



INNOVATE

DATA AND AI/ML EDITION

22 February 2023

Adding machine learning to your software engineering toolbelt

Matthew Coles

Principal Engineer

Amazon Web Services

Agenda

- Understanding the difference between AI, ML and deep learning
- Exploring development workflow similarities
- How to easily get started with AI, ML
- Practical perspective/demo of AI-powered coding companion
- Other AI tools for developers
- Building your own ML models
- Resources to get started

Let's get started!



Understanding the difference between AI, ML and deep learning



Artificial intelligence

Any technique that enables computers to mimic human intelligence using logic, if-then statements, and ML (including deep learning)



Machine learning

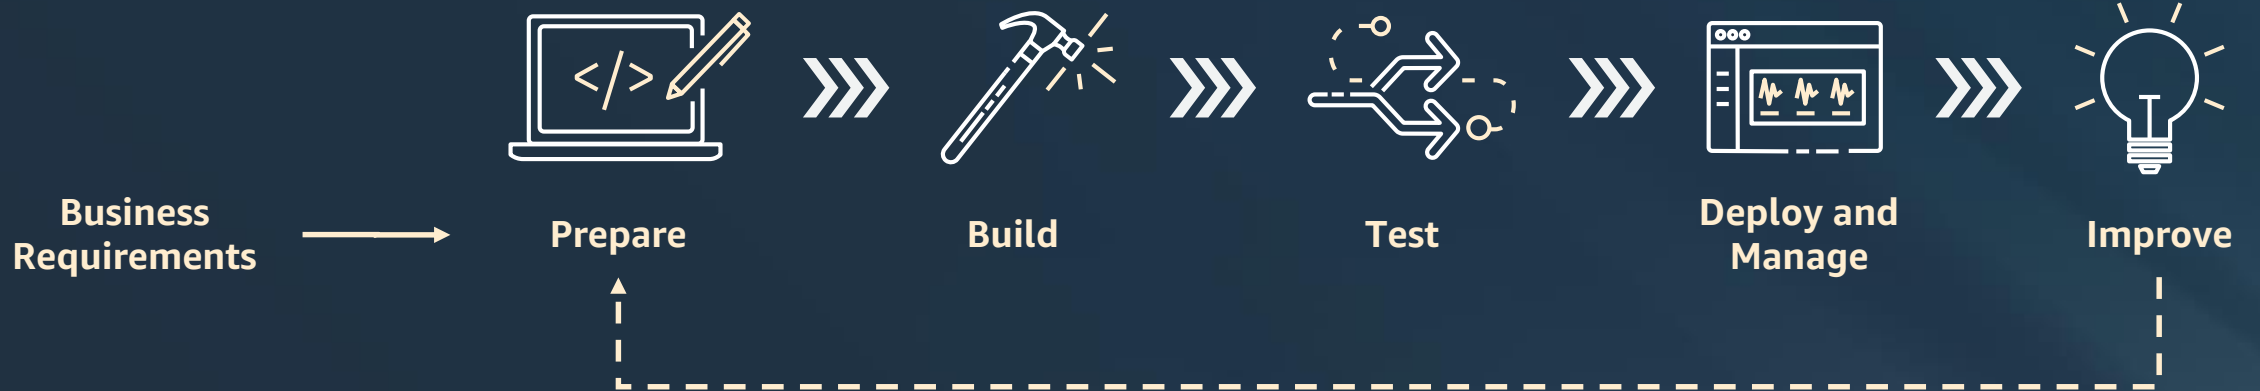
Subset of AI that uses machines to search for patterns in data to build logic models automatically



Deep learning

Subset of ML composed of deeply multi-layered neural networks that perform tasks like speech and image recognition

Exploring development workflow similarities



Exploring development workflow similarities

PREPARE →

SageMaker Ground Truth

Label training data for machine learning

SageMaker Data Wrangler

Aggregate and prepare data for machine learning

SageMaker Processing

Built-in Python, BYO R/Spark

SageMaker Feature Store

Store, update, retrieve, and share features

SageMaker Clarify

Detect bias and understand model predictions

BUILD →

SageMaker Studio Notebooks

Jupyter notebooks with elastic compute and sharing

Built-in and Bring your-own Algorithms

Dozens of optimized algorithms or bring your own

Local Mode

Test and prototype on your local machine

SageMaker Autopilot

Automatically create machine learning models with full visibility

SageMaker JumpStart

Pre-built solutions for common use cases

TRAIN & TUNE →

Managed Training

Distributed infrastructure management

SageMaker Experiments

Capture, organize, and compare every step

Automatic Model Tuning

Hyperparameter optimization

Distributed Training Libraries

Training for large datasets and models

SageMaker Debugger

Debug and profile training runs

Managed Spot Training

Reduce training cost by 90%

DEPLOY & MANAGE →

Managed Deployment

Fully managed, ultra low latency, high throughput

Kubernetes & Kubeflow Integration

Simplify Kubernetes-based machine learning

Multi-Model Endpoints

Reduce cost by hosting multiple models per instance

SageMaker Model Monitor

Maintain accuracy of deployed models

SageMaker Edge Manager

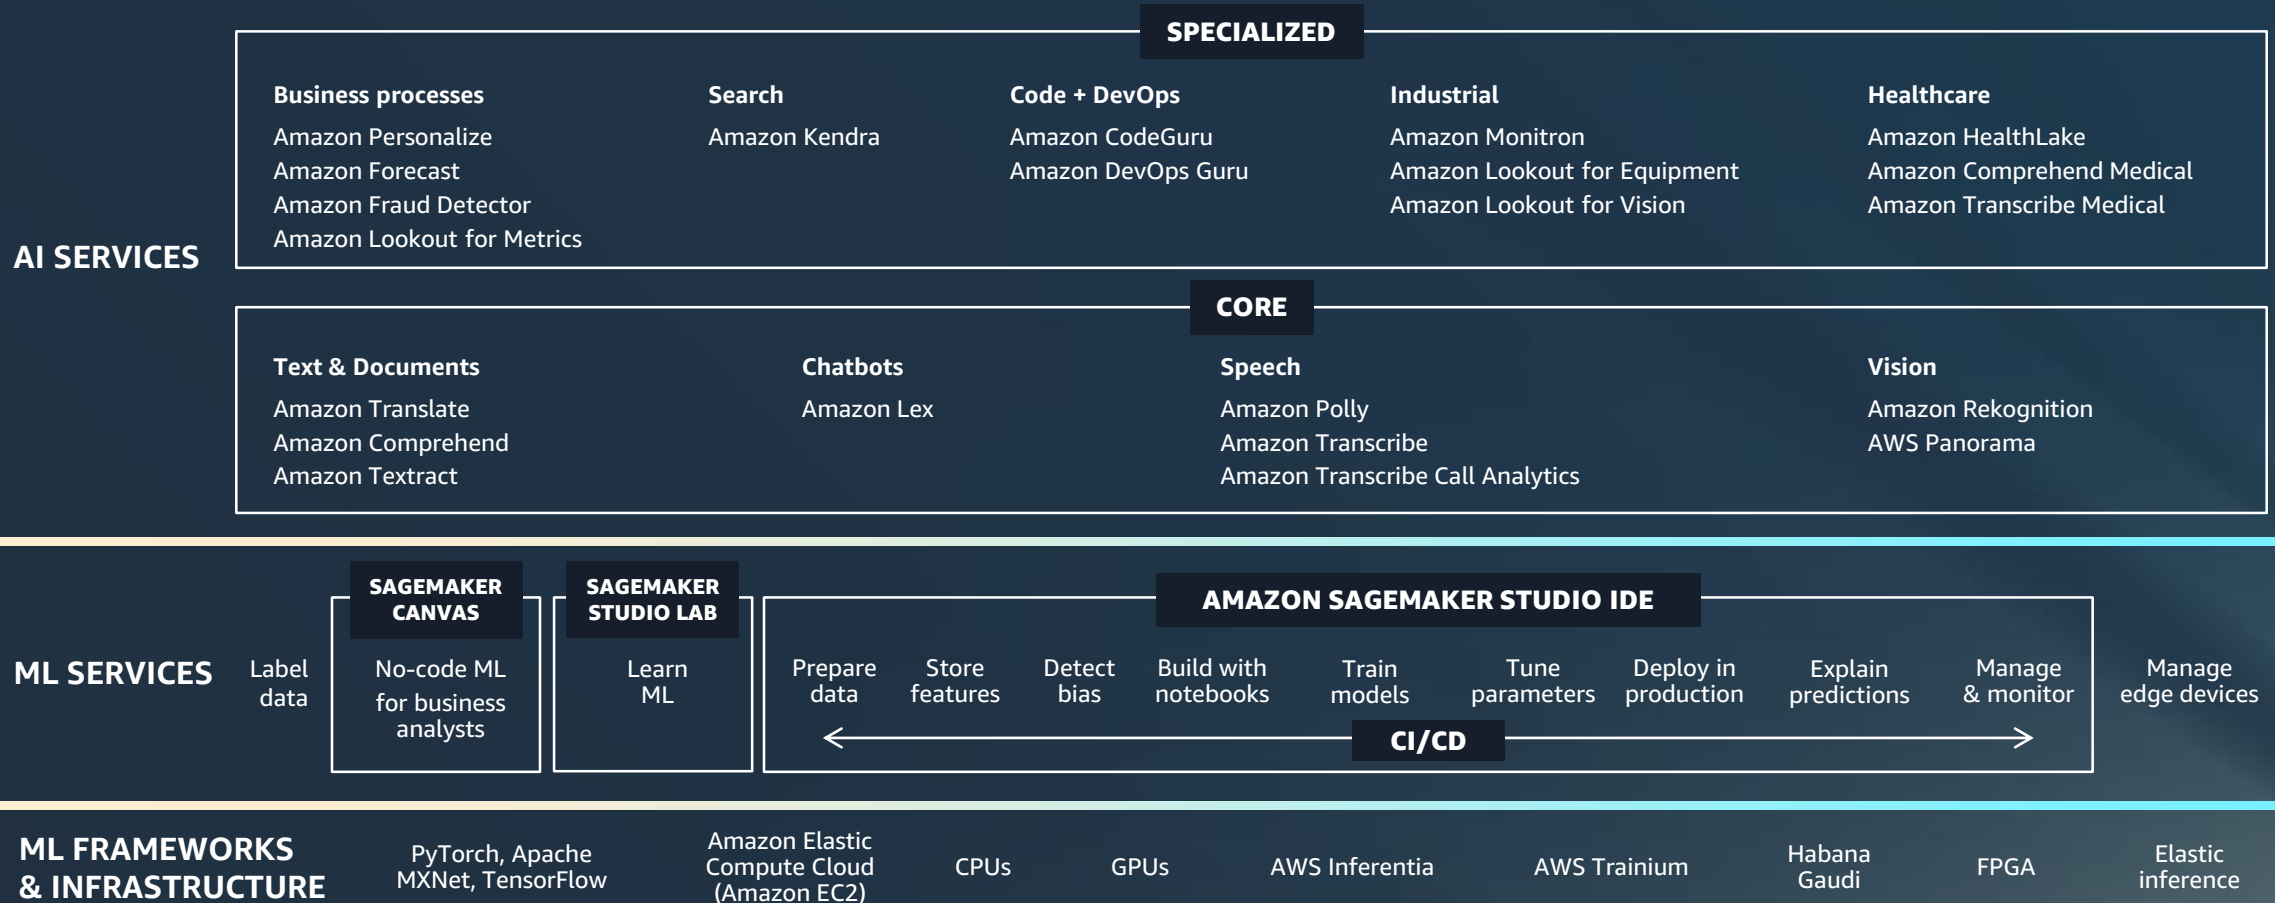
Manage and monitor models on edge devices

SageMaker Pipelines

Workflow orchestration and automation

How to easily get started with ML

Broadest and most complete set of machine learning capabilities



**Dive deep / practical perspective /
demo AI-powered coding
companion**



Exploring an AI-powered coding companion



Multiple programming languages



Multiple frameworks and libraries



Numerous cloud services and APIs



Identifying relevant best practices

Amazon CodeWhisperer: AI-powered coding companion

Provides code recommendations in natural language based on contextual information like prior code and comments

GENERATES:

- Entirely new code based on context
- Code from plain English comments
- Complete functions

Available in all major integrated development environments (IDEs) as an extension

```
# Write a function to upload a file to S3.
def upload_file_to_s3(file_name, bucket_name, object_name):
    """
    Uploads a file to an S3 bucket

    :param file_name: File to upload
    :param bucket_name: Bucket to upload to
    :param object_name: S3 object name. If none then file_name is used
    :return: True if file was uploaded, else False
    """

    # Upload the file
    s3_client = boto3.client('s3',
                             aws_access_key_id=AWS_ACCESS_KEY_ID,
                             aws_secret_access_key=AWS_SECRET_ACCESS_KEY,
                             region_name=AWS_REGION_NAME)

    try:
        s3_client.upload_file(file_name, bucket_name, object_name)
        print(f'File {file_name} uploaded to S3 bucket {bucket_name} as {object_name}')
        return True
    except FileNotFoundError:
        print(f'File {file_name} not found')
```

Amazon CodeWhisperer: AI-powered coding companion

Develop next-gen applications

Accelerate frontend and backend development by empowering developers with automatic code recommendations

Use ML in your applications

Save time and effort by using Amazon CodeWhisperer to generate code to build and train your ML models

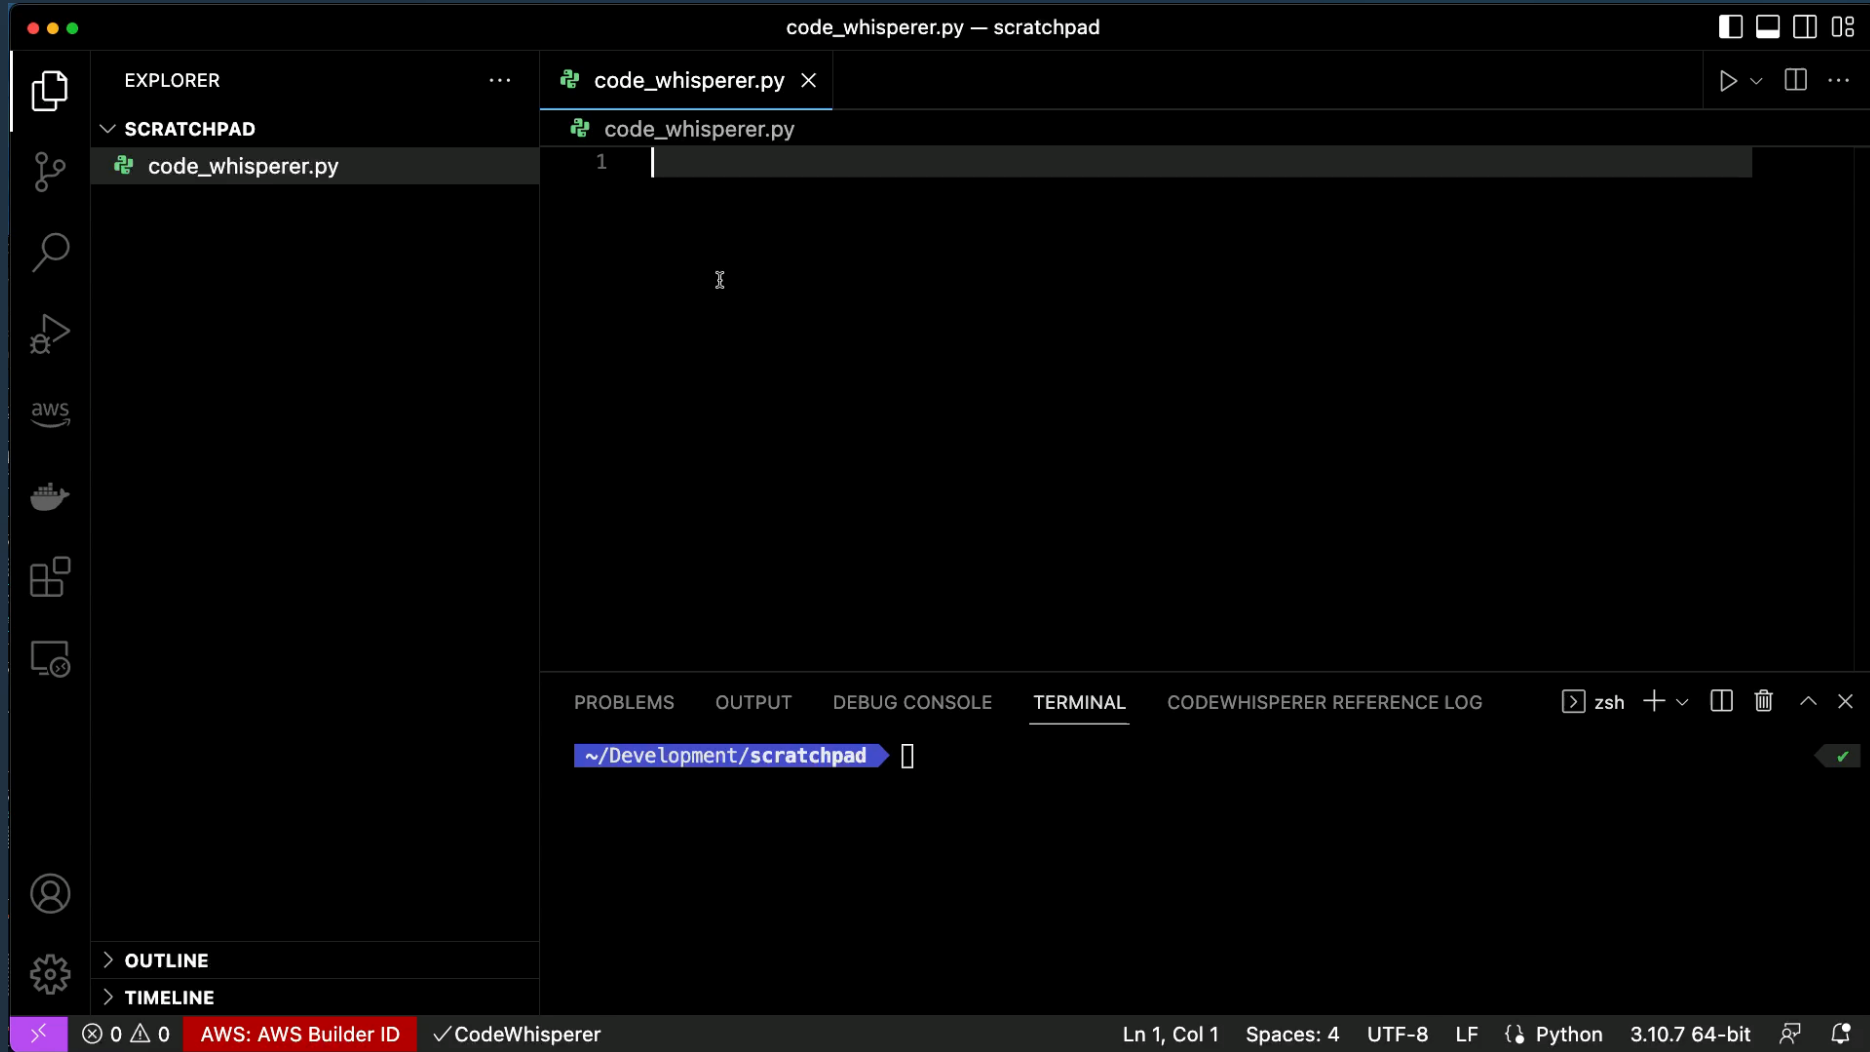
Build applications using AWS services

Speed up the development process with code recommendations for AWS APIs across the most popular services, including Amazon EC2, AWS Lambda, and Amazon S3

Automate unit test generation

- Offload writing repetitive unit test code
- Based on natural language comments, Amazon CodeWhisperer automatically recommends unit test code that matches your implementation code

Amazon CodeWhisperer: ML-powered coding companion



Other AI tools for developers

Amazon Rekognition



OBJECT, SCENE, AND ACTIVITY



CUSTOM LABELS



CONTENT MODERATION



TEXT



PPE DETECTION



FACE DETECTION AND ANALYSIS



FACE COMPARE & SEARCH



CELEBRITY RECOGNITION



VIDEO SEGMENTS



LIVE STREAM VIDEO



PATHING

Amazon Rekognition

The screenshot shows the AWS Management Console Home page. At the top, there is a navigation bar with the AWS logo, a 'Services' menu, a search bar, and a user profile for 'AdministratorAccess/macols-lsengard @ 6694-1750-2288'. The main content area is divided into several sections:

- Console Home**: Includes a 'Reset to default layout' button and an '+ Add widgets' button.
- Recently visited**: A list of services visited recently, including Amazon Rekognition, Amazon Lex, Amazon Textract, Amazon Transcribe, Device Farm, DynamoDB, CodePipeline, CodeStar, IAM, S3, Cloud9, VPC, and CloudWatch. A 'View all services' link is at the bottom.
- Welcome to AWS**: A sidebar with three links: 'Getting started with AWS', 'Training and certification', and 'What's new with AWS?'. Each link has a brief description and an external link icon.
- AWS Health**: Shows 'Open issues' as 0 for the 'Past 7 days' and 'Scheduled changes'.
- Cost and usage**: Displays 'Current month costs' as \$24.94 and 'Forecasted month end costs' as \$37.12. It also lists 'Top costs for current month' with a table:

Service	Cost
AWS Glue	\$7.33
EC2 - Other	\$5.31

At the bottom, there is a footer with 'Feedback', a link to 'Unified Settings', and copyright information: '© 2023, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences'.



Amazon Textract - Reading documents

Accuracy of forms and tables allows customers to reach **80–90% automation levels** on documents that previously needed manual processing.

INPUT → **OUTPUT**

Employment Application

Application Information

Full Name: Jane Doe

Phone Number: 555-0100

Home Address: 123 Any Street, Any Town, USA

Mailing Address: same as above

How did you hear about this position?

Job fair Website Company Employee

Previous Employment History

Start Date	End Date	Employer Name	Position Held	Reason for leaving
1/15/2009	6/30/2011	Any Company	Assistant baker	relocated
7/1/2011	8/10/2013	Example Corp.	Baker	better opp.
8/15/2013	Present	Any Company	head baker	N/A, current

Raw Text

Employment Application Application Information Full Name: Jane Doe

Phone Number: 555-0100 Home Address: 123 Any street, Any Town. USA

Forms

Job fair Website

SELECTED NOT_SELECTED

Table Recognition

7/1/2011 8/10/2013 Example Corp. Baker better opp.

Amazon Textract - Reading documents

The screenshot shows the AWS Management Console Home page. At the top, there is a navigation bar with the AWS logo, a 'Services' menu, a search bar, and a user profile dropdown. The main content area is divided into several sections:

- Console Home**: Includes 'Reset to default layout' and '+ Add widgets' buttons.
- Recently visited**: A list of services including Amazon Textract, Amazon Transcribe, Amazon Rekognition, Device Farm, DynamoDB, CodePipeline, CodeStar, IAM, S3, Cloud9, VPC, CloudWatch, and Amazon Translate.
- Welcome to AWS**: Contains links for 'Getting started with AWS', 'Training and certification', and 'What's new with AWS?'.
- AWS Health**: Shows 'Open issues' as 0 and 'Scheduled changes'.
- Cost and usage**: Displays 'Current month costs' as \$24.94 and 'Top costs for current month' with a table:

Service	Cost
AWS Glue	\$7.33
EC2 - Other	\$5.31

At the bottom, there is a footer with 'Feedback', 'Looking for language selection? Find it in the new Unified Settings', and copyright information for 2023 Amazon Web Services, Inc. or its affiliates.



Amazon Transcribe



Punctuation,
capitalization



Speaker labelling/channel
identification



Word-level
time stamps



Word-level
confidence scores



Alternative
transcriptions



Custom
vocabulary




Vocabulary
Filtering



Custom Language Models
(CLM)



Language ID



PCI/PII content redaction

Amazon Transcribe

aws Services Search [Option+S] Sydney AdministratorAccess/macols-lsengard @ 6694-1750-2288

Console Home [Info](#)

[Reset to default layout](#) [+ Add widgets](#)

Recently visited [Info](#)

- Amazon Transcribe
- Amazon Textract
- Amazon Rekognition
- Device Farm
- DynamoDB
- CodePipeline
- CodeStar
- IAM
- S3
- Cloud9
- VPC
- CloudWatch
- Amazon Translate

[View all services](#)

Welcome to AWS

- [Getting started with AWS](#)
Learn the fundamentals and find valuable information to get the most out of AWS.
- [Training and certification](#)
Learn from AWS experts and advance your skills and knowledge.
- [What's new with AWS?](#)
Discover new AWS services, features, and Regions.

AWS Health [Info](#)

Open issues: **0** (Past 7 days)

Scheduled changes: **0**

Cost and usage [Info](#)

Current month costs: **\$24.94**

Forecasted month end costs: **\$77.12**

Top costs for current month

	AWS Glue	\$7.33
	EC2 - Other	\$5.31

[Feedback](#) [Looking for language selection? Find it in the new Unified Settings](#)

© 2023, Amazon Web Services, Inc. or its affiliates. [Privacy](#) [Terms](#) [Cookie preferences](#)

Amazon Translate



Broad language coverage

4970 X<>Y translation combinations
Language identification



Customizable translation

Customize translation output with custom terminologies and active custom translation



Low latency

< 150ms/sentence on average
< 80ms/conversational/short form



Document batch translation

Translates large set of docx, pptx, xlsx, XML, HTML, and text documents with one API call



Data security

Data ownership
Encryption
Access management



Broad domain coverage

Trained on 11 domains



Broad regional coverage

Available in the following 17 regions



Pay per use

Simple API calls and you pay for what you use

Amazon Comprehend



Documents

Email, chat, social,
phone calls, and more



Amazon Comprehend

Automatically extract
insights from text



Entities
+ Custom entities



Key
Phrases



PII (Personally Identifiable
Information)
Labeling + Detection



Sentiment + Targeted
Sentiment



Document
Classification



Topics



Language



Syntax



Events

Amazon Personalize

User personalization

prime TV shows we think you'll like

Based on titles you have watched and more



Similar items

Customers who viewed this item also viewed

Five similar area rug recommendations are shown in a row. Each item includes a thumbnail image, a title, dimensions, material, color, star rating, and price with a Prime badge.

- Nourison Celestial Modern Abstract Area Rug, 6'7" x 9'7" (7x10), Sealife Multicolor Grey**
★★★★☆ 543
\$192.87 ✓prime
- nuLOOM Traditional Waterfall Vintage Abstract Area Rug, 6' 7" x 9', Blue**
★★★★☆ 7
\$135.99 ✓prime
- Nourison Celestial Modern Abstract Area Rug, 5'3" x 7'3" (5x7), Sealife Multicolor Grey**
★★★★☆ 543
\$98.87 ✓prime
- nuLOOM Alayna Abstract Area Rug, 5' x 7' 5", Blue**
★★★★☆ 71
\$74.89 ✓prime
- Nourison Celestial Modern Abstract Area Rug, 7'10" x 10'6" (8x11), Sealife Multicolor Grey**
★★★★☆ 543
\$250.87 ✓prime

Personalized ranking

Recommended for You in Classics

A carousel of classic audiobook recommendations. The main focus is 'Children Playing Before a Statue of Hercules' by David Sedaris, with a green cover and a statue illustration. Other books visible include 'To Kill a Mockingbird' by Harper Lee, 'Brave New World' by Aldous Huxley, 'The Wealth of Nations' by Adam Smith, and 'Dune' by Frank Herbert.

Amazon Lex

Intent

An action that a user wants performed

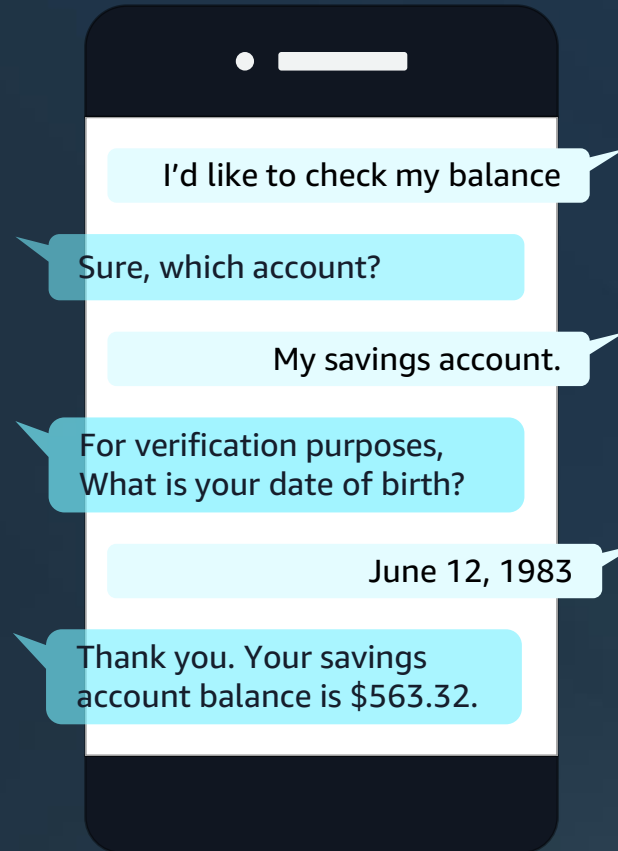
Prompt

Bot message designed to get data needed to complete the requested action

Fulfillment

Completed action based on the user request

● Get account balance



Utterance

User input that triggers an intent

Slot

Input data that is needed to complete the requested action

Getting started with AI guides

- AI guides:

- Amazon CodeWhisperer: https://aws.amazon.com/codewhisperer/#How_to_get_started
- Amazon Rekognition: <https://docs.aws.amazon.com/rekognition/latest/dg/getting-started.html>
- Amazon Textract: <https://docs.aws.amazon.com/textract/latest/dg/getting-started.html>
- Amazon Transcribe: <https://docs.aws.amazon.com/transcribe/latest/dg/getting-started-sdk.html>
- Amazon Comprehend: <https://docs.aws.amazon.com/comprehend/latest/dg/using-api-sync.html>
- Amazon LEX: <https://docs.aws.amazon.com/lex/latest/dg/getting-started.html>
- Amazon Personalize: <https://docs.aws.amazon.com/personalize/latest/dg/getting-started-domain-python.html>
- AWS Amplify: <https://docs.amplify.aws/start/>

- AI workshops:

- Amazon Rekognition: <https://catalog.us-east-1.prod.workshops.aws/workshops/4a2a9a24-071d-4d96-b9be-0cc57b7db434/en-US>
- Amazon Textract: <https://catalog.us-east-1.prod.workshops.aws/workshops/f76c7a73-2080-4670-bd1d-db26a6259142/en-US>
- Amazon Personalize and AWS Amplify: <https://catalog.us-east-1.prod.workshops.aws/workshops/bb080ee8-4722-4290-ac6e-d4cde0a65142/en-US>

AWS Amplify

BUILD



Configure AWS backends fast



Seamlessly connect frontends



CLI



Studio



Libraries

Feature categories



Authentication



Analytics



DataStore



PubSub



Storage



Predictions



API (GraphQL & REST)



Interactions



Functions



Notifications



Geo



Extensibility

Powered by AWS services such as AWS AppSync (GraphQL API) and Amazon Cognito (Authentication)

Building your own ML models



Amazon SageMaker

Integrated Workbench

Capabilities designed specifically for ML, data preparation, experiment management, and workflows

Managed Infrastructure

Designed for ultra low latency and high throughput, automatic scaling, and distributed training

Managed Tooling

Purpose-built from the ground up to work together including auto ML, collaboration, debugger, profiler, bias analyzer, and explainability

The screenshot displays the Amazon SageMaker Studio interface. The main window is titled "SAGEMAKER DEBUGGER" and shows monitoring, profiling, and debugging options. The "Nodes" tab is active, displaying a "Profiling report" with suggestions for improving training time, such as "Choose a larger instance type with more memory" and "Switch from the File to Pipe mode". The "Profiling Rules" panel on the right lists rules like "CPU bottleneck", "Long initialization time", and "Batch size too small". The "Training job summary" table at the bottom provides details about the training job.

Label	Value
Start time	2020-09-10 4:26:50.253
End time	2020-09-10 4:51:59.695
Job duration	1,509.44 seconds
Training loop start	2020-09-10 4:31:32.947
Training loop end	2020-09-10 4:51:59.862

Amazon SageMaker

SageMaker Studio IDE

Label
data



Collect and
prepare data



Store
features



Check
for bias



Visualize in
notebooks



Pick
algorithm



Train
models



Tune
parameters



Deploy in
production



Manage
and monitor



CI/CD

MODEL MANAGEMENT FOR EDGE DEVICES

Getting started with machine learning

ML guide:

Amazon Sagemaker: <https://docs.aws.amazon.com/sagemaker/latest/dg/gs.html>

ML workshop:

Amazon Sagemaker Immersion Day: <https://catalog.us-east-1.prod.workshops.aws/workshops/63069e26-921c-4ce1-9cc7-dd882ff62575/en-US>

AWS Machine learning path: <https://aws.amazon.com/training/learn-about/machine-learning/>



In summary



Visit the Data & AI/ML resource hub

Dive deeper into these resources, get inspired and learn how you can use AI and machine learning to accelerate your business outcomes.

- 6 steps to machine learning success e-book
- 7 leading machine learning use cases e-book
- Machine learning at scale e-book
- Achieving transformative business results with machine learning e-book
- Tackling our world's hardest problems with machine learning e-book
- Accelerating machine learning innovation through security e-book
- ... and more!



<https://bitly.co/FqdC>

[Visit resource hub](#)



AWS Training and Certification

Access the AI & ML learning plan courses built by AWS experts on AWS Skill Builder

- Get started with digital self-paced, on-demand training and ramp-up guides to help you grow your technical skills
- Learn how to apply machine learning, artificial intelligence, and deep learning to unlock new insights and value in your role
- Take the steps today, towards validating your expertise with an AWS Certified Machine Learning – Specialty Certification



<https://bit.ly/3FnxDH7>

Learn your way [explore.skillbuilder.aws](https://skillbuilder.aws) »



Thank you for attending AWS Innovate – Data & AI/ML Edition

We hope you found it interesting! A kind reminder to **complete the survey**.
Let us know what you thought of today's event and how we can improve the event experience for you in the future.



aws-apj-marketing@amazon.com



twitter.com/AWSCloud



facebook.com/AmazonWebServices



youtube.com/user/AmazonWebServices



slideshare.net/AmazonWebServices



twitch.tv/aws

Thank you!

