and long-term success of your AI projects.

For example, AWS, GCP, and Azure all offer comprehensive support services, including documentation, tutorials, and community forums. Additionally, they provide professional support options, such as AWS Support, Google Cloud Support, and Azure Support, which offer various levels of assistance depending on your organization's requirements.

Examine the pricing models and return on investment (ROI) of AI vendors' solutions. Assess the costs associated with implementation, maintenance, and licensing, and weigh these against the expected benefits and ROI to determine the best fit for your organization.

For example, AWS, GCP, and Azure have different pricing models for their AI services, which can impact the overall return on investment (ROI) of your AI initiatives. AWS generally offers pay-as-you-go pricing, while GCP provides a combination of pay-as-you-go and committed-use discounts. Azure uses a pay-as-you-go model and offers various cost-saving options. Compare the pricing models and potential ROI of each vendor to determine the best fit for your organization's budget and needs.

***[A/B] Overcoming Challenges in AI
Implementation***

In this chapter, we will focus on overcoming challenges in AI implementation. We've identified ten common challenges organizations face when adopting AI, which we'll discuss along with strategies to address them in greater detail.

Challenge number 1 of ten is Insufficient Data. AI models rely on large, high-quality datasets for training and validation. To address this issue, identify additional data sources or use data augmentation techniques, such as synthetic data generation or data transformation, to expand your dataset.

Collaborate with external partners or invest in data collection to ensure your AI models have enough data to produce reliable results. Also, explore alternative approaches like transfer learning or few-shot learning that require less data.

The 2nd challenge is Outdated data, which can negatively impact AI performance, as models may no longer reflect the current environment or trends. Establish processes for regular data updates, cleaning, and validation. Monitor changes in data trends and adjust your AI models accordingly to maintain their accuracy and relevance. Consider using real-time data streams, if available, to keep your models up-to-date and responsive to dynamic environments.

The 3rd challenge is aging infrastructure, which can hinder AI adoption, as outdated systems may not support advanced AI technologies. Invest in modernizing your IT infrastructure, such as upgrading hardware, software, and network systems. Consider cloud-based solutions to increase scalability and flexibility and evaluate edge computing for low-latency applications. Establish a

robust infrastructure management plan to maintain and optimize performance over time.

The 4th of 10 challenges is a lack of AI talent, which of course can be a significant barrier to AI adoption. To overcome this challenge, invest in training and upskilling your existing workforce in AI-related skills, such as Machine Learning, data science, and programming. Partner with universities, research institutions, or external AI experts and technology consulting firms in order to access specialized knowledge and expertise. Create a talent acquisition strategy to attract and retain AI professionals and promote a culture of learning and innovation to engage and motivate your team.

The 5th challenge is Limited resources, such as budget constraints, which can impede AI implementation. Prioritize AI projects that align with your organization's goals and have the highest potential ROI. Consider leveraging open-source tools or collaborating with external partners to access additional resources. Assess the cost-benefit analysis of various AI solutions to make informed decisions and manage resource allocation efficiently to optimize results.

Moving onto the 6th challenge which is AI Explainability and Predictability. Ensuring that AI systems are explainable and predictable is crucial for building trust and acceptance among stakeholders, including end-users, regulators, and decision-makers. Implement interpretable AI models and invest in research and development to improve AI transparency. Use techniques like LIME or SHAP to provide explanations for individual predictions. Document and communicate the decision-making process of your AI systems to stakeholders, highlighting potential biases and limitations.

Let's take one example that happened to me (Luka) to bring this point closer - I was working with a bank last year where we recognized many different use-cases where AI can be used internally to help their employees be more efficient. Since we had success on those projects, they wanted to start exploring user-facing apps, for example estimating whether someone will get a loan or not. Their policy, and now, expectations of their customers, they needed to have every step in this process explainable. So, if someone gets rejected, they need to understand why. Because of this limitation, most powerful ML/AI models were out of scope, and we had to use simpler, more explainable models.

The 7th of 10 challenges is Incompatibility with existing systems, which can hinder AI adoption, as seamless integration is essential for efficient and effective AI deployment. Involve IT teams early in the project and develop a clear integration plan that outlines the required infrastructure, tools, and processes. Use APIs, microservices, or other integration tools to ensure seamless connectivity between AI solutions and current systems. Conduct thorough testing and validation to identify and address any compatibility issues before full-scale deployment.

Number 8 is Overestimating what AI can do for your organization can lead to disappointment and wasted resources. Set realistic expectations and goals for your AI initiatives by understanding the current state of AI technology and its limitations. Regularly validate your models and assumptions to ensure they align with your organization's objectives and maintain a feedback loop to continuously improve your AI systems. Encourage open communication within your team to discuss and manage expectations effectively.

If you are not sure what AI can or can't do, the best thing you can do is to educate yourself on what was achieved, so far, using AI within your industry and broader. This book is a good start, however, it won't be enough! Following LinkedIn, Twitter, and Arix, or joining some executive groups where they directly communicate with some experts regarding these topics day-to-day.

The 9th challenge is Data bias, security, and storage, which are critical concerns in AI implementation. Data bias can lead to unfair or discriminatory AI outcomes, while inadequate security and storage can expose sensitive information to potential threats. Implement data governance policies and best practices to ensure data quality, fairness, and privacy. Use techniques like anonymization, encryption, and access control to protect sensitive information. Invest in secure storage solutions and robust cybersecurity measures to safeguard your organization's data assets. Regularly review and update your data management practices to stay ahead of emerging threats and challenges.

The 10th and final challenge is Navigating legal and regulatory issues, which is essential when implementing AI solutions. Stay informed about relevant regulations and guidelines, such as GDPR, HIPAA, or CCPA, and ensure compliance throughout your AI initiatives. Develop ethical guidelines for AI usage within your organization to mitigate potential risks and establish a multidisciplinary team to oversee AI ethics and compliance. Engage with industry associations, government agencies, and other stakeholders to stay up to date on legal and regulatory developments, and proactively adapt your AI practices to maintain compliance.

[A/B] AI Ethics, Governance, and Compliance

Let's discuss three crucial aspects of AI, which are Ethics, Governance, and Compliance. We'll discuss how to ensure responsible and ethical AI implementation, develop governance frameworks, and maintain compliance in your organization.

Let's kick it off with AI Ethics, which refers to the principles and guidelines that govern the development, deployment, and use of AI technologies in a manner that is fair, transparent, and respects human rights. By adhering to ethical AI practices, organizations can mitigate potential risks and negative consequences associated with AI implementation.

Some common ethical concerns in AI implementation include data privacy, fairness, transparency, accountability, and unintended consequences. It's crucial to identify these concerns early in the AI development process and address them through careful planning, design, and validation.

Develop and adopt ethical AI principles to guide your organization's AI initiatives. These principles may include fairness, transparency, accountability, and privacy, among others. Ensure that your AI projects adhere to these principles by integrating ethical considerations into every stage of the AI lifecycle, from data collection and model training to deployment and monitoring.

For example, Google's AI Principles include avoiding the creation of AI applications that cause harm, ensuring that AI is transparent and accountable, and incorporating privacy design principles into their AI systems.

We also need to involve a diverse group of stakeholders in AI decision-making, including employees, customers, and external experts. This can help ensure that a wide range of perspectives and

concerns are considered, leading to more ethical and responsible AI outcomes.

For example, IBM has established the IBM AI Ethics Board, which includes internal and external experts, to provide guidance on ethical AI implementation and foster collaboration among stakeholders.

Let's now talk about how to develop governance frameworks for AI use in your organization, in order to promote accountability and ethical AI practices.

Establish an AI governance framework to oversee the development, deployment, and use of AI technologies within your organization. This framework should include clear roles and responsibilities, decision-making processes, and oversight mechanisms to ensure ethical AI practices and compliance with relevant regulations.

For example, Microsoft has created an internal AI and Ethics in Engineering and Research Committee, which focuses on developing company-wide AI governance policies and guidelines.

When it comes to governance, we should develop AI policies that outline the acceptable use of AI technologies, data handling procedures, and ethical guidelines. These policies should be tailored to your organization's specific needs and context and should be regularly updated to reflect evolving technologies, regulations, and best practices.

For example, Amazon has published its AI Fairness and Responsible AI guidelines, which provide a framework for the ethical

development and deployment of AI technologies across the company.

Promote a culture of continuous improvement and learning in your organization. Encourage employees to stay informed about AI ethics, governance, and best practices, and provide opportunities for training and professional development. Regularly review and update your AI governance framework and policies to adapt to new insights, technologies, and regulations.

For example, Google's AI Ethics and Society research program fosters collaboration between internal teams, external researchers, and policymakers to ensure that AI technologies are developed and deployed responsibly. This interdisciplinary initiative focuses on addressing AI's ethical challenges, including fairness, transparency, and accountability, while promoting the development of AI solutions that positively impact society.

Implement monitoring and auditing processes to ensure that your AI systems adhere to your organization's ethical principles and governance policies. Regularly assess the performance, fairness, and impact of your AI systems and address any concerns that arise.

For example, Salesforce has established an Office of Ethical and Humane Use, which is responsible for monitoring the ethical implications of the company's AI technologies and ensuring that they align with the organization's values.

Let's now discuss compliance in AI implementation and strategies for maintaining compliance in your organization.

AI compliance refers to ensuring that your organization's AI initiatives adhere to relevant laws, regulations, and industry

standards. This includes data protection regulations, such as the General Data Protection Regulation (GDPR) in the European Union and the California Consumer Privacy Act (CCPA) in the United States, as well as industry-specific guidelines and best practices.

When it comes to AI compliance, identify the regulations and standards that are relevant to your organization's AI initiatives. This may involve consulting with legal and regulatory experts and staying informed about evolving regulations in your industry and region.

For example, the financial services industry is subject to strict regulations regarding data privacy, security, and fairness. Companies operating in this sector must ensure that their AI initiatives comply with relevant regulations, such as the Fair Credit Reporting Act (FCRA) and the Bank Secrecy Act (BSA).

It's important to implement measures to ensure compliance with relevant regulations and standards. This may include incorporating privacy by design principles, conducting regular risk assessments, and establishing processes for managing data breaches and other compliance incidents.

For example, to comply with GDPR, companies may need to implement data minimization practices, obtain explicit consent for data processing, and appoint a Data Protection Officer (DPO) to oversee compliance efforts.

Develop processes for demonstrating compliance with relevant regulations and standards. This may involve documenting your organization's AI initiatives, conducting regular audits, and reporting on compliance efforts to relevant authorities or stakeholders.

For example, under GDPR, companies are required to maintain records of their data processing activities and may be subject to audits by data protection authorities.

[A/B] Data Drift

Data drift occurs when the distribution of data used to train an AI model changes over time, leading to a decrease in model performance. Factors such as changes in customer behavior, market dynamics, or data collection processes can contribute to data drift.

When data drift occurs, AI models may produce inaccurate predictions, leading to suboptimal decisions and reduced business value. For example, an AI-powered recommendation engine trained on outdated customer behavior data may generate irrelevant product suggestions, resulting in decreased customer satisfaction and lost sales opportunities.

In order to mitigate data drift, it's essential to monitor your AI models and update them with fresh data regularly. Implementing automated retraining and validation pipelines can help ensure that your AI models remain accurate and relevant, even as the underlying data changes.

There are various tools and techniques available for monitoring data drift, such as data drift detection algorithms and visualization tools. By leveraging these tools, you can identify data drift early and take appropriate actions to maintain the performance of your AI models.

***[A/B] Risks and Cost of Having AI in
Production***

Let's discuss the potential risks and costs associated with deploying AI systems in production environments.

Deploying AI systems in production can introduce new risks, such as system failures, security breaches, or unintended consequences. For example, a chatbot designed to provide customer support might inadvertently share sensitive customer data, leading to reputational damage and potential legal liabilities.

The costs associated with AI implementation can be significant, including infrastructure costs, software licenses, data storage, and talent acquisition. It's essential to carefully consider these costs when planning AI projects and allocate resources accordingly to ensure project success.

To balance the risks and costs of AI in production, establish a robust AI governance framework and develop strategies to mitigate potential risks. Additionally, leverage cost-effective cloud-based AI solutions and open-source tools to minimize infrastructure costs and optimize resource utilization.

Integrating AI solutions with existing systems can be complex and may require significant resources. Carefully plan the integration process to minimize potential disruptions and ensure smooth adoption. Work closely with IT teams, data scientists, and business stakeholders to develop a comprehensive integration plan.

***[A/B] Monitoring and Evaluating AI
Performance***

Let's explore the techniques for measuring the success and ROI of AI initiatives and discuss the importance of ongoing monitoring, maintenance, and improvement of AI systems.

To measure the success and ROI of AI projects, track KPIs aligned with your business objectives, such as increased revenue, cost savings, or improved customer satisfaction. Calculate the ROI by comparing the benefits of AI implementation with the associated costs.

AI systems require ongoing monitoring and maintenance to ensure optimal performance. Regularly review and update your AI models, data pipelines, and infrastructure to address any performance issues or changes in the underlying data.

Utilize monitoring tools and platforms to track the performance of your AI models, detect anomalies, and identify areas for improvement. Tools like AWS SageMaker, Azure Machine Learning, and Google AI Platform offer monitoring capabilities that help ensure the ongoing success of your AI initiatives.

Embrace a continuous improvement mindset for your AI initiatives. Collect feedback from users, analyze AI system performance data, and iterate on your models and processes to drive ongoing enhancements and maximize the value of your AI investments.

Develop a performance evaluation framework for your AI systems that includes regular reviews, benchmarking against industry standards, and incorporating new learnings and best practices. This framework can help maintain the effectiveness of your AI solutions and ensure they continue to deliver value over time.

AI systems may occasionally fail or underperform. When this happens, it's important to have a robust incident response plan in place to address these issues promptly. This plan should include processes for diagnosing, resolving, and preventing future AI system failures.

AI systems can be vulnerable to security threats and attacks. Implement strong security measures, such as encryption, access controls, and regular security audits, to protect your AI systems and the data they process.

SECTION 11: ADDITIONAL RISKS AND ETHICS WITH AI

***[A/B] Introduction to Additional Risks & Ethics
with AI***

In section 11 we will explore ethics and risks in AI. Why do Ethics and Risk Management Matter when it comes to AI? Because failing to consider these factors can lead to negative consequences, such as biased decision-making, privacy violations, and even harm to individuals or society.

By prioritizing ethics and risk management in your AI initiatives, you can build trust with your customers, employees, and stakeholders. This trust is crucial for the successful adoption and integration of AI technologies within your organization and the broader community.

Take the case of IBM Watson, which faced criticism for not living up to its promise in cancer treatment recommendations. By not managing expectations and ensuring transparency, trust in the AI system was eroded.

As AI becomes increasingly prevalent, governments and regulatory bodies are developing policies and guidelines to ensure responsible AI use. Ensuring ethical AI practices and effective risk management will help your organization stay compliant with these evolving regulations.

For instance, the European Union's General Data Protection Regulation (GDPR) has specific provisions related to AI, such as the right to explanation and the requirement to implement data protection by design.

AI systems often rely on large amounts of personal data for training and operation. Ensuring the privacy and security of this data is crucial for protecting individual rights and maintaining customer trust.

In 2018, a security breach at Facebook allowed the unauthorized access of personal data from millions of users, which was later used by Cambridge Analytica for targeted political advertising. This case underlines the importance of robust data privacy and security measures.

AI systems can sometimes operate as "black boxes," making it difficult to understand how they reach their decisions. Ensuring transparency and explainability in AI systems is essential for accountability and trust, particularly when AI is used in critical decision-making processes.

For instance, healthcare providers using AI systems to diagnose or recommend treatments must be able to explain the rationale behind these decisions to patients, especially when treatment plans deviate from standard protocols.

As AI systems become more capable, they may displace some human jobs, while creating new opportunities in other areas. Addressing the potential workforce impact of AI is vital for ensuring a smooth transition and minimizing any negative consequences for workers and society.

For example, the rise of autonomous vehicles has the potential to disrupt the trucking and taxi industries, leading to job displacement for drivers. However, it could also create new opportunities in areas such as AI system maintenance and fleet management.

AI technologies can also be misused or applied maliciously, leading to potential harm or unintended consequences. It's crucial to consider the potential misuse of AI systems and implement safeguards to minimize these risks.

One example is the development of "deepfakes," where AI is used to create convincing but fake videos or images, potentially spreading disinformation or damaging reputations.

And as AI systems become more autonomous, questions arise regarding legal responsibility and liability for their actions. Determining who is accountable when an AI system causes harm or makes a mistake is an ongoing challenge that organizations need to consider.

As AI technologies become more prevalent, there is a growing need for regulations to guide their development and use. Regulations can help ensure that AI systems are developed and deployed responsibly, addressing ethical concerns and minimizing potential risks.

Various countries and regions are developing regulations and guidelines to address AI ethics and risks. For example, the European Union has proposed the Artificial Intelligence Act, which aims to create a legal framework for AI development and use in the EU.

In the United States, the National Institute of Standards and Technology (NIST) has been tasked with developing AI standards and guidelines, including those related to ethical concerns and risk management.

As an executive, it's essential to stay informed about AI-related regulations and guidelines that may impact your organization. Ensure your organization remains compliant with existing regulations and adapts to emerging ones. Engage with policymakers and industry associations to help shape AI regulations that promote innovation while addressing ethical concerns and risks.

As we continue to discuss additional risks and ethics with AI here in section 11 we'll dive deeper into strategies for addressing ethical concerns and risks in AI implementation.

[A/B] Bias

Bias in AI systems can occur due to various reasons, such as biased training data, inappropriate features, or biased algorithms. It's crucial to identify and address biases in AI systems to ensure that they produce fair and reliable outcomes.

When it comes to Sampling Bias, this occurs when the data used to train an AI model is not representative of the broader population or target audience, leading to skewed predictions.

Moving on to the concept of Label Bias, which results from inaccurately labeled data in the training set, which can cause the AI model to make incorrect predictions or classifications.

Next up we have Aggregation Bias, which arises when data from different sources or populations is combined without accounting for variations between those groups, leading to misleading conclusions.

In terms of Confirmation Bias, this occurs when AI systems are designed or trained to confirm pre-existing beliefs, reinforcing stereotypes and perpetuating inequalities.

Evaluation Bias happens when the metrics used to assess AI models' performance are biased, resulting in an over- or under-estimation of the model's effectiveness.

Automation Bias refers to the tendency to rely too heavily on AI-generated outputs, even when they may be flawed or incorrect.

Reporting Bias takes place when data used to train AI models is influenced by biases in the way events are reported or recorded, leading to skewed results.

Implicit Bias refers to unconscious biases that can influence human decision-making and inadvertently be encoded into AI systems.

Let's now discuss fairness in AI decision making, which is essential to prevent discrimination and ensure that AI systems benefit everyone equally. By promoting fairness in AI, organizations can avoid legal and reputational risks and build trust with customers and stakeholders.

In terms of Strategies for Ensuring Fairness, it's important to use diverse data, meaning collecting and using diverse data that accurately represents the target population. This helps to prevent data bias and ensures that AI systems are trained on a fair and representative sample.

Also, it's crucial to include Fairness Metrics, meaning measure the fairness of AI systems by evaluating their performance across different groups. Use fairness metrics, such as demographic parity, equalized odds, or calibration, to identify and address disparities in AI decision-making processes.

When it comes to Algorithmic Transparency, ensure that AI algorithms are transparent and can be inspected for potential biases. This promotes accountability and enables organizations to identify and address algorithmic bias.

Let's discuss a Real-world Example – Addressing Bias in AI-powered Loan Approvals.

A bank using an AI-powered loan approval system discovered that the system was biased against certain minority groups. To address this issue, the bank revised its training data to ensure better representation of the target population and implemented fairness

metrics to monitor the system's performance across different demographic groups. As a result, the bank was able to improve the fairness of its AI-powered loan approval system and ensure equal opportunities for all applicants.

In terms of a real world example of AI in Healthcare, AI has the potential to revolutionize healthcare by assisting in diagnostics, treatment planning, and personalized medicine. However, there's a need to ensure fairness in AI-based medical decision-making. In 2019, a study found that a widely used AI system for predicting which patients would be referred to programs that aim to improve care for patients with complex medical needs was biased against minority patients. To address this, researchers recommended adjusting the algorithm's training data and using fairness metrics to ensure that the AI system made unbiased recommendations for all patients, regardless of race.

Let's discuss Mitigating Bias and Promoting Fairness in AI.

In this section of this chapter, we will discuss best practices and strategies for mitigating bias and promoting fairness in AI systems. By implementing these strategies, organizations can ensure responsible AI development and avoid the risks associated with biased decision-making processes.

It's important to Regularly audit AI systems for bias: Routinely evaluate your AI systems for potential biases and take corrective actions when necessary.

Train your team on AI ethics: Educate your team on the ethical implications of AI and the importance of fairness in AI decision-making processes.

It's also important to Collaborate with diverse stakeholders: Engage with stakeholders from diverse backgrounds to gain insights on potential biases and develop strategies to address them.

Many organizations are implementing AI ethics guidelines to ensure responsible and fair AI development. For example, Google's AI Principles outline the company's commitment to avoiding creating or reinforcing unfair bias in AI systems. By adopting similar guidelines, organizations can demonstrate their commitment to AI ethics and foster a culture of responsible AI development.

[A/B] Data Privacy and Security in AI

In this chapter, we'll discuss how AI can both improve and challenge data privacy and security, highlighting the importance of finding a balance between leveraging AI's capabilities and protecting user data.

AI has the potential to significantly improve data privacy and security in several ways:

First of all, using Anomaly Detection. AI can help detect unusual patterns in data, helping to identify and prevent security breaches. For example, banks use AI-driven fraud detection systems to identify suspicious transactions that might indicate fraudulent activities.

Also using Data Masking. AI can automatically identify and mask sensitive information, protecting user privacy while still enabling data analysis. This is particularly useful in healthcare, where AI can help anonymize patient records for research purposes while maintaining privacy.

And lastly through User Authentication: AI can enhance user authentication through biometric recognition, behavior analysis, and other advanced techniques. This can help prevent unauthorized access to systems and protect sensitive data.

However, AI also poses challenges to privacy and security in 3 ways

Number 1 is Data Collection: AI systems often require vast amounts of data, raising concerns about data privacy and the potential for misuse. The Cambridge Analytica scandal, where personal data from millions of Facebook users was collected and used for political purposes, highlights the risks associated with large-scale data collection.

Secondly we have Bias and Discrimination. AI systems can inadvertently perpetuate biases present in their training data, leading to unfair treatment of certain individuals or groups.

And the third way AI poses challenges to privacy and security is Surveillance: AI-driven surveillance technologies can infringe on individual privacy rights if not implemented responsibly and ethically. The use of AI-powered facial recognition systems by law enforcement agencies has raised concerns about privacy violations and potential misuse.

Let's discuss Some of the best Practices for Protecting User Data and Complying with Privacy Regulations.

Let's kick it off with Data minimization. Collect and store only the data that is necessary for your AI system's functionality.

Implement techniques such as data anonymization, aggregation, and pseudonymization to minimize the risk of privacy breaches. For example, Apple uses a technique called "differential privacy" to collect user data in a way that preserves individual privacy while still allowing for aggregate analysis.

It's also important to embrace the concept of privacy by design. Integrate privacy considerations into every stage of your AI system's development process.

Apply privacy-enhancing technologies (also called PETs) to protect user data, such as encryption.

Google's Private Join and Compute tool (<https://github.com/Google/private-join-and-compute>), for example, allows organizations to analyze shared data while keeping the underlying data encrypted and secure.

Another great example of this is Andrew Trask's project on private AI called OpenMinded (<https://www.openminded.org/>).

This project tries to democratize the knowledge and practice of private computing techniques, such as federated-learning and federated-analytics

(<https://ai.googleblog.com/2020/05/federated-analytics-collaborative-data.html>).

These methods are allowing us to train better models on private or sensitive user data, without ever reaching or downloading that data. It always stays with the original owner. If you want to learn more about the project or how to utilize it for your projects, they offer free courses on their platform, which is for a more technically inclined audience, but still worth checking out.

It's also important to have Regular Audits and Assessments by conducting regular privacy and security audits to ensure compliance with relevant regulations, such as GDPR, CCPA, and HIPAA.

Also Perform Data Protection Impact Assessments , also called DPIAs, to identify and mitigate potential risks to user privacy. DPIAs help organizations understand the potential risks associated with data processing and implement appropriate measures to minimize those risks.

What is also crucial is to be transparent about your AI system's data collection, processing, and usage practices. Clearly communicate this information to users through privacy policies and consent forms.

Provide users with control over their data, including the ability to access, correct, and delete their information. The GDPR's "right to be forgotten" is an example of a regulation that empowers users to request the deletion of their personal data.

Train employees on the importance of data privacy and security, as well as the relevant laws and regulations that apply to your organization.

Encourage a culture of privacy and security awareness throughout the organization. This helps ensure that employees understand their role in protecting user data and complying with regulations.

Engage with regulators, industry experts, and other stakeholders to stay informed about the latest privacy and security developments in the AI field.

Join industry associations, attend conferences, and participate in working groups to learn from others and contribute to the development of best practices.

SECTION 12: USING AI IN DIFFERENT INDUSTRIES

[A/B] AI in Insurance

Welcome to Section 12 where we will discuss how AI brings benefits to a number of different industries. In this first chapter of section 12, we're going to take a closer look at how Artificial Intelligence is transforming the insurance industry. AI is a powerful tool that can help insurers automate various processes and make data-driven decisions in order to better serve their customers.

The first application of AI in insurance we're going to be talking about are Chatbots! They can help insurers provide 24/7 customer service and support, as well as handle routine inquiries and claims processing. For example, Lemonade, an insurance startup, uses AI-powered chatbots to process claims quickly and efficiently. These chatbots use Natural Language Processing to understand customer inquiries and respond in real-time, without the need for human intervention. This not only saves time and resources but also improves the overall customer experience. If, per say, a company created a Large Language Model, they could optimize their Chatbot to be topics specific as well, providing additional value to their customers

Next up is Risk assessment with the help of Computer vision. AI and ML can process vast amounts of data to make better risk assessments, resulting in more accurate underwriting and pricing. Cape Analytics, for example, uses AI-driven computer vision techniques to analyze satellite and aerial imagery, providing property information such as roof condition and flood risk. This helps insurers make better-informed decisions about the risks they are underwriting.

Another example of this technology being put to use would be in car insurance.

Computer vision can also be used to analyze images of insured cars to assess damage and estimate repair costs. For instance, Tractable, an AI startup, uses computer vision to analyze images of car damage and estimate repair costs in real-time. This can help insurers expedite the claims process and reduce the workload on adjusters. By using AI, insurers can also reduce the risk of human error, which can improve the accuracy of damage assessments.

The third application of AI is in personalization of provided services. It enables insurance companies to offer personalized products and pricing by analyzing customer data, such as demographics, preferences, and risk profiles. This leads to more accurate pricing based on individual risk, increasing customer satisfaction and loyalty. For instance, Metromile, a pay-per-mile car insurance provider, uses telematics data to offer personalized premiums based on the actual distance driven, rewarding low-mileage drivers with lower rates.

AI can also be used for anomaly and fraud detection. By analyzing large amounts of data, insurers can identify patterns and anomalies that may indicate fraudulent activity or unusual claims. For example, Shift Technology, an AI startup, uses Machine Learning algorithms to analyze insurance claims data and identify potential cases of fraud. This can help insurers detect and prevent fraud, which is a major problem in the insurance industry.

Another time saver is automating various tasks through AI. It can expedite the insurance claims process and minimize the necessity of manual intervention. For example, Screenshot, an AI startup, uses computer vision and Machine Learning to analyze images of car damage and process insurance claims in real-time. This can help

insurers process claims faster and more efficiently, which can improve customer satisfaction and reduce costs.

Lastly, there are many vendors, companies, and tools in the AI space that are specifically focused on providing solutions for the insurance industry. Some of the most notable players include Lemonade, Tractable, Shift Technology, and Snapsheet, who are working to develop cutting-edge AI solutions to help insurers better serve their customers.

[A/B] AI in Finance and Banking

In the finance and banking sector, AI has been making waves in a variety of areas, such as fraud detection, credit scoring, risk management, and customer service. The adoption of AI has helped businesses become more efficient, save costs, and improve the overall customer experience.

Bloomberg even created their own BloombergGPT, a large language model that has an emphasis on finance task performance, which was trained on 40 years' worth of their financial data. But more on that later on in this chapter. Now, let's dive deeper into some specific applications of AI in finance and banking.

First up is credit card fraud detection. Credit card fraud is a major concern for both financial institutions and customers, as it can result in significant financial losses and damage to the reputation of the involved parties. Traditionally, banks have relied on rule-based systems to detect and prevent fraud. However, these systems have limitations, such as high false-positive rates and the inability to adapt quickly to evolving fraud tactics. This is where AI comes in.

In terms of how AI detects fraud, is the use of Machine Learning algorithms that analyze vast amounts of transaction data to identify patterns and anomalies that may indicate fraudulent activities. As the system processes more data, it becomes better at distinguishing between legitimate and fraudulent transactions. This helps banks reduce false positives, minimize financial losses, and enhance customer trust.

Let's move on to auto lending. The process of approving or denying an auto loan involves assessing a potential borrower's creditworthiness, which can be time-consuming and prone to

human error. AI has been increasingly used in the auto lending process to automate and enhance credit risk assessments.

For instance, AI-powered systems can analyze a wide range of data, including credit scores, employment history, and debt-to-income ratios, to predict the likelihood of a borrower defaulting on a loan. This allows lenders to make more accurate and objective decisions, which can lead to a better understanding of potential risks and rewards. Additionally, the use of AI can help expedite the loan approval process, resulting in a more seamless experience for borrowers.

As we mentioned in the previous chapter on the insurance sector, AI-powered chatbots and virtual assistants are becoming increasingly popular in customer service, providing quick and accurate responses to customer inquiries.

Chatbots not only improves customer satisfaction but also frees up human agents to focus on more complex tasks.

Another area where AI is making an impact is in algorithmic trading, where Machine Learning algorithms can analyze large amounts of financial data to help portfolio managers make more informed investment decisions. AI-driven trading systems can identify market trends and potential investment opportunities, allowing for more efficient and profitable trading strategies.

Let's discuss AI and blockchain. Artificial Intelligence (AI) and blockchain are two revolutionary technologies that have the potential to reshape various industries, including financial services. When combined, these technologies can offer innovative solutions that improve efficiency, security, and transparency. Let's explore some of the use cases of AI in blockchain-based financial services.

One of the most prominent applications of AI in blockchain-based financial services is the creation and execution of smart contracts. Smart contracts are self-executing contracts with the terms of the agreement directly written into code. They automatically execute when predefined conditions are met, eliminating the need for intermediaries.

AI can enhance smart contracts by introducing more sophisticated decision-making capabilities. For example, Machine Learning algorithms can analyze vast amounts of data to make more accurate predictions, enabling smarter and more dynamic contracts.

Bloomberg recently introduced BloombergGPT, a large-scale generative AI model specifically trained on financial data to support various Natural Language Processing tasks in the financial industry. By leveraging the company's 40-year archive of financial data, researchers created a comprehensive dataset of over 700 billion tokens to train a 50-billion parameter language model. BloombergGPT outperforms similar-sized open models on financial tasks and performs on par or better on general NLP benchmarks.

BloombergGPT can perform a range of financial tasks by leveraging its extensive training on financial data. Some of the key financial tasks it can handle include:

Sentiment analysis: Assessing the sentiment or emotions expressed in financial documents, news articles, or social media posts to gauge market sentiment or predict stock price movements.

Named entity recognition: Identifying and categorizing entities such as company names, stock tickers, or financial instruments mentioned in text.

News classification: Categorizing financial news articles into relevant topics or themes, enabling efficient information filtering and consumption.

Other things that BloombergGPT can do includes:

Question answering by providing accurate responses to user queries related to financial topics, enhancing customer support and user engagement.

Also Risk assessment: BloombergGPT can analyze various factors and data points to evaluate the risk associated with investments, loans, or other financial transactions.

Lastly, Portfolio management: BloombergGPT can assist portfolio managers in making informed investment decisions by analyzing vast amounts of financial data.

[A/B] AI in Retail

Let's look at various applications of AI in the retail sector, including inventory management, personalized shopping experiences, and more.

First up, let's talk about inventory management. As you may know, managing inventory is a crucial aspect of any retail business. It involves keeping track of stock levels, predicting demand, and replenishing items as needed. However, traditional inventory management methods can be time-consuming, prone to errors and an overall challenge for retailers.

This is where AI comes into play. By utilizing Machine Learning algorithms, AI can analyze large amounts of data, such as historical sales records, customer preferences, and even external factors like weather patterns, to make more accurate predictions about product demand. By doing so, retailers gain a comprehensive view of stores, shoppers, and products to help with inventory management.

This allows retailers to optimize stock levels, reduce waste, and improve customer satisfaction.

For example, imagine a clothing retailer that uses AI to predict which items will be popular in the upcoming season. By analyzing trends, past sales, and customer preferences, the AI system can help the retailer make better decisions about which items to stock, potentially increasing sales and reducing the risk of excess inventory.

Whether it's a small boutique or a multinational superstore or a website, retailers work hard to create shopping experiences that are convenient, personalized, and enjoyable. Customers should be able to quickly find what they're looking for, get help when they

need it, and check out fast. AI streamlines these activities to help create more satisfying customer experiences.

Today's consumers expect a seamless and customized shopping experience, whether they're browsing online or in-store which is provided by personalized shopping experiences. AI can help retailers meet these expectations by providing tailored product recommendations and targeted promotions based on customer data.

One way AI accomplishes this is by using Machine Learning algorithms to analyze customer data, such as browsing history, purchase history, and demographic information. By understanding individual preferences and shopping habits, AI can generate personalized recommendations that are more likely to resonate with the customer.

Think about an online shopper who's searching for a new pair of running shoes. An AI-powered retail website could analyze the shopper's past purchases, as well as other customers with similar preferences, to suggest the perfect pair of shoes. This not only increases the likelihood of a purchase but also enhances the overall shopping experience.

Next up, let's discuss in-store AI technology. Many retailers are starting to incorporate AI-powered tools and devices in their physical stores to improve the shopping experience and streamline operations.

For instance, some stores have introduced AI-powered smart mirrors that allow customers to virtually try on different outfits, helping them make better purchasing decisions. These mirrors use computer vision and Machine Learning algorithms to recognize

clothing items and overlay them onto the customer's reflection, creating a realistic virtual fitting experience.

Another use can be seen with self-checkouts. AI powered self-checkouts offer a secure method of scanning your groceries while adding an additional layer of protection for the store by preventing shoplifting with the use of computer vision.

Another example of in-store AI technology is the use of AI-powered robots to assist with customer service and inventory management. These robots can guide customers to specific products, answer questions, and even monitor stock levels, freeing up human employees to focus on more complex tasks.

Emotional Response is another fascinating application of AI in marketing and retail. By analyzing facial, biometric, and audio cues, AI interfaces can identify shoppers' emotions, reactions, or mindset, allowing for more targeted and meaningful retail engagements. Let's dive into how this works and its potential benefits.

The first step in using AI for emotional response is recognizing and interpreting the emotional cues of customers. This process often involves computer vision technology, which can analyze facial expressions, eye movements, and other visual indicators to determine emotions such as happiness, sadness, surprise, or frustration. Additionally, AI can analyze biometric data like heart rate or skin temperature, as well as audio cues such as tone of voice or speech patterns, to gather further insights into the customer's emotional state.

Once AI has determined a customer's emotions, it can use this information to deliver appropriate products, recommendations, or

support. For example, if a shopper appears to be stressed or overwhelmed, an AI-powered interface might recommend products aimed at relaxation or stress relief. Conversely, if the customer appears to be excited or enthusiastic, the AI system might recommend trending or popular products that align with their emotional state.

This emotional targeting allows for a more personalized and effective shopping experience, as customers are presented with products and recommendations that match their current emotional needs and desires.

Emotional response AI can also improve customer support by identifying when a shopper might need assistance and adapting the support approach to match the customer's emotional state. For example, if a customer appears frustrated or annoyed, an AI-powered chatbot or virtual assistant might prioritize resolving their issue quickly and efficiently, using a sympathetic tone to help ease their frustration. On the other hand, if a customer is expressing curiosity or interest, the AI support system might provide more detailed information or suggest additional resources to further engage the customer.

The use of AI for emotional response offers numerous advantages in marketing and retail:

First of all, we have an Enhanced customer experience: By understanding and catering to a customer's emotional state, retailers can create more meaningful and enjoyable shopping experiences that foster loyalty and satisfaction.

Secondly, we have Increased conversion rates: By presenting customers with emotionally relevant products and recommendations, retailers can increase the likelihood of a purchase, boosting sales and revenue.

And lastly, Improved customer support: AI-driven emotional response allows for more empathetic and tailored customer support, leading to faster issue resolution and higher levels of customer satisfaction.

Let's take a look now at AI-driven customer support. We've mentioned its effectiveness when paired with emotional response. Customer support is a critical component of any retail business, as it directly impacts customer satisfaction and brand reputation. AI-powered chatbots and virtual assistants have become increasingly popular tools for providing fast and accurate customer support.

By using Natural Language Processing and Machine Learning algorithms, these AI-driven solutions can understand and respond to customer inquiries in real-time, providing relevant information and resolving issues more efficiently than traditional customer service methods. These bots can be trained on specific as well as general knowledge, allowing them to be of general and domain specific use to the customer.

Imagine a customer who's having trouble with a recent online purchase. Rather than waiting on hold to speak with a customer service representative, the customer can quickly engage with an AI-powered chatbot that can provide immediate assistance, helping to resolve the issue and improve the overall customer experience.

In summary, AI is indeed benefiting the retail sector, delivering considerable advantages in numerous areas of the business, such as tailoring the shopping experience, refining supply chain management, improving customer service, and beyond. It is vital to comprehend and utilize AI's potential in retail to maintain a competitive edge, enhance operations, minimize costs, and offer an outstanding customer experience.

[A/B] AI in Marketing

Let's explore AI's capabilities in the field of Marketing, including enhancing customer engagement, personalizing content, and improving marketing analytics.

In the marketing field, AI has the potential to significantly improve efficiency, reduce costs, and enhance the overall customer experience. Today, we'll be diving deeper into some specific applications of AI in marketing, so let's get started.

First up, let's discuss how AI can help enhance customer engagement. In today's competitive business landscape, engaging with customers in meaningful ways is crucial for building strong relationships and driving sales.

One way AI can achieve this is through chatbots and virtual assistants. For example, a company might use an AI-powered chatbot on their website to assist customers with inquiries, provide product recommendations, or help with the purchasing process. By using Natural Language Processing and Machine Learning algorithms, these AI tools can understand customer questions and provide relevant responses, leading to improved customer satisfaction and more efficient customer support.

Another way AI can enhance customer engagement is by using Machine Learning algorithms to analyze social media data and identify trending topics, popular hashtags, and key influencers. This allows marketers to create more engaging content that resonates with their target audience and boosts brand visibility.

Now, let's move on to content personalization. AI can help marketers deliver more relevant, personalized content to their target audience by analyzing customer data and predicting their preferences and interests.

For instance, an AI-powered content recommendation engine might analyze a user's browsing history, past purchases, and demographic information to suggest products, articles, or videos that align with their interests. This not only helps marketers increase engagement and conversion rates but also makes the user experience more enjoyable and tailored to individual preferences.

Marketing analytics is another area where AI can make a significant impact. By using Machine Learning algorithms, AI can analyze vast amounts of data, such as customer demographics, online behavior, and purchase history, to generate valuable insights and help marketers make data-driven decisions.

For example, AI can help marketers identify high-value customer segments, predict customer lifetime value, and optimize marketing campaigns for better ROI. Additionally, AI-powered tools can analyze ad performance data to provide recommendations on ad creative, targeting, and bidding strategies, leading to more effective and cost-efficient marketing efforts.

AI-based tools are transforming the way marketers approach A/B testing and campaign optimization, particularly for paid search campaigns. By leveraging Machine Learning algorithms, these tools enable businesses to save time, reduce costs, and create high-performing ad content while maximizing ROI. Let's explore how this process works and provide some examples of its application.

A/B testing is a popular marketing technique used to compare two or more variations of an advertisement or webpage to determine which one performs better. Traditionally, this process could be time-consuming and resource-intensive, as marketers had to

manually create multiple versions of an ad, analyze performance data, and make adjustments accordingly.

AI-based tools simplify and streamline A/B testing by using Machine Learning algorithms to analyze ad performance data in real-time and automatically adjust ad elements, such as headlines, images, or call-to-action buttons, to improve performance. For example, an AI-powered A/B testing tool might analyze click-through rates, conversion rates, or other performance metrics to identify the most effective ad variation and allocate more budget to that version.

Another advantage of AI-based tools in paid search campaigns is the ability to dynamically allocate budgets across channels to maximize ROI. Instead of using a fixed budget for each channel or campaign, AI-driven platforms can analyze performance data in real-time and distribute budget accordingly. For instance, if an AI tool detects that a particular campaign is generating a high ROI, it may automatically shift budget from lower-performing campaigns to further invest in the successful one.

This dynamic budget allocation allows marketers to make the most of their advertising spend and ensure that they're investing in the most effective channels and campaigns.

AI-based tools can also optimize thousands of micro-campaigns every minute, making minor adjustments to targeting, bidding, or ad creative to improve performance. This level of granular optimization would be virtually impossible for human marketers to achieve manually, given the vast number of variables and data points involved.

Furthermore, AI-powered platforms can use Machine Learning algorithms to identify and expand on successful strategies, such as targeting new keywords, audiences, or geographic locations that demonstrate high potential for ROI.

[Slide 9: Retaining Customer Behavior Data for Future Targeting]

By leveraging AI-based tools, marketers can also collect and analyze valuable data on customer behavior, which can be used for future targeting efforts. For example, an AI-driven platform might analyze browsing and purchase history, demographic information, or online behavior patterns to create detailed customer profiles. These profiles can then be used to inform future marketing campaigns, ensuring that ads are reaching the right audience and driving ROI.

Before we wrap up, let's touch on a few more ways AI is being utilized in the marketing field. AI can be used for sentiment analysis, where Machine Learning algorithms analyze customer feedback, reviews, and social media comments to gauge public opinion and sentiment towards a brand, product, or campaign. This helps marketers understand customer perceptions and adjust their strategies accordingly.

AI can also be used for predictive analytics, allowing marketers to forecast future trends, consumer behavior, and market conditions. This enables businesses to stay ahead of the curve and adapt their marketing strategies proactively.

In conclusion, AI is revolutionizing the world of marketing, offering significant benefits across various aspects of the field, from enhancing customer engagement to personalizing content, and improving marketing analytics. As a business executive,

understanding and leveraging the power of AI in marketing is crucial for staying competitive, streamlining operations, reducing costs, and providing an exceptional customer experience.

[A/B] AI in Supply Chain

Let's discuss the various applications of AI in supply chains, such as improving efficiency, enhancing forecasting, and optimizing logistics.

In the realm of supply chain management, AI has the potential to significantly improve efficiency, reduce costs, and enhance the overall effectiveness of supply chain operations.

One of the most valuable applications of AI in the supply chain is to improve visibility and responsiveness across the entire chain. This is achieved by gathering and analyzing past and current data from multiple connected devices and systems, such as SRM (Supplier Relationship Management) software, CRM (Customer Relationship Management) and ERP (Enterprise Resource Planning) systems, and business intelligence solutions.

Let's dive deeper into some specific applications of AI in supply chain management.

First up, let's discuss how AI can help improve supply chain efficiency. Supply chain operations involve a wide range of tasks, such as procurement, production, warehousing, transportation, and more. AI can help streamline and optimize these processes, leading to cost savings and improved performance.

For instance, AI-powered systems can analyze historical and real-time data to optimize production schedules, taking into account factors such as demand fluctuations, equipment downtime, and workforce availability. This can help businesses reduce production lead times, minimize inventory levels, and better meet customer expectations.

Another area where AI can improve efficiency is in warehouse management. By using robotics and computer vision technology, AI can automate tasks such as picking and packing, leading to increased productivity, reduced labor costs, and fewer errors.

Now, let's move on to demand forecasting, a critical aspect of supply chain management. Accurate demand forecasts enable businesses to plan their inventory levels, production schedules, and logistics more effectively. AI can significantly improve demand forecasting by using Machine Learning algorithms to analyze vast amounts of data, such as historical sales data, market trends, and external factors like weather or seasonal events.

By leveraging AI, businesses can generate more accurate and granular demand forecasts, helping them to better anticipate fluctuations in customer demand and optimize their supply chain operations accordingly.

Another area where AI can make a significant impact is in logistics. Managing the transportation of goods from production facilities to distribution centers and customers is a complex and costly process. AI can help optimize logistics by analyzing data on factors such as delivery routes, fuel consumption, vehicle maintenance, and traffic conditions to identify opportunities for cost savings and efficiency improvements.

For example, AI-powered systems can optimize delivery routes to minimize fuel consumption and travel time while ensuring timely deliveries. Additionally, AI can be used to predict vehicle maintenance needs, allowing businesses to schedule servicing proactively and minimize downtime.

One of the critical roles of AI in supply chain management is ensuring on-time product delivery. Advanced AI-based GPS tools enable better navigation and route optimization for transportation and fleet management. These tools process driver, vehicle, and customer data through Machine Learning algorithms to determine the most efficient route for product delivery.

Another way AI increases fleet efficiency is by monitoring vehicle performance and predicting maintenance needs. By analyzing data from sensors and onboard diagnostics, AI can identify potential issues before they become critical, reducing downtime and maintenance costs.

As we come to the end of this chapter, let's explore a few more examples of how AI is transforming supply chain management.

One notable use case is supplier risk assessment. By leveraging Machine Learning algorithms, AI can examine a wide range of supplier data, including financial stability, past performance, and quality metrics. This comprehensive analysis allows businesses to pinpoint potential risks and make well-informed procurement decisions.

For example, an AI system might flag a supplier with declining financial health or consistently late deliveries, prompting the company to reevaluate its relationship or search for alternative suppliers.

Another significant application of AI in supply chain management is inventory management. Maintaining optimal inventory levels is essential to prevent stock outs or overstocking, both of which can be costly. AI can help businesses achieve this balance by examining factors such as customer demand, production lead times, and

carrying costs, along with external factors like market trends or seasonal fluctuations. With this information, AI-powered systems can generate data-driven recommendations for optimal inventory levels, ensuring that businesses can efficiently meet customer demand while minimizing costs.

Additionally, AI can assist in creating dynamic inventory replenishment strategies, adjusting stock levels based on real-time demand patterns or supply chain disruptions. This proactive approach helps businesses maintain the right amount of inventory at the right time, improving overall supply chain efficiency and responsiveness.

In summary, AI is transforming the landscape of supply chain management by delivering substantial advantages in numerous areas, such as increasing efficiency, improving demand forecasting, optimizing logistics, and beyond.

SECTION 13: THE FUTURE OF AI

***[A/B] Overview of 13 Emerging AI
Technologies***

In this chapter, we'll explore an overview of 13 emerging AI technologies that are shaping the future of various industries. By staying informed about these trends, organizations can prepare for and capitalize on the potential of these transformative technologies.

Let's start with Generative AI, which refers to AI models that can create new data, such as images, text, or music, by learning patterns from existing data. Applications include generating artwork, producing realistic video game environments, and creating personalized marketing content.

Number 2 is Digital twinning, which involves creating a digital replica of a physical object or system, allowing for real-time monitoring and optimization. Industries such as manufacturing, healthcare, and smart cities can benefit from digital twins to improve efficiency and predict potential issues.

Moving on to number 3, which is Private computing, which focuses on the development of AI techniques that protect user data while still enabling Machine Learning. Technologies such as federated learning and differential privacy help to balance AI's capabilities with the need for privacy and security.

Number 4 is Personalization and Custom AI Models. AI is increasingly being used to create highly personalized experiences and products for targeted sub-groups of users. Custom AI models can be tailored to specific industries or use cases, resulting in more accurate predictions and improved user satisfaction.

Number 5 is better human-in-the-loop training techniques involve incorporating human feedback into the AI training process to improve model performance. This approach can help create more

accurate and reliable AI systems by leveraging human expertise to guide model development.

Number 6: Low-code and no-code AI platforms enable non-experts to build and deploy AI applications without extensive programming knowledge. This democratizes access to AI, empowering more individuals and organizations to harness its capabilities.

Number 7: Zero-shot and few-shot learning techniques enable AI models to learn new tasks with little or no training data. This can dramatically reduce the amount of data and time required to train AI systems, making them more adaptable and efficient.

Number 8: AI replicas can generate realistic video, voice, or text content by mimicking real-world entities. Examples include deepfake technology, AI-generated voices for virtual assistants, and AI-generated text for chatbots or news articles.

Number 9: AI is increasingly being used to automate and optimize business processes, such as customer support, inventory management, and financial analysis. By streamlining these processes, organizations can increase efficiency, reduce costs, and improve decision-making.

Number 10: As AI becomes more prevalent, countries and organizations are working on developing regulations and ethical guidelines for AI technologies. This includes addressing issues such as bias, fairness, transparency, and data privacy.

Number 11: AI-driven customer analytics and behavior analysis are enabling the creation of highly personalized products and services. By understanding individual preferences and needs, businesses can deliver more relevant and engaging experiences.

Number 12: AI is being used in education to enhance learning experiences, improve assessment methods, and personalize learning paths. Applications include adaptive learning platforms, AI-driven tutoring systems, and virtual reality learning environments.

And lastly number 13, which is AI-enabled hardware, such as cameras, monitors, and mobile devices, integrates AI capabilities directly into the devices we use daily. These advancements can lead to more intelligent and efficient devices, enabling new features and applications across various industries.

[A/B] Preparing for the Future of AI

Let's discuss how organizations can prepare for and capitalize on the potential of emerging AI technologies.

Stay informed. Keep up-to-date with the latest AI research, trends, and developments by following industry news, attending conferences, and engaging with AI communities. Staying informed will help you make better decisions regarding the adoption of new AI technologies in your organization.

Next Invest in AI talent by hiring skilled professionals and providing ongoing training to your existing workforce. Building a strong AI team and infrastructure will enable your organization to capitalize on new AI technologies and stay ahead of the competition.

Also Encourage a culture of experimentation and adaptability within your organization. Be open to testing new AI technologies and adapting your strategies as needed. This mindset will help your organization stay agile in the face of rapidly evolving AI technologies.

Moving on to Collaborating with other organizations, academic institutions, and AI vendors to stay informed about new technologies and learn from shared experiences. Forming partnerships can help you access resources, knowledge, and expertise that will accelerate your AI initiatives.

Lastly, ensure that your organization adheres to ethical guidelines and best practices when implementing AI technologies. This includes addressing issues such as bias, fairness, transparency, and data privacy. Responsible AI implementation will help build trust and minimize potential negative impacts on society.

In conclusion, By staying informed and preparing for these trends, organizations can capitalize on the transformative potential of AI technologies.

***[A/B] Where We See AI Being Used/Deployed
and How?***

In this chapter, we'll explore various platforms and environments where AI can be deployed and used. We'll discuss local deployment, AI in the cloud, browser-based AI, AI on the edge, mobile devices, IoT devices, and Raspberry Pi implementations.

Let's start with local deployment, which refers to running AI models on a user's local machine or on-premises server. This approach offers advantages such as low latency, data privacy, and control over resources. Local deployment is often used in industries where data privacy and security are paramount, such as healthcare and finance.

Next we have Cloud-based AI deployment, which leverages the power of cloud computing to run AI models. This approach offers scalability, flexibility, and cost savings, as resources can be easily adjusted based on demand. Cloud-based AI is popular in industries such as e-commerce, marketing, and customer support, where large-scale data processing and real-time insights are critical.

Moving on to browser-based AI deployment, which utilizes web technologies like JavaScript and WebAssembly to run AI models directly in the user's browser. This approach enables AI applications to run without requiring installation, offering improved user experience and accessibility. Browser-based AI is often used for applications like image editing, language translation, and content generation.

In terms of Edge AI deployment, this involves running AI models on devices located close to the data source, such as cameras, sensors, or other IoT devices. This approach offers reduced latency, increased data privacy, and reduced bandwidth requirements.

Edge AI is commonly used in applications like autonomous vehicles, smart cities, and industrial automation.

Also we have Deploying AI models on mobile devices, which allows for on-device processing and real-time insights. This approach offers benefits such as improved user experience, reduced data transmission, and increased data privacy. AI on mobile devices is popular in applications like facial recognition, voice assistants, and augmented reality.

In terms of AI deployment in IoT meaning internet of things devices, this enables smart decision-making and automation in connected devices. This approach offers improved efficiency, increased functionality, and better user experiences. AI in IoT devices is commonly used in applications like smart home automation, energy management, and health monitoring.

Lastly, we have Raspberry Pi, which is a low-cost, compact computing platform that can be used for AI deployment. Raspberry Pi-based AI applications offer benefits such as affordability, accessibility, and educational value. AI on Raspberry Pi is often used in do it yourself projects, robotics, and educational settings.