AI-based Log Analysis using Modern NLP Techniques

Dinu Ionuț-Vlăduț, PhD Candidate
Email: ionut.dinu@unitbv.ro

Work experience:

– 5 years @ Siemens
– 1 year @ Freelance

Universitatea
Transilvania
din Brașov
FACULTATEA DE INGINERIE ELECTRICĂ
ȘI ȘTIINȚA CALCULATOARELOR

1. What is SOC?

- A central team that oversees and manages an organization's security stance.

- This unit usually comprises security experts tasked with detecting, addressing, and reducing security threats.

- In essence, the SOC team ensures that the organization functions securely.

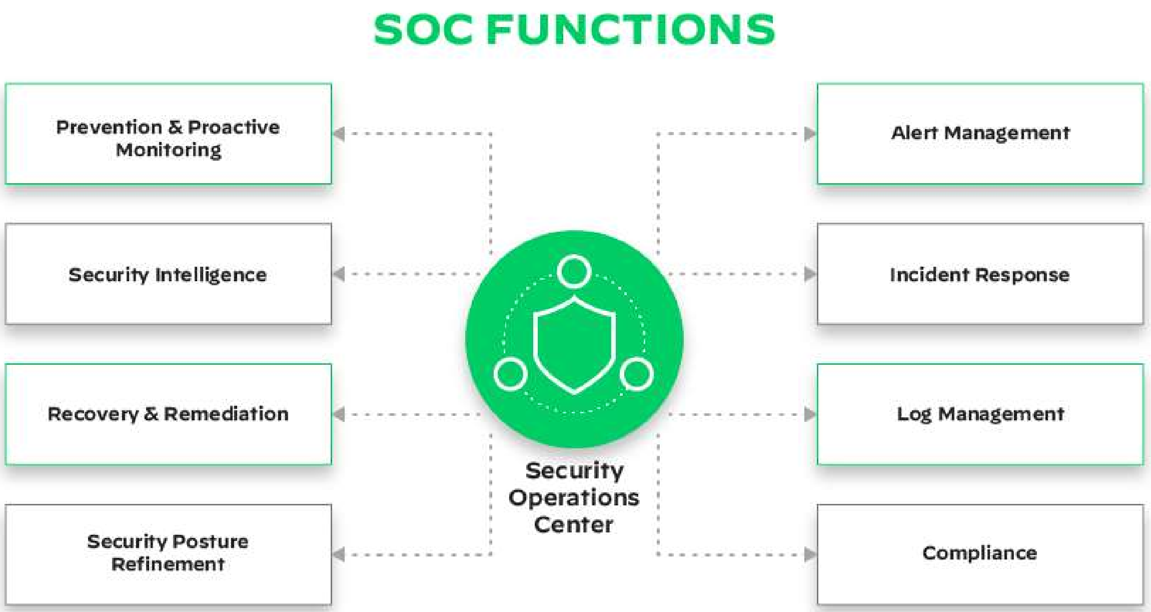

2. How does a SOC help a business?

– Identify an alert as potentially malicious and open an incident.

– Investigate the root cause and impact of the incident.

– Recommend mitigation options to isolate and remove a threat.

– Constantly adapt to process, visibility, and technology to improve in real time as incidents occur.

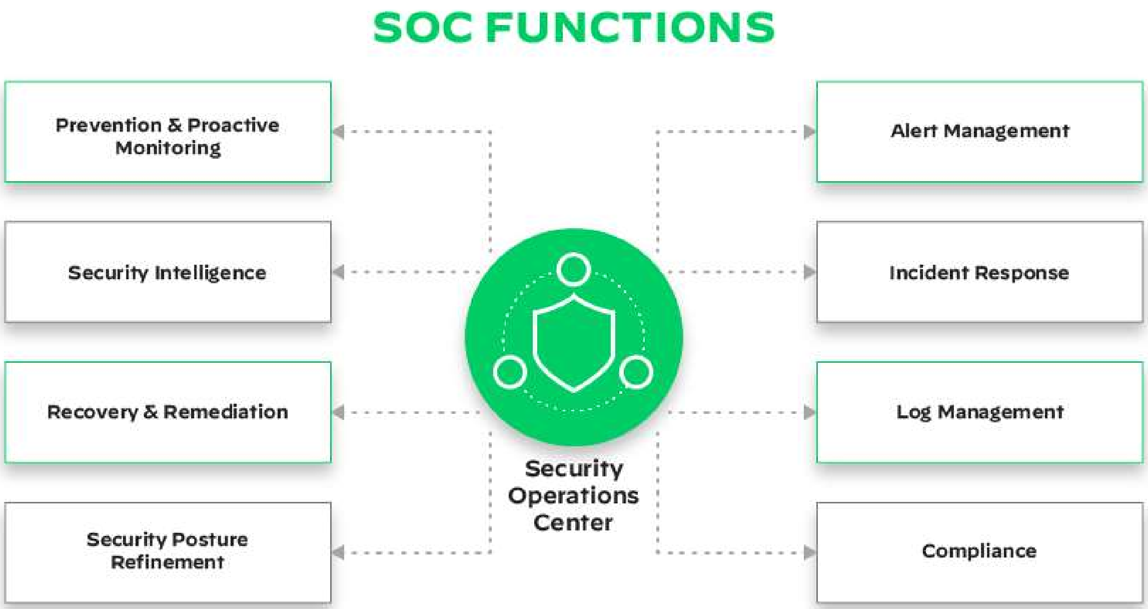

3. Core features of a SOC

– 24/7 monitoring organization networks and systems for compromising activity or other threats.

– High level of detail while assessing a breach, from logs, traffic patterns and firewalls to Intrusion Detection Systems (IDS) and application security.

– Utilizing tools such as Security Information and Event Management (SIEM), Intrusion Detection and Prevention Systems (IDS/IPS), Endpoint Detection and Response (EDR)

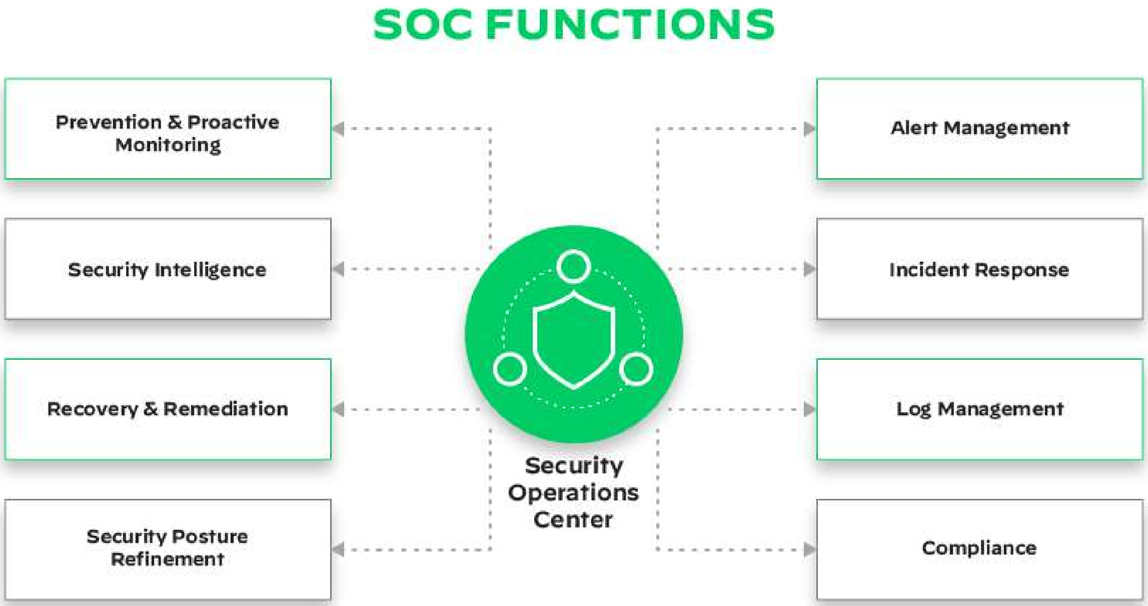
SOC FUNCTIONS

Prevention & Proactive Monitoring

Security Intelligence

Recovery & Remediation

Security Posture Refinement

Security Operations Center

Alert Management

Incident Response

Log Management

Compliance

4. Security Information and Event Management in a SOC

- SIEM solutions help businesses keep track of their applications security information while analyzing their logs in real time, detecting anomalous usage patterns.

- It collects various information from multiple data-sources such as network devices, applications, user activity and it uses multiple analytics techniques to assess the data and provide a quick response.

- SIEMs allow businesses to respond in a fast manner to security incidents and apply the correct countermeasures to mitigate the threat.

- Benefits of using SIEM:
  - Centralized data-source, aggregating and displaying data from multiple sources into a single place
  - Quick response time and detailed analysis
  - High degree of transparency and correlation between multiple types of data

4.1. Splunk Security

- Splunk offers a popular SIEM solution, which handles both security of a application and of a network.
- Ingests and analyzes machine-generated data from multiple data-sources
- Real-time analysis & monitoring
- Event correlation and Risk-based Alerting
- Flexible deployment options (cloud, on-prem, hybrid, etc.)

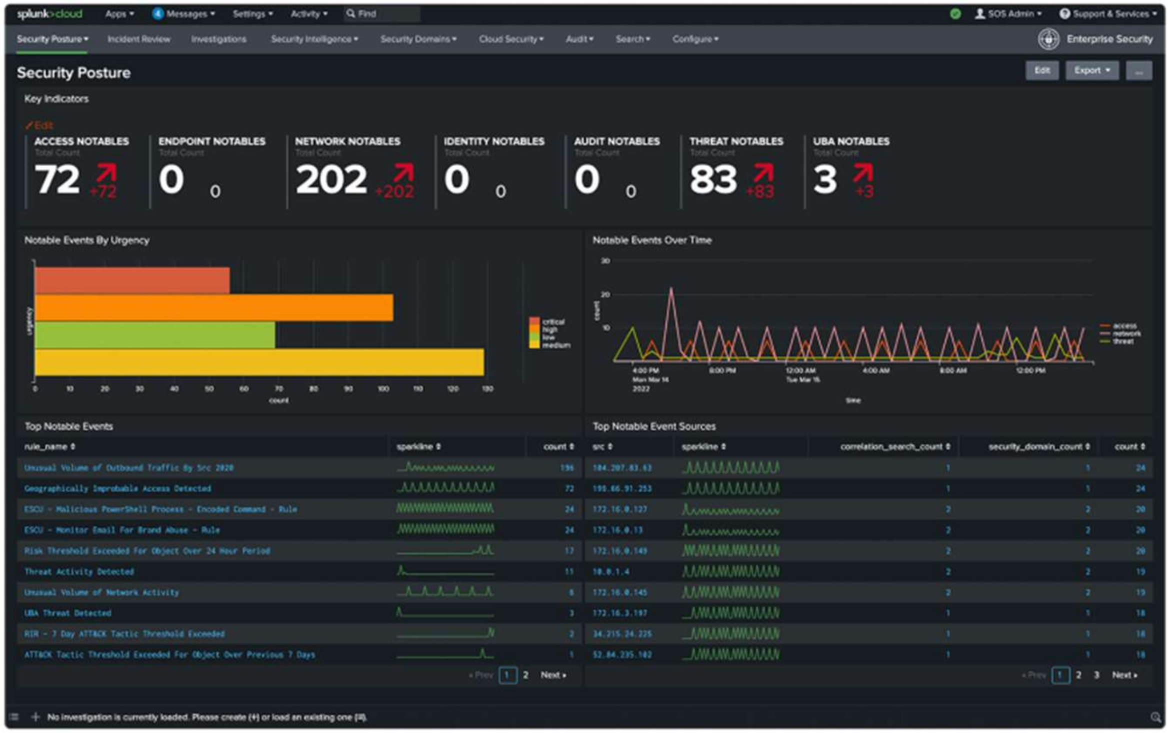

Security Posture

Key Indicators

✓Edit

ACCESS NOTABLES	ENDPOINT NOTABLES	NETWORK NOTABLES	IDENTITY NOTABLES	AUDIT NOTABLES	THREAT NOTABLES	UBA NOTABLES
Total Count	Total Count	Total Count	Total Count	Total Count	Total Count	Total Count
72 ↗ +72	0 0	202 ↗ +202	0 0	0 0	83 ↗ +83	3 ↗ +3

Top Notable Events

rule_name ⇕	sparkline ⇕	count ⇕
Unusual Volume of Outbound Traffic By Src 2020		196
Geographically Improbable Access Detected		72
ESCU - Malicious PowerShell Process - Encoded Command - Rule		24
ESCU - Monitor Email For Brand Abuse - Rule		24
Risk Threshold Exceeded For Object Over 24 Hour Period		17
Threat Activity Detected		11
Unusual Volume of Network Activity		6
UBA Threat Detected		3
RIR - 7 Day ATT&CK Tactic Threshold Exceeded		2
ATT&CK Tactic Threshold Exceeded For Object Over Previous 7 Days		1

« Prev 1 2 Next »

Top Notable Event Sources

src ⇕	sparkline ⇕	correlation_search_count ⇕	security_domain_count ⇕	count ⇕
104.207.83.63		1	1	24
199.66.91.253		1	1	24
172.16.0.127		2	2	20
172.16.0.13		2	2	20
172.16.0.149		2	2	20
10.0.1.4		2	2	19
172.16.0.145		2	2	19
172.16.3.197		1	1	18
34.215.24.225		1	1	18
52.84.235.102		1	1	18

« Prev 1 2 3 Next »

No investigation is currently loaded. Please create (+) or load an existing one (☰).

4.2. LogRhythm

- LogRhythm is an enterprise-class platform that seamlessly combines SIEM, log management, file integrity monitoring and machine analytics with host and network forensics in a unified Security Intelligence Platform.
- Log Management & Data aggregation
- AI-based Event Correlation and Analytics
- Compliance reports

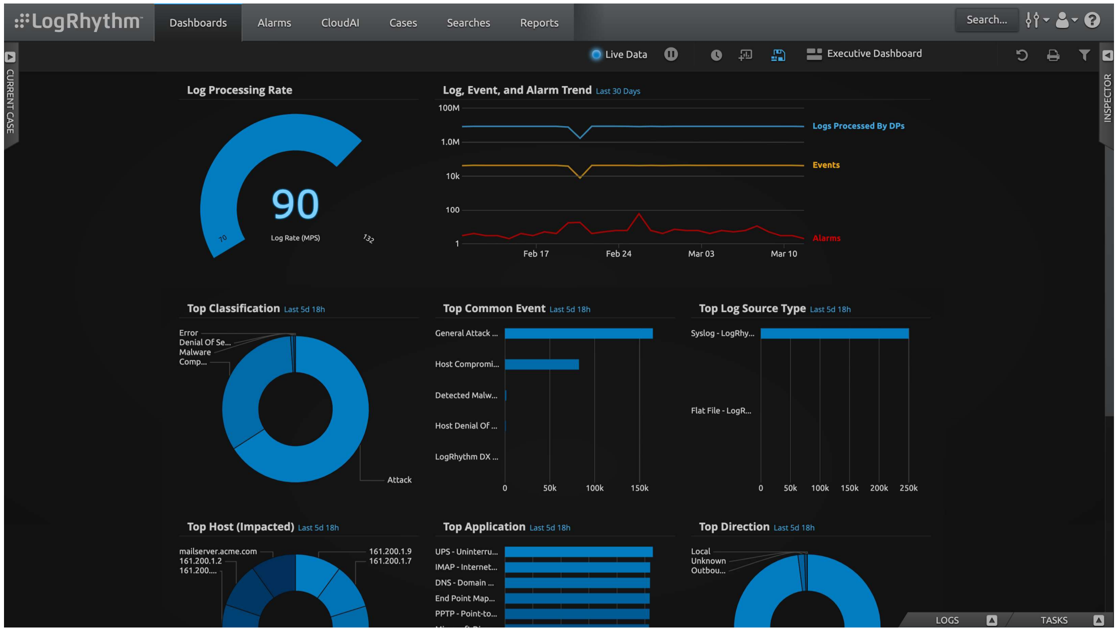

4.3. IBM QRadar SIEM

- It is a modular architecture that facilitates the detection and prioritization of threats. It supports multiple logging protocols and offers various configuration-side options, as well as high-end analytics.
- Advanced Threat Detection
- Incident Investigation & Forensics
- Threat Intelligence Integration
- Enhanced SOC Efficiency

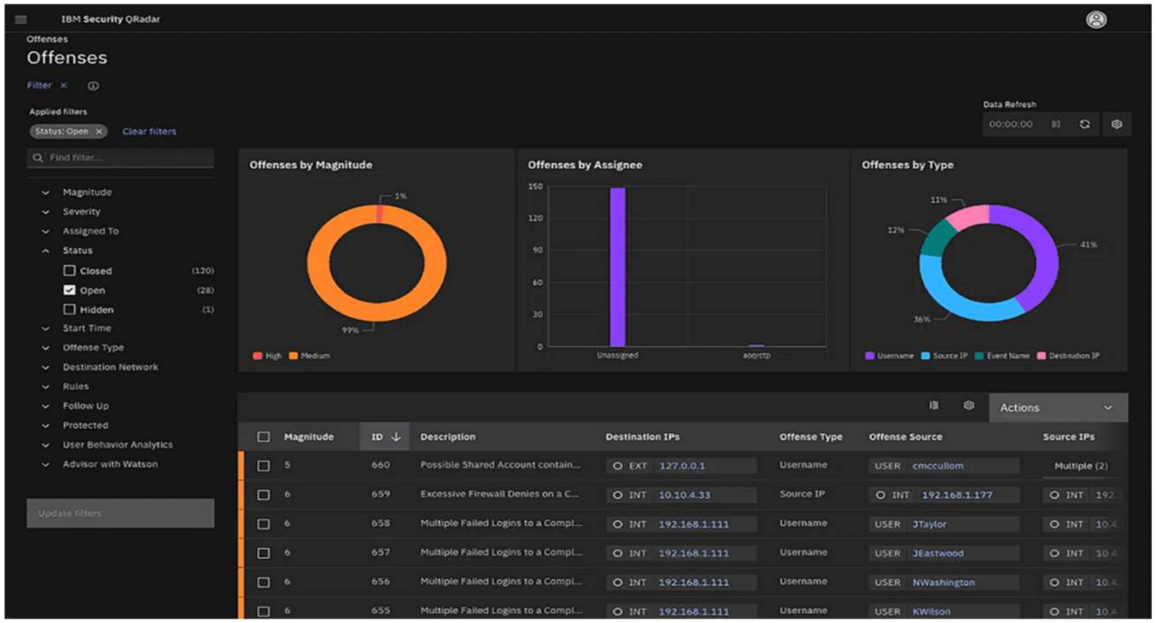

☐	Magnitude	ID ↓	Description	Destination IPs	Offense Type	Offense Source	Source IPs
☐	5	660	Possible Shared Account contain...	○ EXT 127.0.0.1	Username	USER cmccullom	Multiple (2)
☐	6	659	Excessive Firewall Denies on a C...	○ INT 10.10.4.33	Source IP	○ INT 192.168.1.177	○ INT 192.
☐	6	658	Multiple Failed Logins to a Compl...	○ INT 192.168.1.111	Username	USER JTaylor	○ INT 10.4
☐	6	657	Multiple Failed Logins to a Compl...	○ INT 192.168.1.111	Username	USER JEastwood	○ INT 10.4
☐	6	656	Multiple Failed Logins to a Compl...	○ INT 192.168.1.111	Username	USER NWashington	○ INT 10.4
☐	6	655	Multiple Failed Logins to a Compl...	○ INT 192.168.1.111	Username	USER KWilson	○ INT 10.4

5. Endpoint Detection and Response in a SOC

- Endpoint Detection and Response (EDR) is a cybersecurity approach that focuses on continuously monitoring endpoint devices—such as laptops, desktops, servers, mobile devices, and IoT gadgets—to detect, investigate, and respond to security threats in real time.

- EDR solutions deploy lightweight agents on endpoints that continuously collect and analyze telemetry data (e.g., process executions, network connections, file activities, user logins) to establish a baseline of normal behavior.

- When anomalous or suspicious activity is detected, the EDR tool generates alerts and provides forensic details that help security teams understand the attack chain and respond quickly.

- Responses can be automated (such as isolating a compromised device or quarantining suspicious files) or performed manually by security analysts.

5.1. CrowdStrike Falcon

- A cloud-native solution known for its rapid threat detection and response capabilities, providing real-time and historical visibility into endpoint activity.
- Using behavioral analytics and machine learning, Falcon identifies both known threats and sophisticated, emerging attacks such as fileless malware and zero-day exploits.
- Falcon integrates with CrowdStrike Threat Intelligence, enriching alerts with context about adversaries, tactics, techniques, and procedures (TTPs).

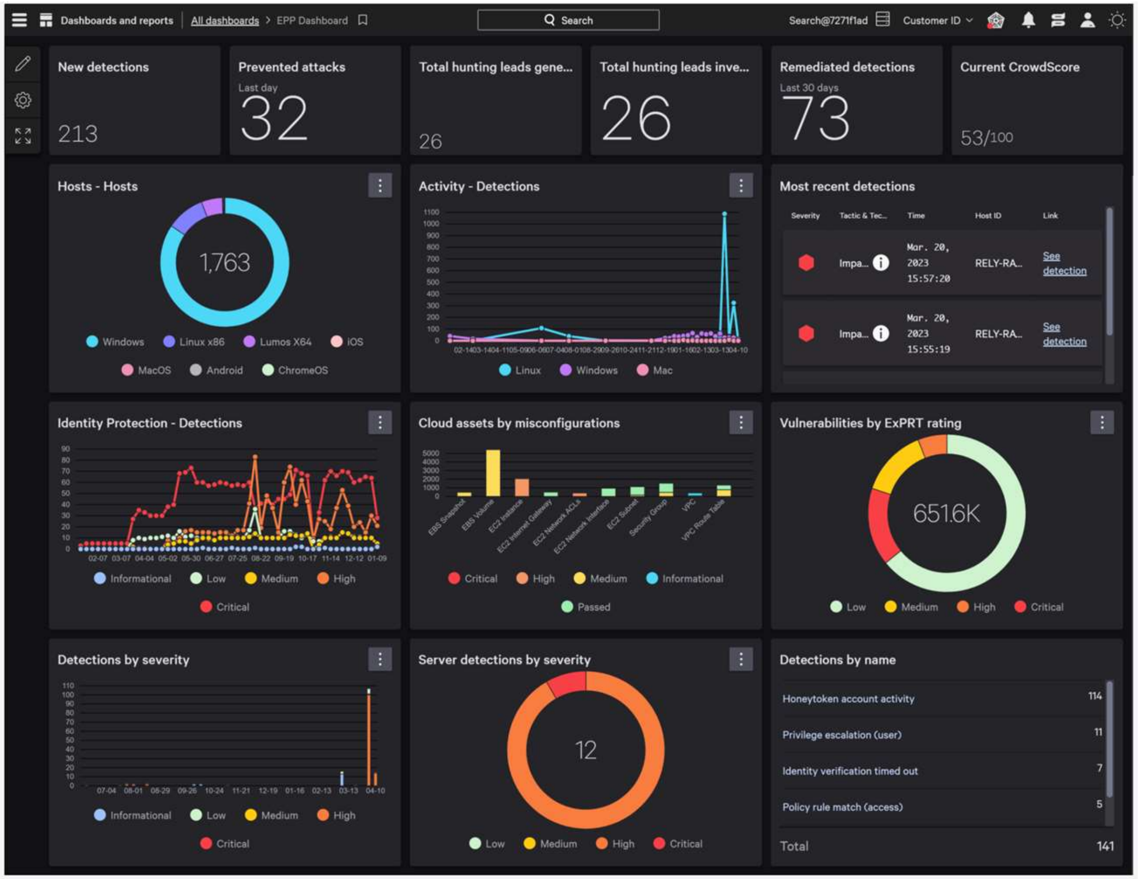

New detections 213

Prevented attacks (Last day) 32

Total hunting leads gene... 26

Total hunting leads inve... 26

Remediated detections (Last 30 days) 73

Current CrowdScore 53/100

Most recent detections

Severity	Tactic & Tec...	Time	Host ID	Link
●	Impa...	Mar. 20, 2023 15:57:20	RELY-RA...	See detection
●	Impa...	Mar. 20, 2023 15:55:19	RELY-RA...	See detection

Detections by name

Honeytoken account activity	114
Privilege escalation (user)	11
Identity verification timed out	7
Policy rule match (access)	5
Total	141

5.2. Microsoft Defender for Endpoint

- Integrates advanced threat protection, behavioral analytics, and automated remediation as part of Microsoft’s comprehensive security suite.

- Know your adversaries with more than 78 trillion daily signals from multiple sources, including the largest clouds, security organizations, 1.5 billion devices, internet graphs, and more than 10 thousand experts in 72 countries.

- Use built-in, security-specific generative AI to rapidly investigate and respond to incidents, prioritize alerts, and learn new skills. Copilot is now embedded in Microsoft Defender XDR for Copilot customers.

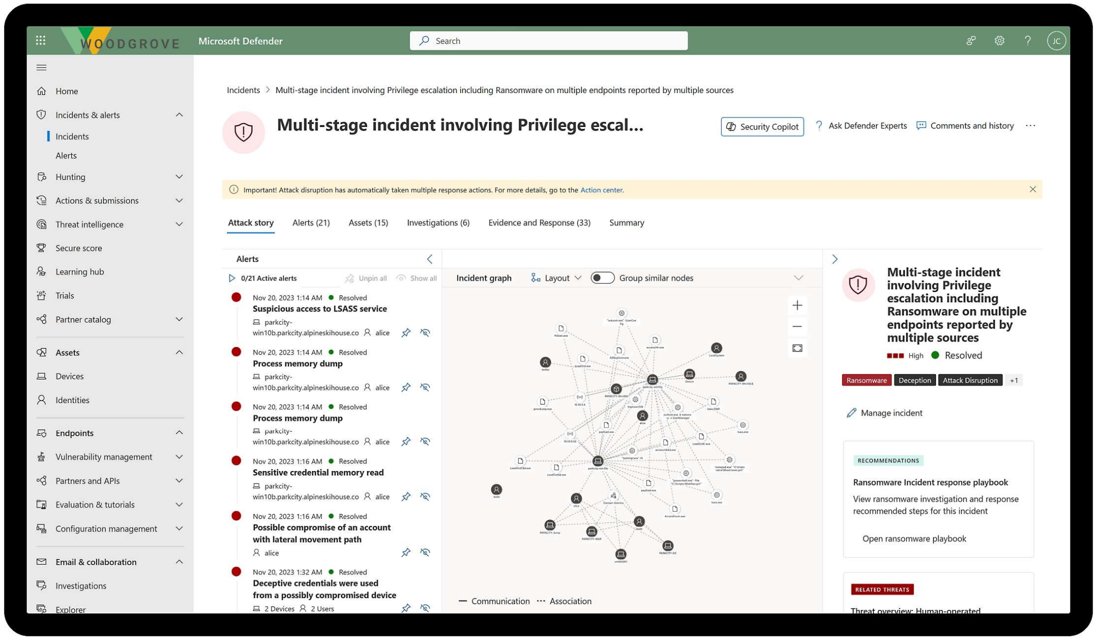

Incidents > Multi-stage incident involving Privilege escalation including Ransomware on multiple endpoints reported by multiple sources

Multi-stage incident involving Privilege escal...

Important! Attack disruption has automatically taken multiple response actions. For more details, go to the Action center.

Nov 20, 2023 1:14 AM ● Resolved
Suspicious access to LSASS service
parkcity-win10b.parkcity.alpineskihouse.co alice

Nov 20, 2023 1:14 AM ● Resolved
Process memory dump
parkcity-win10b.parkcity.alpineskihouse.co alice

Nov 20, 2023 1:14 AM ● Resolved
Process memory dump
parkcity-win10b.parkcity.alpineskihouse.co alice

Nov 20, 2023 1:16 AM ● Resolved
Sensitive credential memory read
parkcity-win10b.parkcity.alpineskihouse.co alice

Nov 20, 2023 1:16 AM ● Resolved
Possible compromise of an account with lateral movement path
alice

Nov 20, 2023 1:32 AM ● Resolved
Deceptive credentials were used from a possibly compromised device
2 Devices 2 Users

Multi-stage incident involving Privilege escalation including Ransomware on multiple endpoints reported by multiple sources
■■■ High ● Resolved

Ransomware Incident response playbook
View ransomware investigation and response recommended steps for this incident

Threat overview: Human-operated

5.3. Carbon Black

- Focuses on detecting advanced threats through continuous monitoring and behavioral analysis on endpoints.

- Carbon Black EDR continuously records and stores comprehensive endpoint activity data, so that security professionals can hunt threats in real time and visualize the complete attack kill chain.

- Precise control, extensive flexibility, and granular customization to adapt your security strategy to evolving threats and confidently prevent attacks.

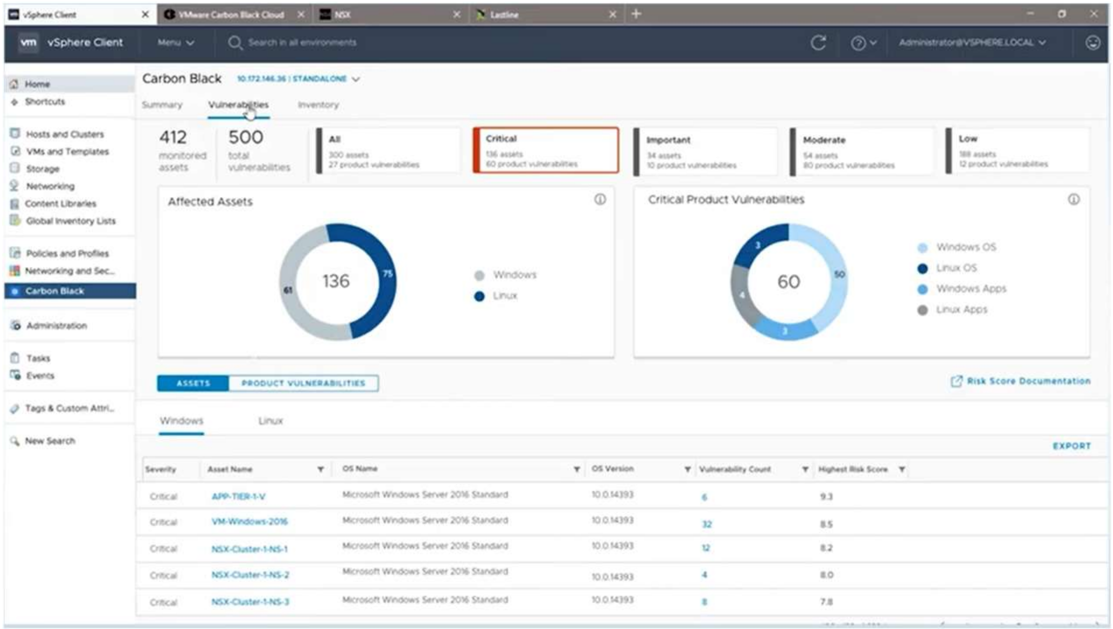

Carbon Black 10.172.146.36 | STANDALONE

412 monitored assets

500 total vulnerabilities

All
300 assets
27 product vulnerabilities

Critical
136 assets
60 product vulnerabilities

Important
34 assets
10 product vulnerabilities

Moderate
54 assets
80 product vulnerabilities

Low
188 assets
12 product vulnerabilities

Severity	Asset Name	OS Name	OS Version	Vulnerability Count	Highest Risk Score
Critical	APP-TIER-1-V	Microsoft Windows Server 2016 Standard	10.0.14393	6	9.3
Critical	VM-Windows-2016	Microsoft Windows Server 2016 Standard	10.0.14393	32	8.5
Critical	NSX-Cluster-1-NS-1	Microsoft Windows Server 2016 Standard	10.0.14393	12	8.2
Critical	NSX-Cluster-1-NS-2	Microsoft Windows Server 2016 Standard	10.0.14393	4	8.0
Critical	NSX-Cluster-1-NS-3	Microsoft Windows Server 2016 Standard	10.0.14393	8	7.8

6. Incident Detection and Response in a SOC

- Incident Detection and Response (IDR) tools are designed to quickly detect, analyze, and remediate security incidents across an organization’s environment. They come in various forms—from SIEM and EDR to SOAR, NDR, and XDR—each focusing on different aspects of the incident lifecycle.

- Incident Response Steps:
  - 1: Early detection.
  - 2: Analysis.
  - 3: Prioritization.
  - 4: Notification.
  - 5: Containment and forensics.
  - 6: Recovery.
  - 7: Incident review.

- Let's imagine the scenario of a hospital getting breached. This means that an outside attacker or insider threat has unauthorized access to private information such as patient social security numbers, medical records, or payment plan data. This private data can then be used by attackers to perform identity theft or sold on secondary markets, making it a lucrative mission target. IDS/IPS Systems help to mitigate these threats.

7. The impact of a SOC in everyday's world

- SOCs also play a crucial role in addressing the persistent cybersecurity skills gap, maximizing the effectiveness of your human talent. They serve as central hubs for incident response coordination, ensuring that when an attack occurs, your team is primed for a swift and effective countermove

- By continuously monitoring a wide array of sources – from network traffic to cloud service logs – SOCs create a dynamic defense system that evolves in real-time.

- By preventing or swiftly mitigating security breaches, SOCs safeguard organizations from the astronomical expenses associated with data loss, system downtime, and reputational damage.

- Operating 24/7/365, SOCs serve as tireless sentinels, keeping a constant watch over an organization's digital landscape.

- With a dedicated team of experts armed with well-honed procedures, SOCs transform potential chaos into a coordinated, efficient response. Unlike ad-hoc approaches, SOC teams follow meticulously crafted plans, ensuring a swift and methodical reaction to threats

8. Artificial Intelligence

- Artificial intelligence (AI) is a set of technologies that enable computers to perform a variety of advanced functions, including the ability to see, understand and translate spoken and written language, analyze data, make recommendations, and more.

- AI requires specialized hardware and software for writing and training machine learning algorithms. No single programming language is used exclusively in AI, but Python, R, Java, C++ and Julia are all popular languages among AI developers.

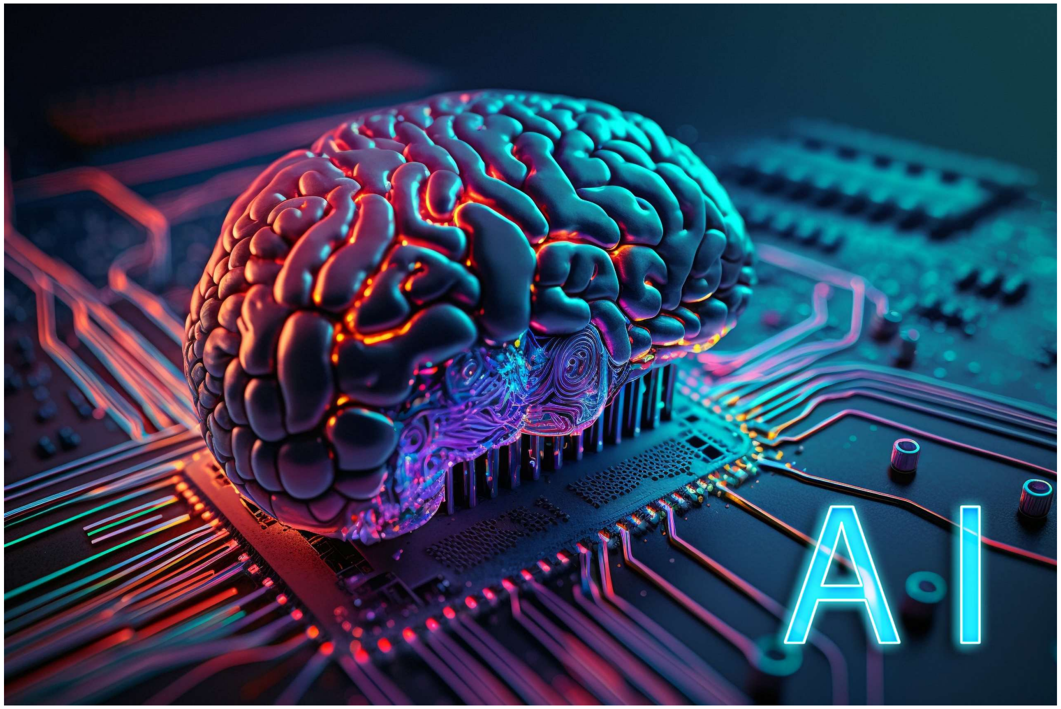

9. Types of Artificial Intelligence

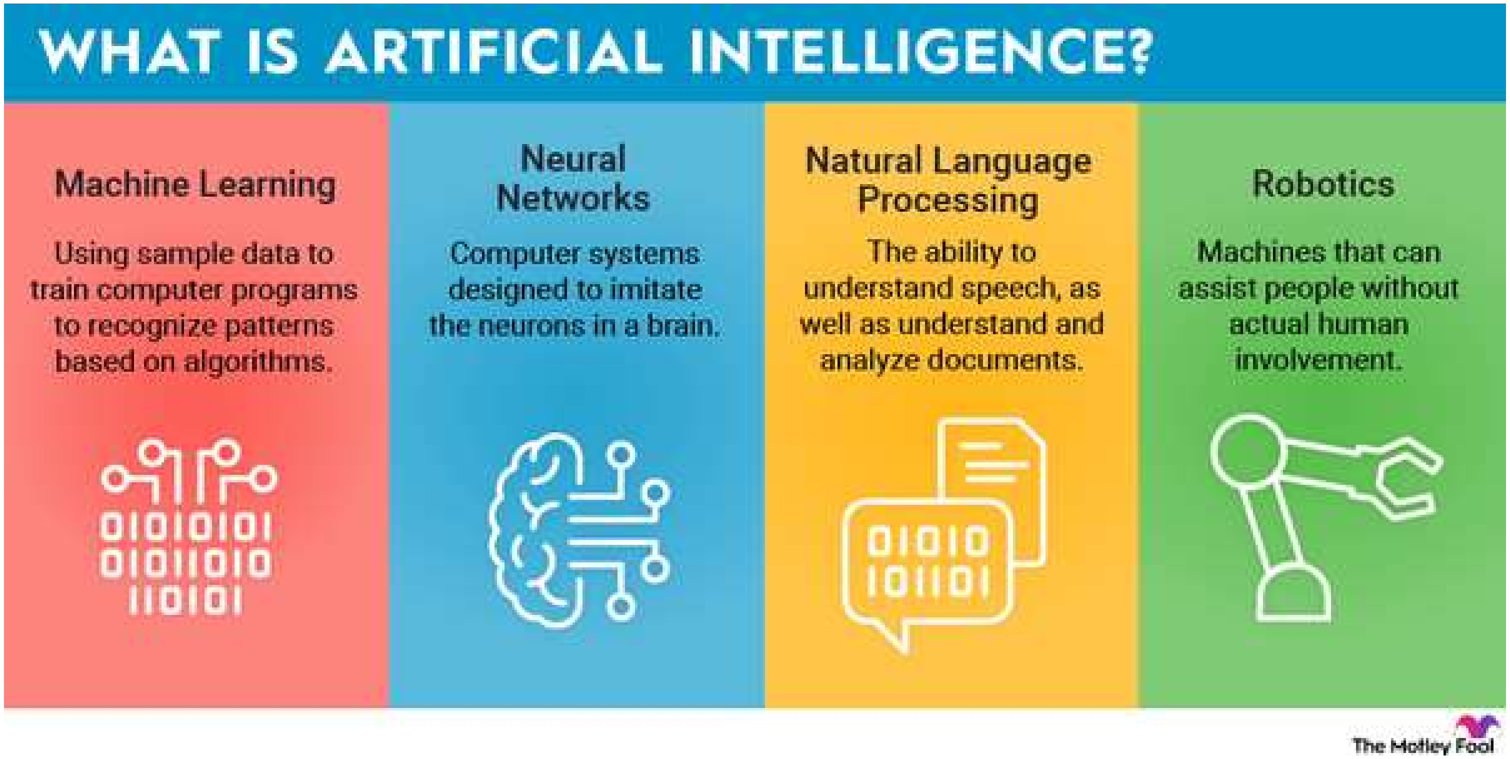

WHAT IS ARTIFICIAL INTELLIGENCE?

Machine Learning

Using sample data to train computer programs to recognize patterns based on algorithms.

Neural Networks

Computer systems designed to imitate the neurons in a brain.

Natural Language Processing

The ability to understand speech, as well as understand and analyze documents.

Robotics

Machines that can assist people without actual human involvement.

The Motley Fool

10. How does an AI model learn?

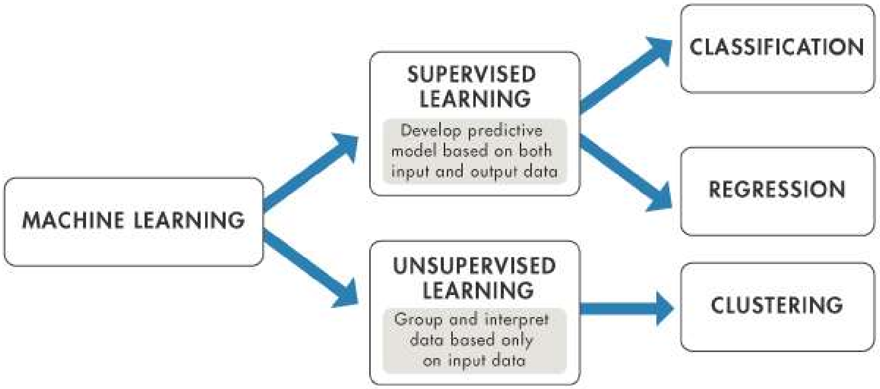
MACHINE LEARNING
SUPERVISED LEARNING
Develop predictive model based on both input and output data
UNSUPERVISED LEARNING
Group and interpret data based only on input data
CLASSIFICATION
REGRESSION
CLUSTERING

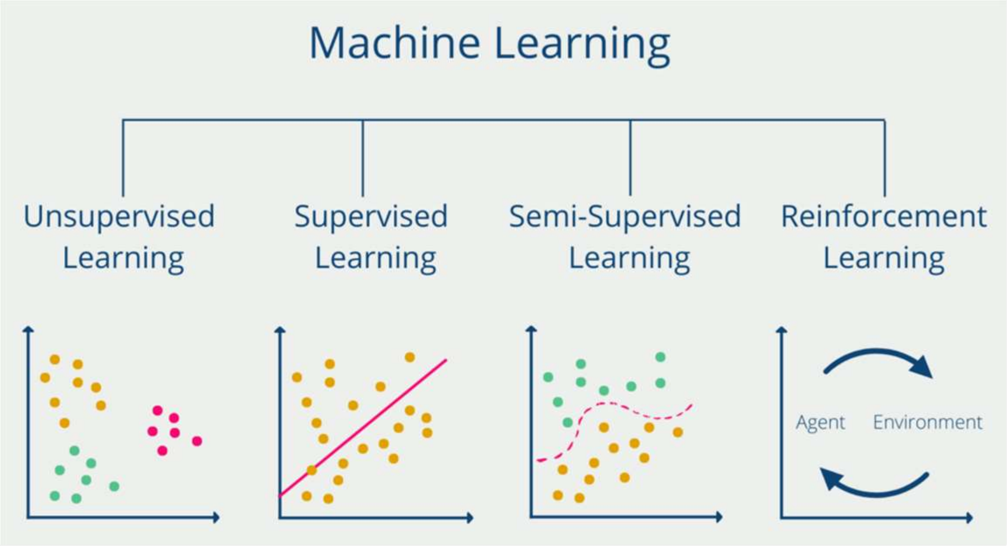

10.1. Supervised learning

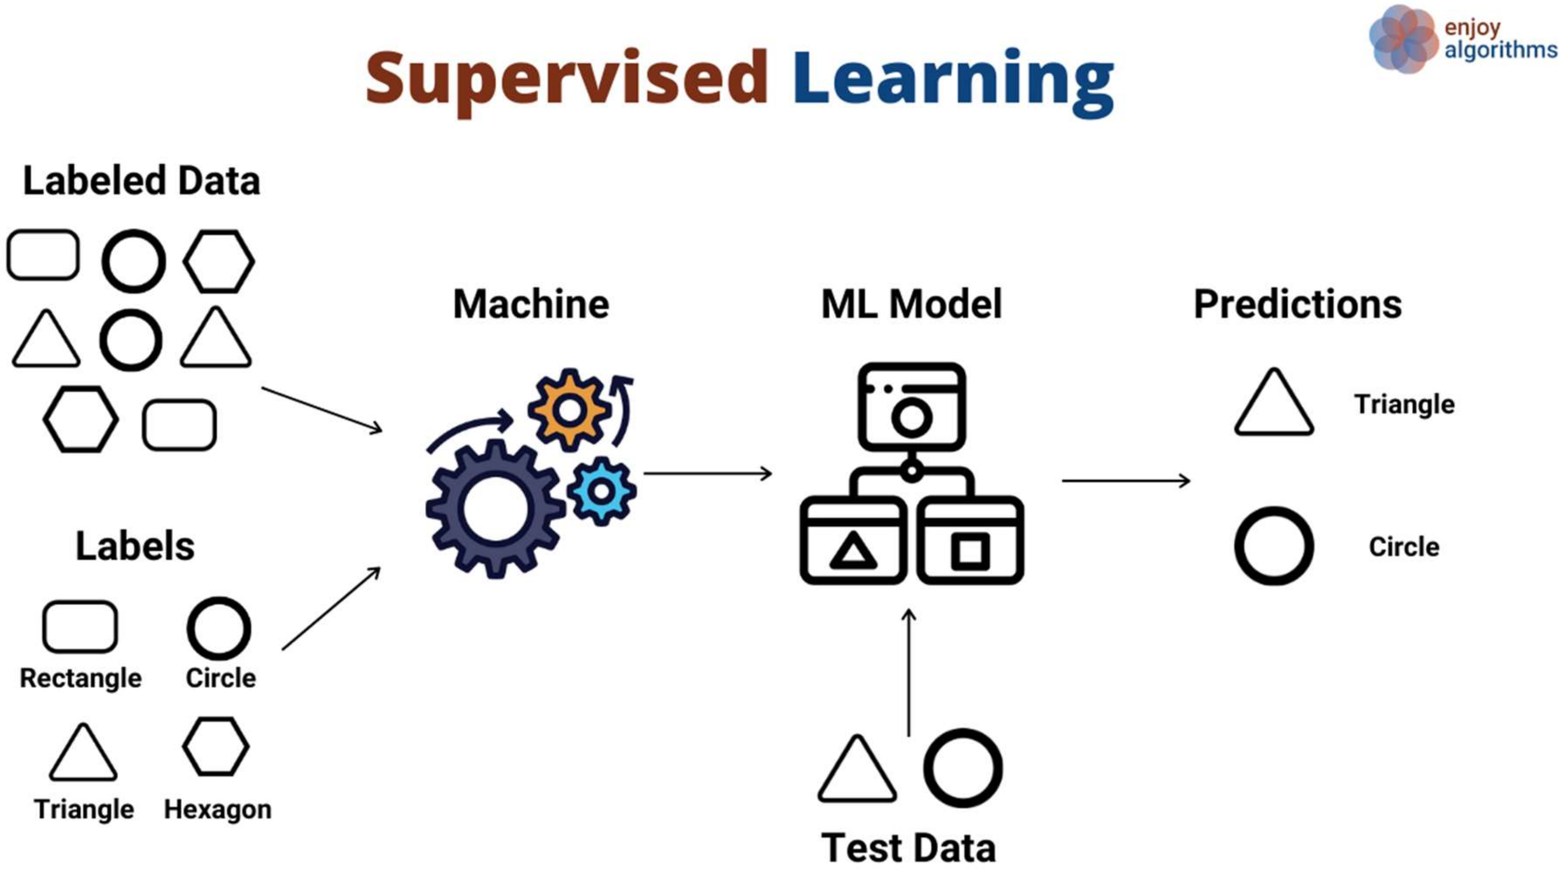

10.2. Unsupervised learning - Iris dataset

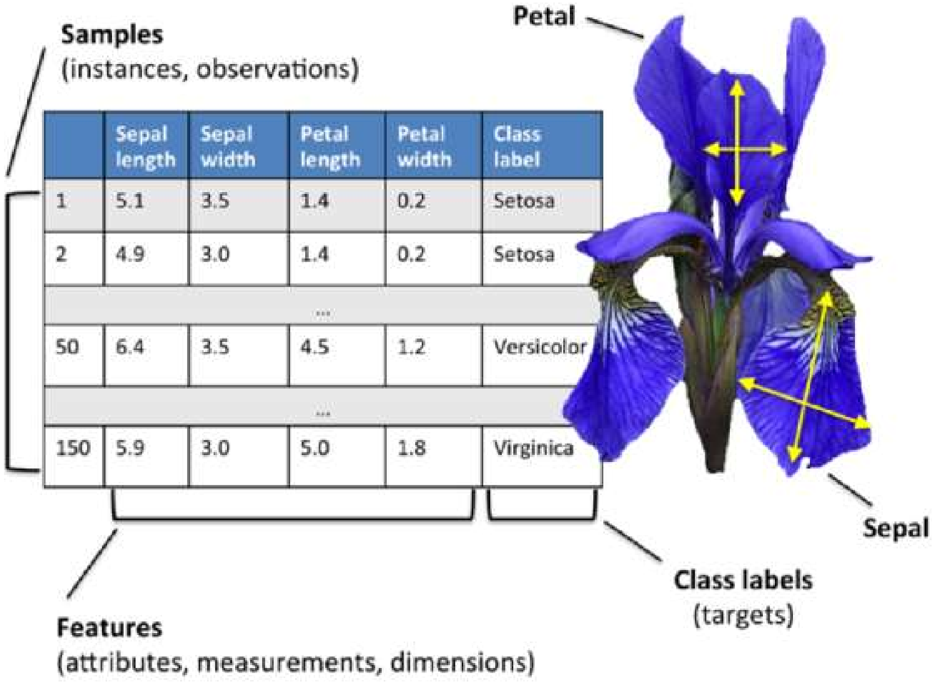
Samples
(instances, observations)
Petal
Sepal
Class labels
(targets)
Features
(attributes, measurements, dimensions)

	Sepal length	Sepal width	Petal length	Petal width	Class label
1	5.1	3.5	1.4	0.2	Setosa
2	4.9	3.0	1.4	0.2	Setosa
			...		
50	6.4	3.5	4.5	1.2	Versicolor
			...		
150	5.9	3.0	5.0	1.8	Virginica

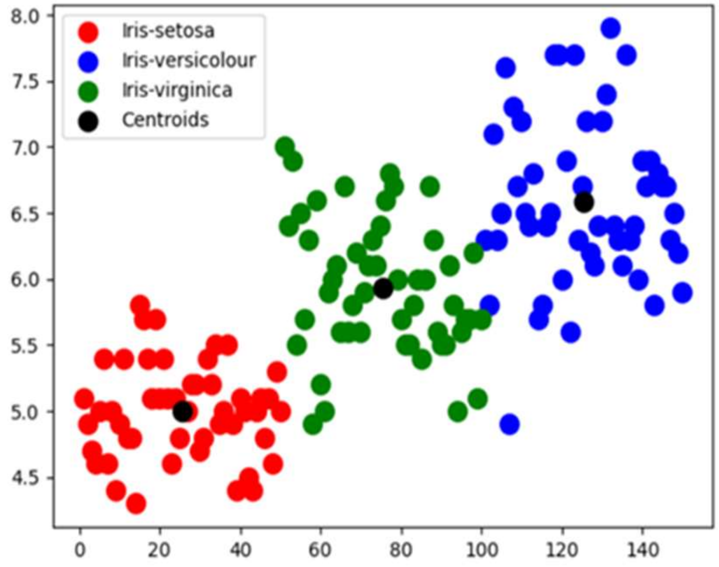

10.3. Reinforcement learning

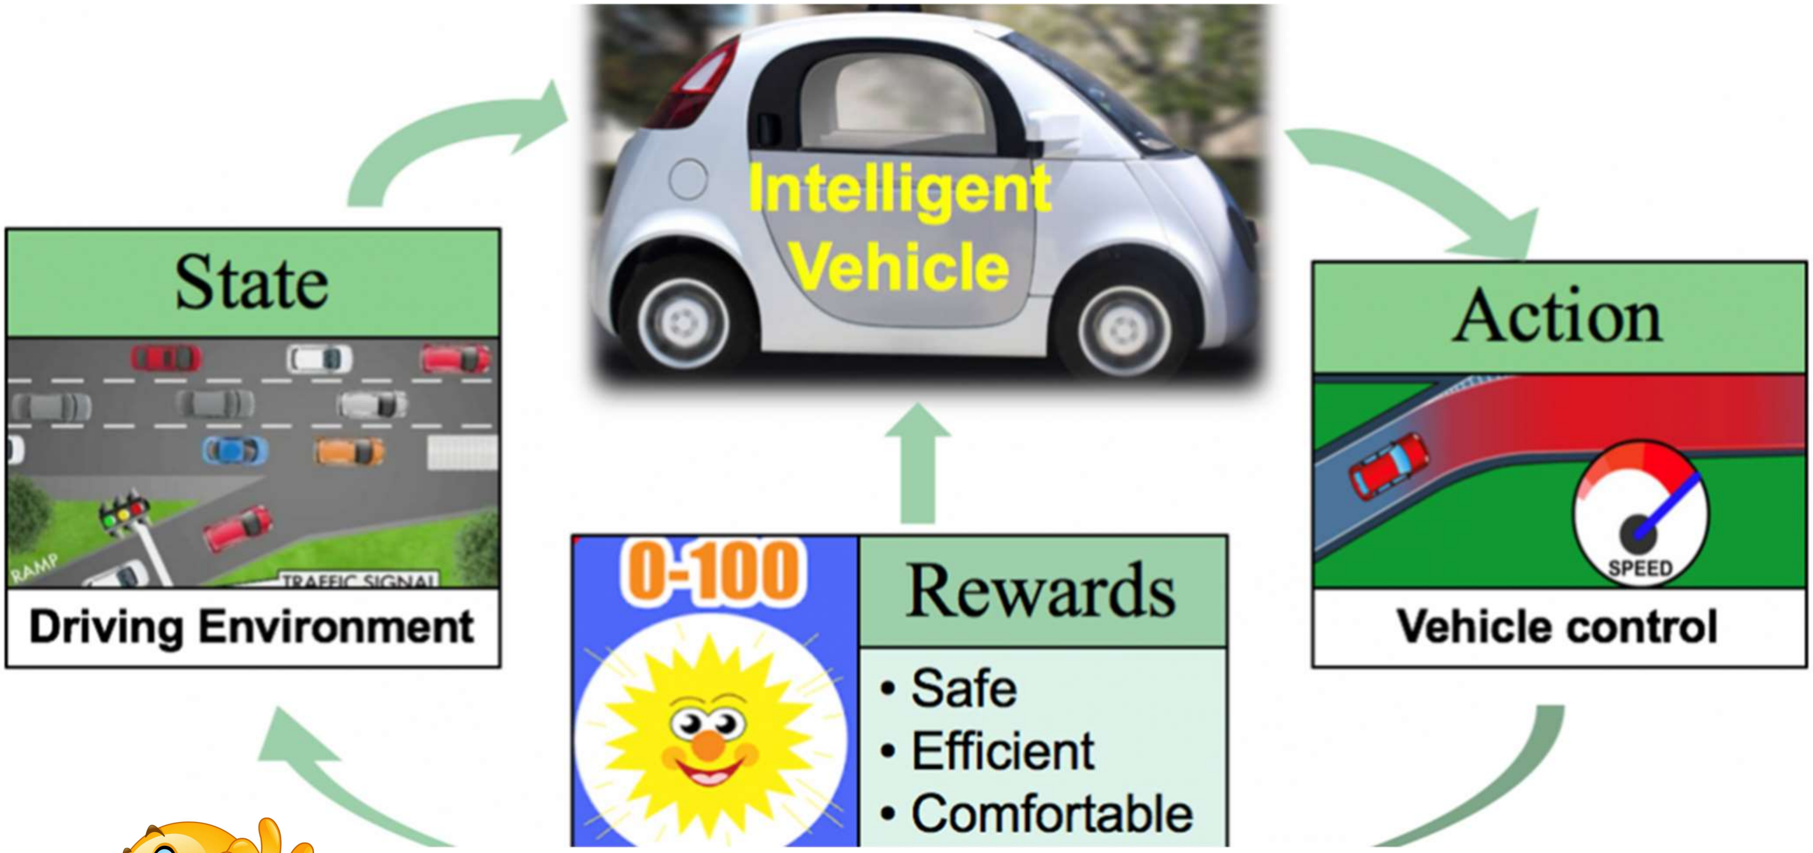

11. Natural Language Processing

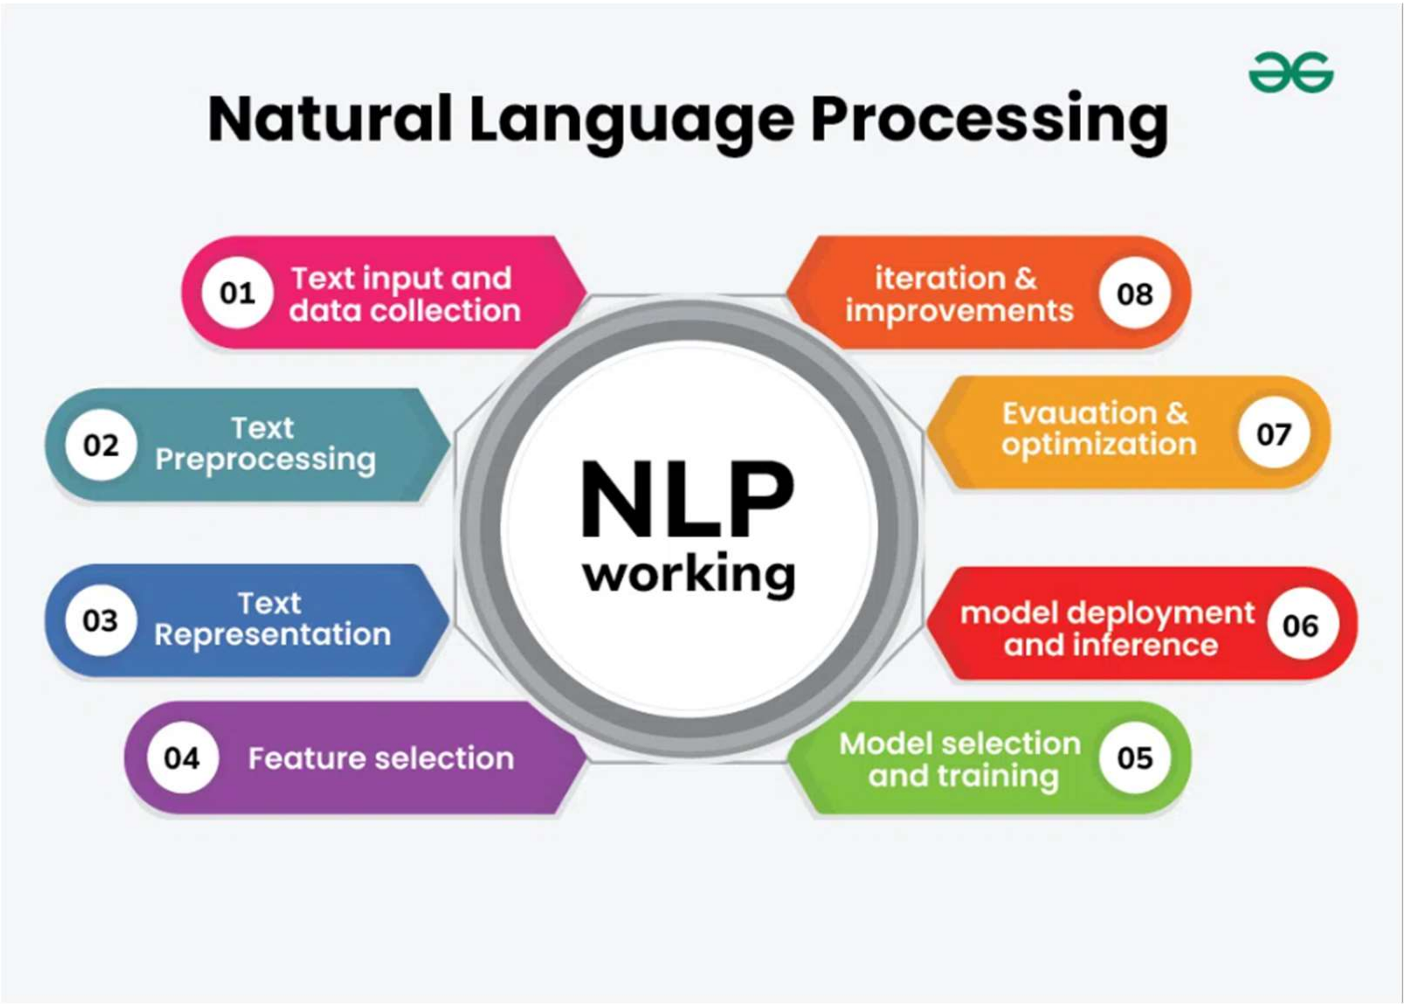
Natural Language Processing
NLP
working
01 Text input and data collection
02 Text Preprocessing
03 Text Representation
04 Feature selection
05 Model selection and training
06 model deployment and inference
07 Evauation & optimization
08 iteration & improvements

11.1. Natural Language Processing (NLP)

1. Text Processing and Preprocessing In NLP
    - Tokenization: Dividing text into smaller units, such as words or sentences.
    - Stemming and Lemmatization: Reducing words to their base or root forms.
    - Stopword Removal: Removing common words (like “and”, “the”, “is”) that may not carry significant meaning.
    - Text Normalization: Standardizing text, including case normalization, removing punctuation, and correcting spelling errors.
2. Syntax and Parsing in NLP
    - Part-of-Speech Tagging: Assigning parts of speech to each word in a sentence (e.g., noun, verb, adjective).
    - Constituency Parsing: Breaking down a sentence into its constituent parts or phrases (e.g., noun phrases, verb phrases).
3. Semantic Analysis:
    - Named Entity Recognition (NER): Identifying and classifying entities in text, such as names of people, organizations, locations, dates, etc.
    - Word Sense Disambiguation (WSD): Determining which meaning of a word is used in a given context.
    - Coherence Resolution: Identifying when different words refer to the same entity in a text (e.g., “he” refers to “John”).
4. Text Classification in NLP
    - Sentiment Analysis: Determining the sentiment or emotional tone expressed in a text (e.g., positive, negative, neutral).
    - Topic Modeling: Identifying topics or themes within a large collection of documents.
    - Spam Detection: Classifying text as spam or not spam.

11.2. Natural Language Processing (NLP)

5. Language Generation
   - Machine Translation: Translating text from one language to another.
   - Text Summarization: Producing a concise summary of a larger text.
   - Text Generation: Automatically generating coherent and contextually relevant text.

6. Speech Processing
   - Speech Recognition: Converting spoken language into text.
   - Text-to-Speech (TTS) Synthesis: Converting written text into spoken language.

7. Question Answering
   - Retrieval-Based QA: Finding and returning the most relevant text passage in response to a query.
   - Generative QA: Generating an answer based on the information available in a text corpus.

8. Sentiment and Emotion Analysis in NLP
   - Emotion Detection: Identifying and categorizing emotions expressed in text.
   - Opinion Mining: Analyzing opinions or reviews to understand public sentiment toward products, services, or topics.

11.3. But how does NLP work? By using embeddings!

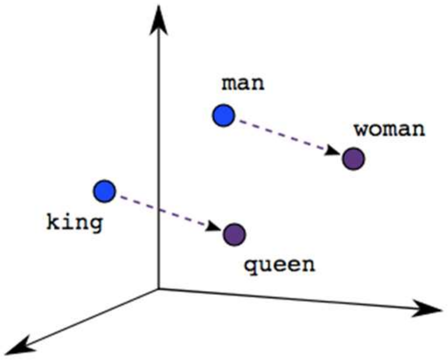

Male-Female

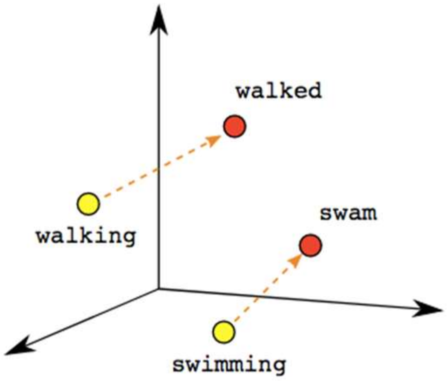

Verb Tense

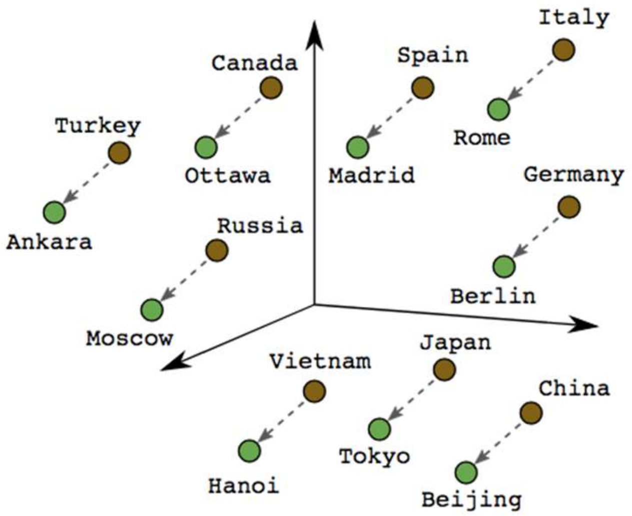

Country-Capital

12. Large Language Models (LLMs)

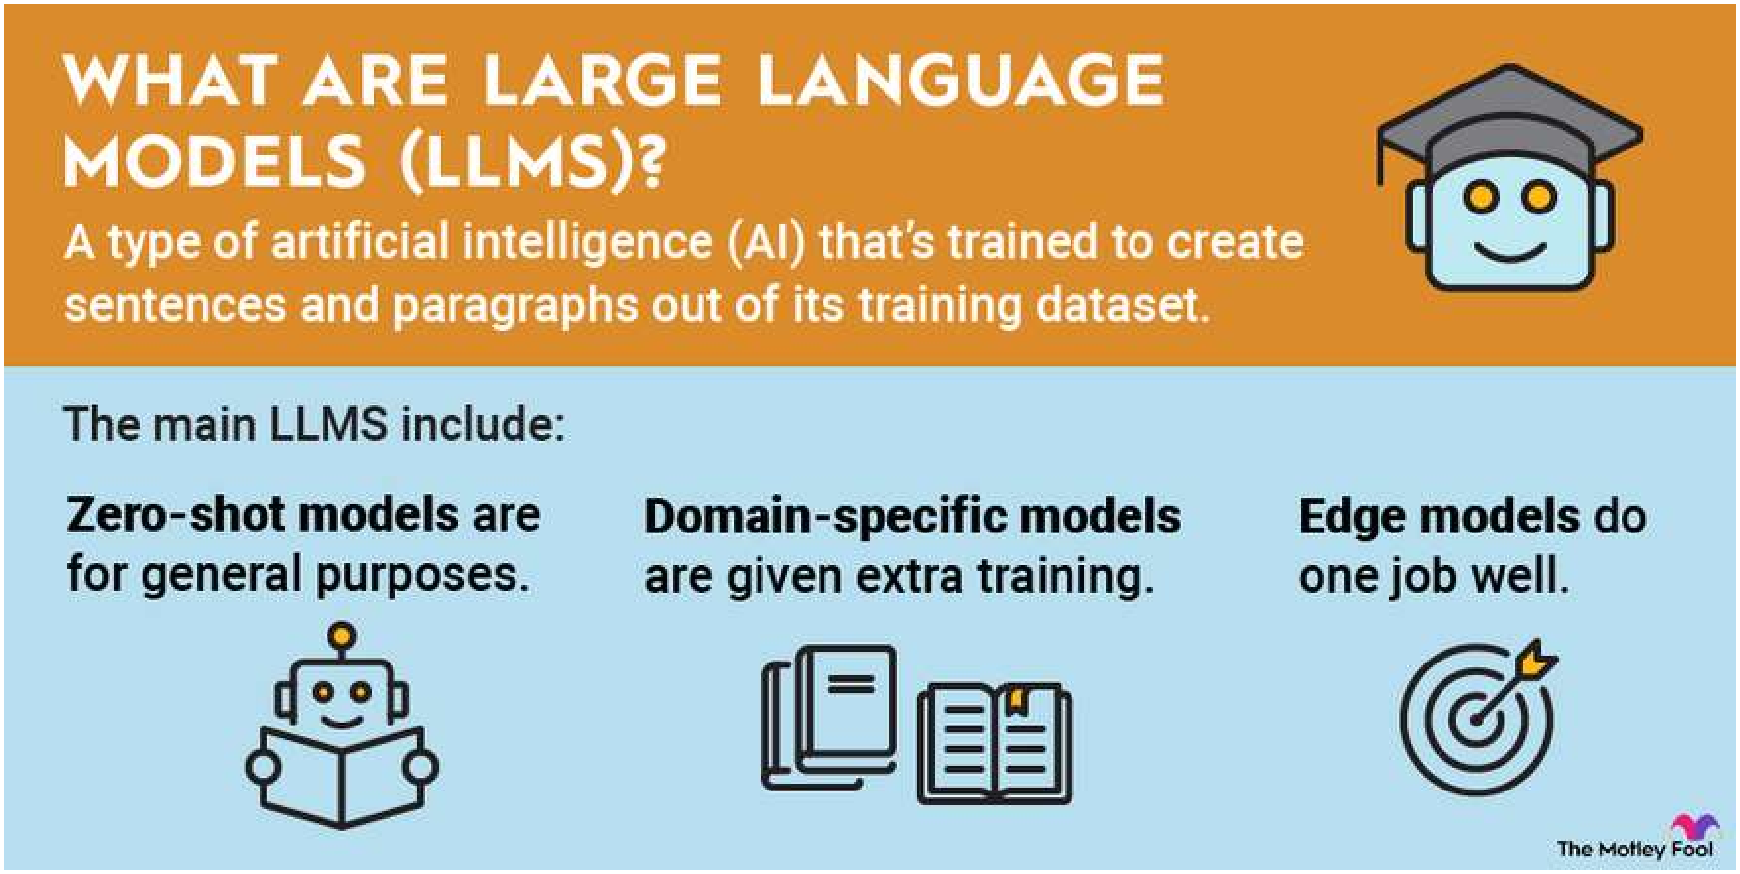

WHAT ARE LARGE LANGUAGE MODELS (LLMS)?

A type of artificial intelligence (AI) that's trained to create sentences and paragraphs out of its training dataset.

The main LLMS include:

Zero-shot models are for general purposes.

Domain-specific models are given extra training.

Edge models do one job well.

The Motley Fool

13. Log analysis using AI & NLP

- Log analysis using AI involves applying machine learning and advanced analytics to sift through vast amounts of log data—from application logs and system events to security and network logs—in order to automatically detect anomalies, identify emerging threats, and generate actionable insights.

- This approach can transform what is traditionally a manual and time-consuming process into an automated, scalable, and more accurate one.

- AI algorithms learn what "normal" behavior looks like by analyzing historical log data. Once a baseline is established, they can flag deviations—such as unexpected spikes in activity or unusual error patterns—which may indicate security incidents or performance issues.

- By correlating logs from multiple sources, AI can uncover patterns that might not be evident when looking at logs in isolation. This helps in piecing together the sequence of events during an incident, thereby improving forensic investigations and incident response.

- NLP techniques can be applied to unstructured log data (e.g., free-text error messages) to extract meaningful information, standardize data formats, and even predict potential issues.

13.1. Example of automatic security issues analysis and countermeasure recommendation using NLP and LLMs

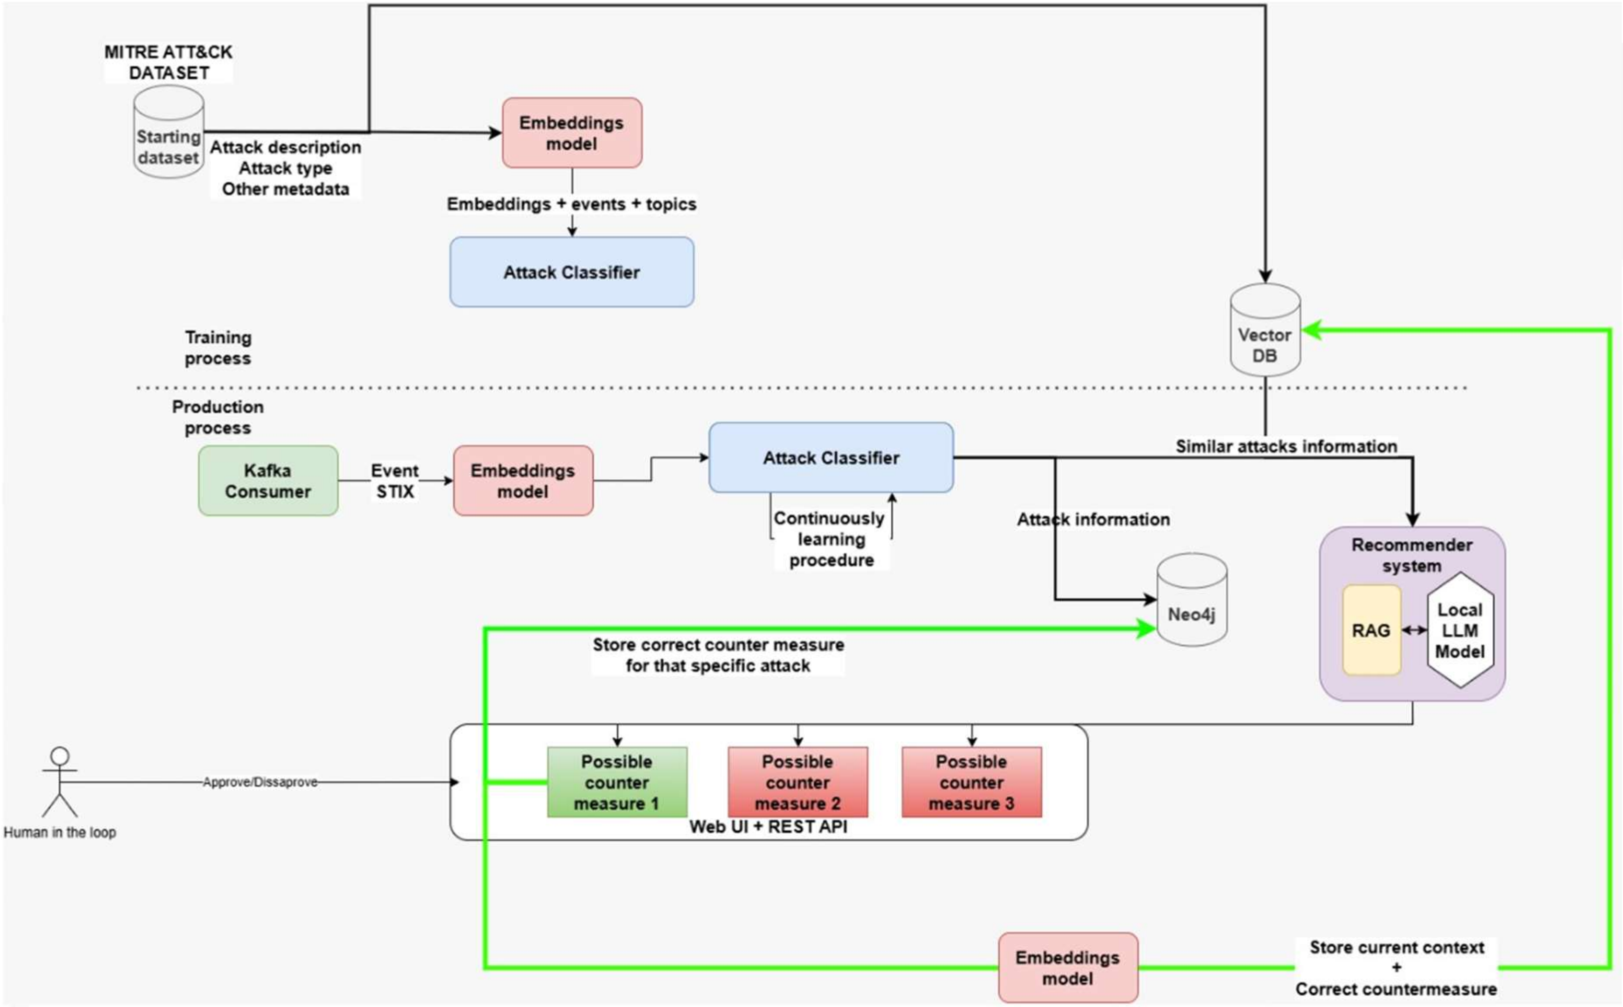

14.1. Fine-tuning a NLP model for classification

- Fine-tuning adapts a pretrained model to a specific task with a smaller specialized dataset. This approach requires far less data and compute compared to training a model from scratch, which makes it a more accessible option for many users.

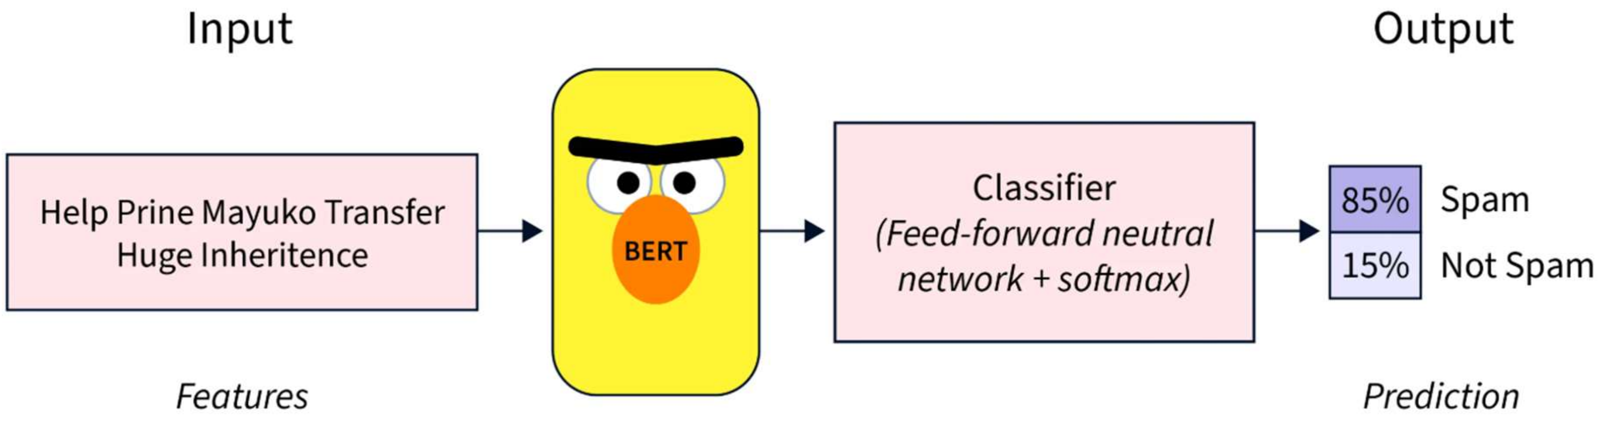
Input
Output
Help Prine Mayuko Transfer Huge Inheritence
BERT
Classifier
(Feed-forward neutral network + softmax)
85%
Spam
15%
Not Spam
Features
Prediction

- Tokenization is where the magic begins. BERT uses **subword tokenization**, meaning it can handle even out-of-vocabulary words by breaking them into smaller chunks.

- Padding and Truncation are key factors in terms of model performance

- Avoid class imbalance in the training dataset, this can be a common pitfall

14.2. Fine-tuning a NLP model for classification

- Imagine you have a model that’s already learned a lot from huge amounts of text data. Now, you want to adapt it to summarize news articles. You don’t need to re-teach the model basic grammar or how to recognize common words — it already knows that! What you *do* need to teach it are specific details about summarization or the particular style of the data you’re working with.

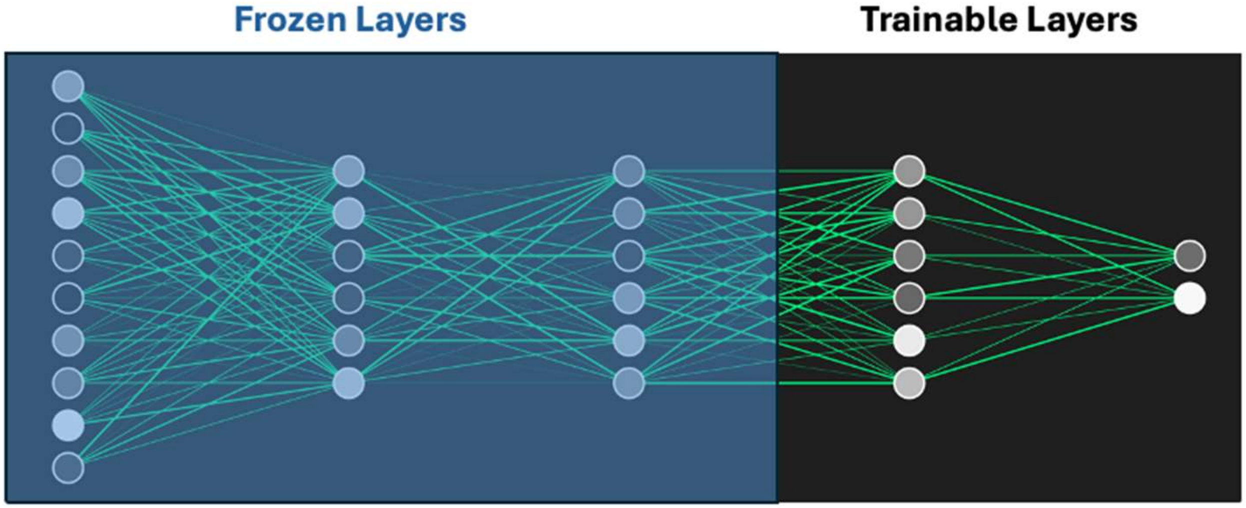

- You might ask: “When should I freeze layers?”
  - For smaller datasets (e.g., <10k samples), freeze the lower layers.
  - For domain-specific tasks, keep all layers trainable to adapt BERT’s embeddings.

14.3. Fine-tuning a LLM model

- Fine-tuning a Large Language Model help the user to benefit the most out of a complex neural network like the one used for such models.

- Improving the performance and answer accuracy of a model for a specific task can greatly increase the usability of LLMs.

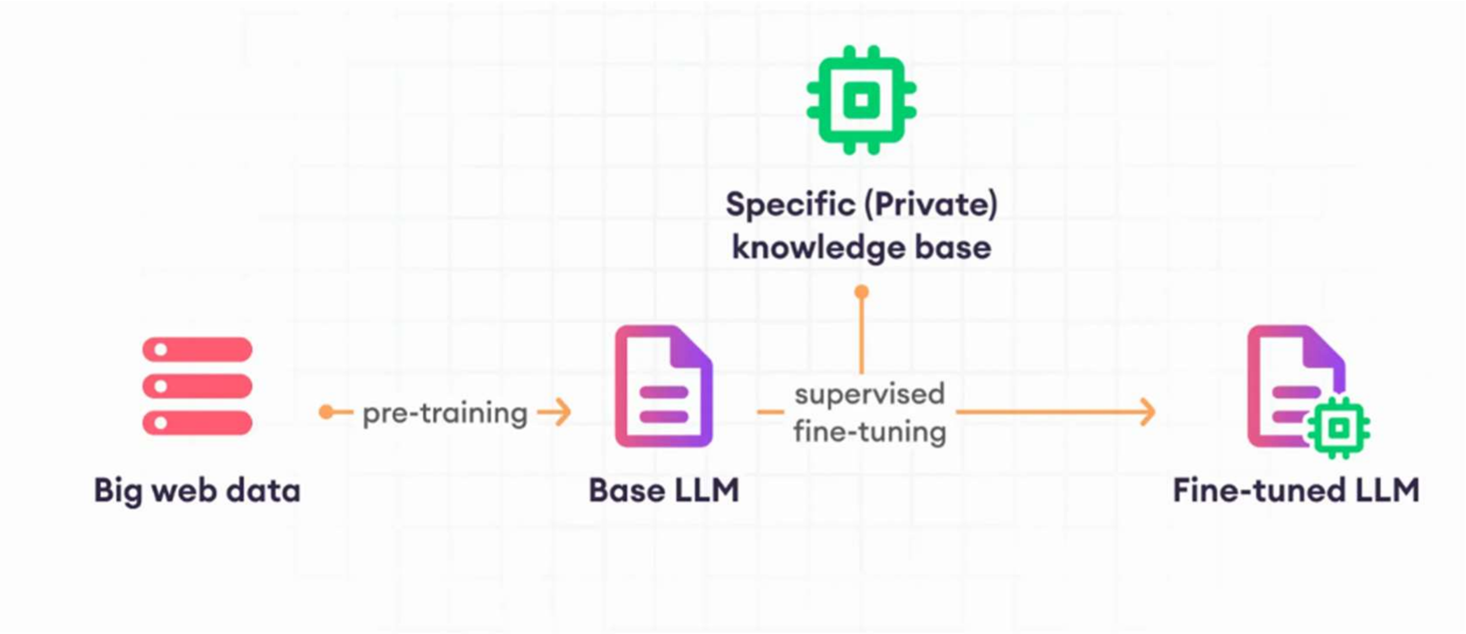

14.3. Fine-tuning a LLM model

- During the fine-tuning phase, when the model is exposed to a newly labeled dataset specific to the target task, it calculates the error or difference between its predictions and the actual labels.

- The model then uses this error to adjust its weights, typically via an optimization algorithm like gradient descent.

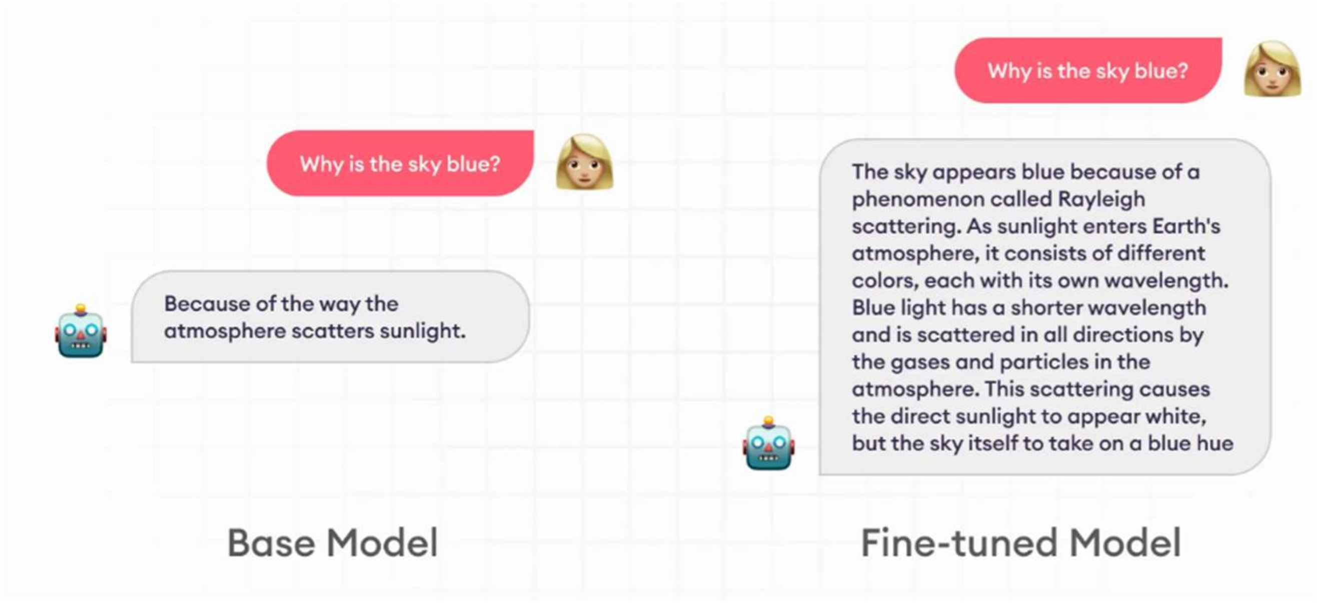

Why is the sky blue?

Because of the way the atmosphere scatters sunlight.

Why is the sky blue?

The sky appears blue because of a phenomenon called Rayleigh scattering. As sunlight enters Earth's atmosphere, it consists of different colors, each with its own wavelength. Blue light has a shorter wavelength and is scattered in all directions by the gases and particles in the atmosphere. This scattering causes the direct sunlight to appear white, but the sky itself to take on a blue hue

14.4. Fine-tuning a LLM model

- The fastest, easiest and cheapest method of “fine-tuning” a LLM model is to adopt a Retrieval Augmented Generation (RAG) architecture

- Using a Knowledge Base we can present new information from where the model can infer new knowledge while generating a response

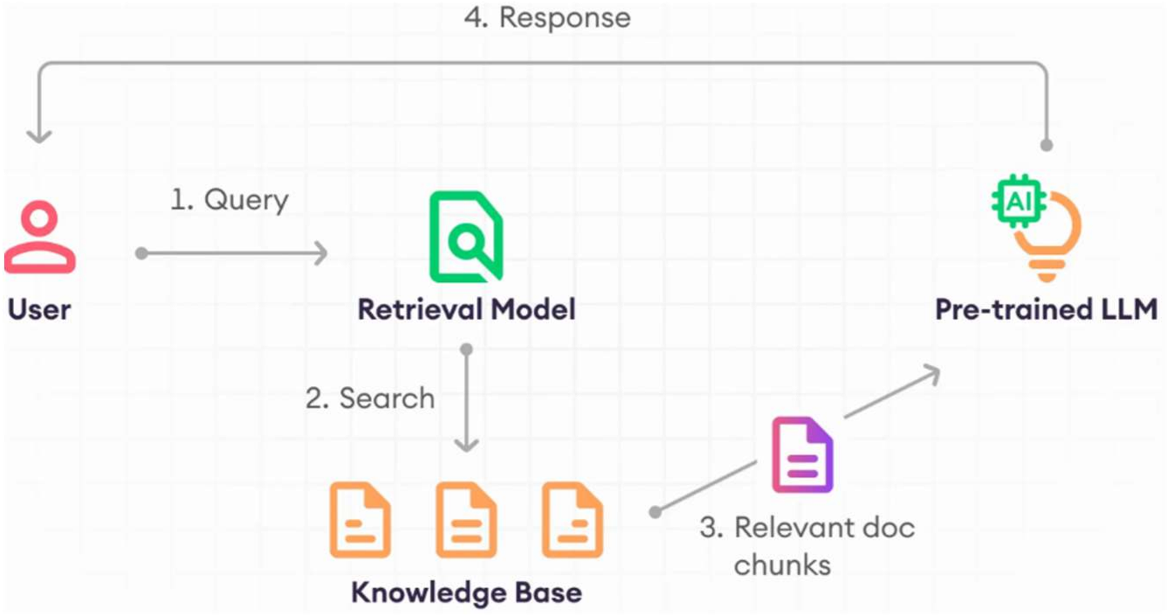

14.5. Fine-tuning a LLM model

- There are multiple types of fine-tuning a LLM, but firstly we have to focus on creating the dataset

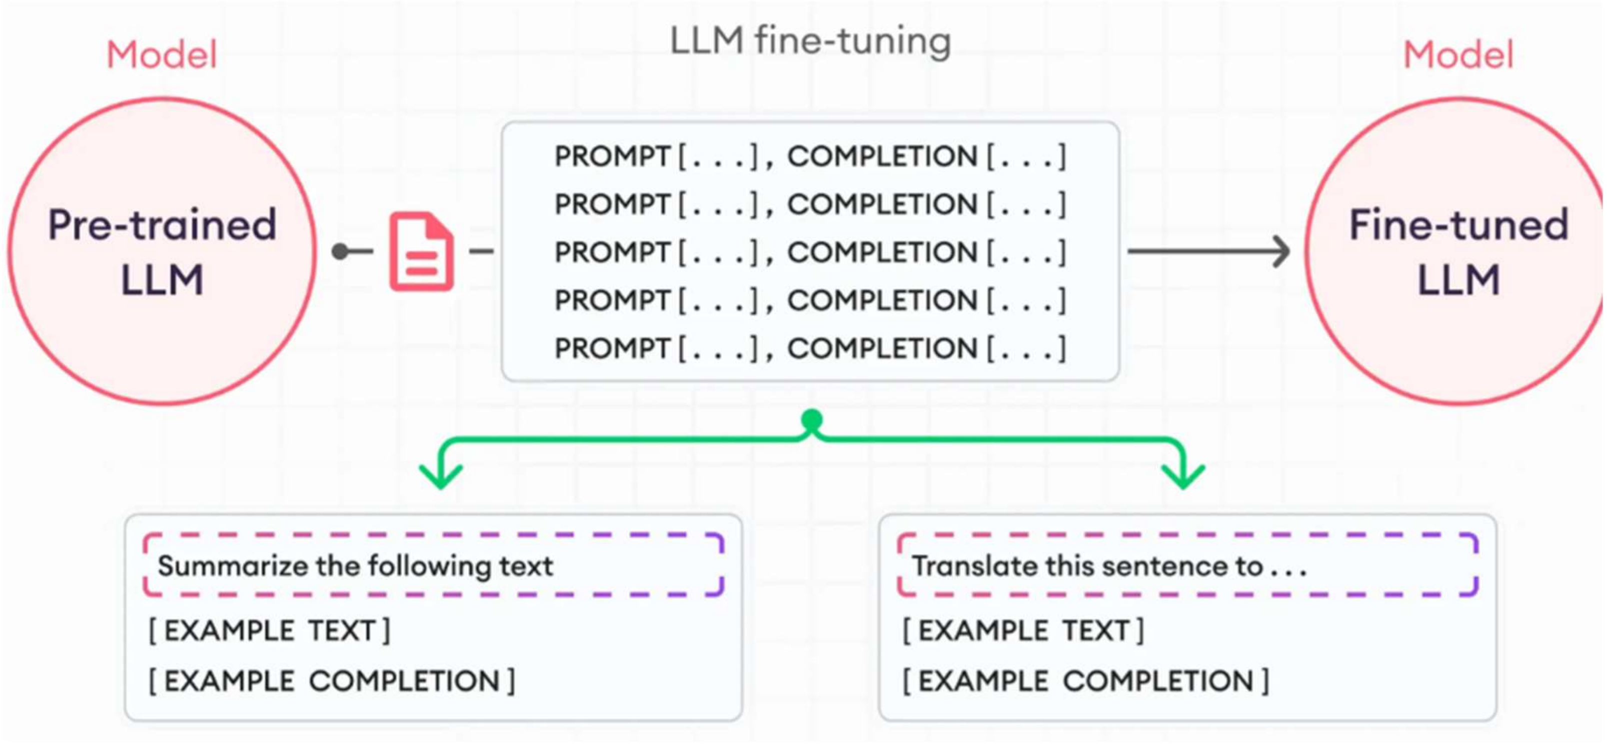

LLM fine-tuning

Model

Pre-trained LLM

PROMPT [. . .], COMPLETION [. . .]
PROMPT [. . .], COMPLETION [. . .]
PROMPT [. . .], COMPLETION [. . .]
PROMPT [. . .], COMPLETION [. . .]
PROMPT [. . .], COMPLETION [. . .]

Model

Fine-tuned LLM

Summarize the following text
[EXAMPLE TEXT]
[EXAMPLE COMPLETION]

Translate this sentence to . . .
[EXAMPLE TEXT]
[EXAMPLE COMPLETION]

14.6. Fine-tuning a LLM model

- Using various methods to train a model can impact the GPU memory print and the final performance (Mistral 7B example)

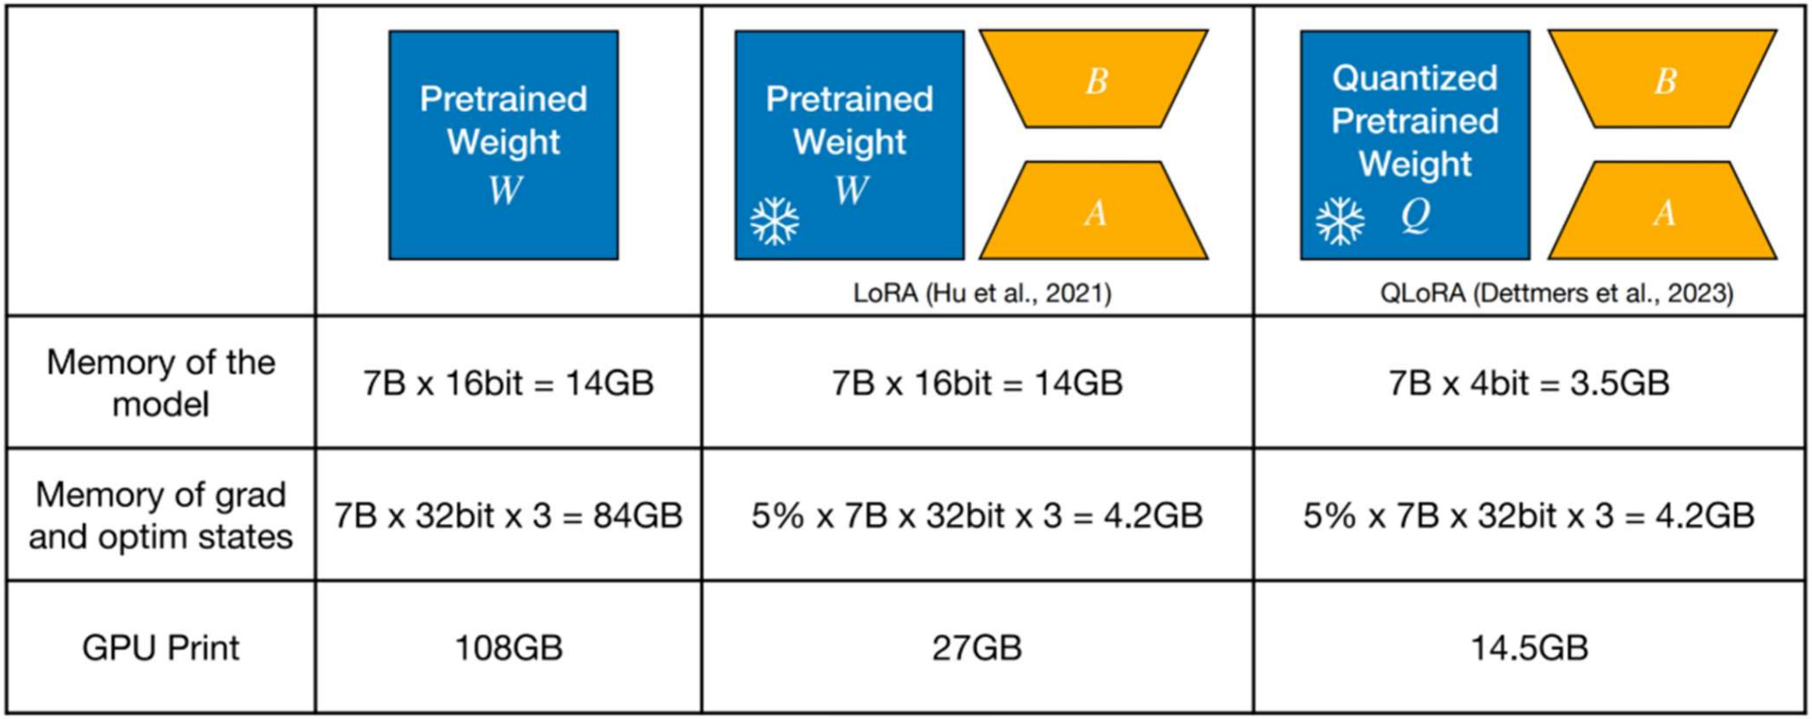

	Pretrained Weight W	Pretrained Weight W B A LoRA (Hu et al., 2021)	Quantized Pretrained Weight Q B A QLoRA (Dettmers et al., 2023)
Memory of the model	7B x 16bit = 14GB	7B x 16bit = 14GB	7B x 4bit = 3.5GB
Memory of grad and optim states	7B x 32bit x 3 = 84GB	5% x 7B x 32bit x 3 = 4.2GB	5% x 7B x 32bit x 3 = 4.2GB
GPU Print	108GB	27GB	14.5GB

14.7. Fine-tuning a LLM model

- Ways to fine-tune a model:

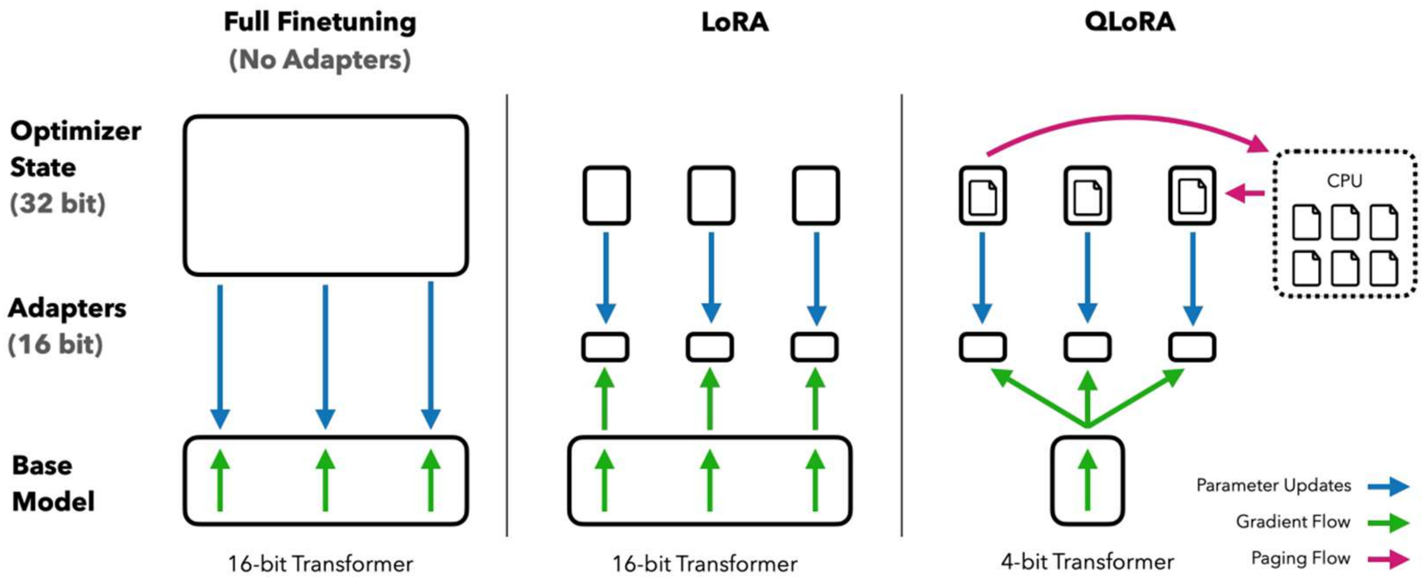

14.8. Fine-tuning a LLM model

– Possible results after fine-tuning a model (e.g. ViGGO dataset - data-to-text generation dataset in the video game domain)

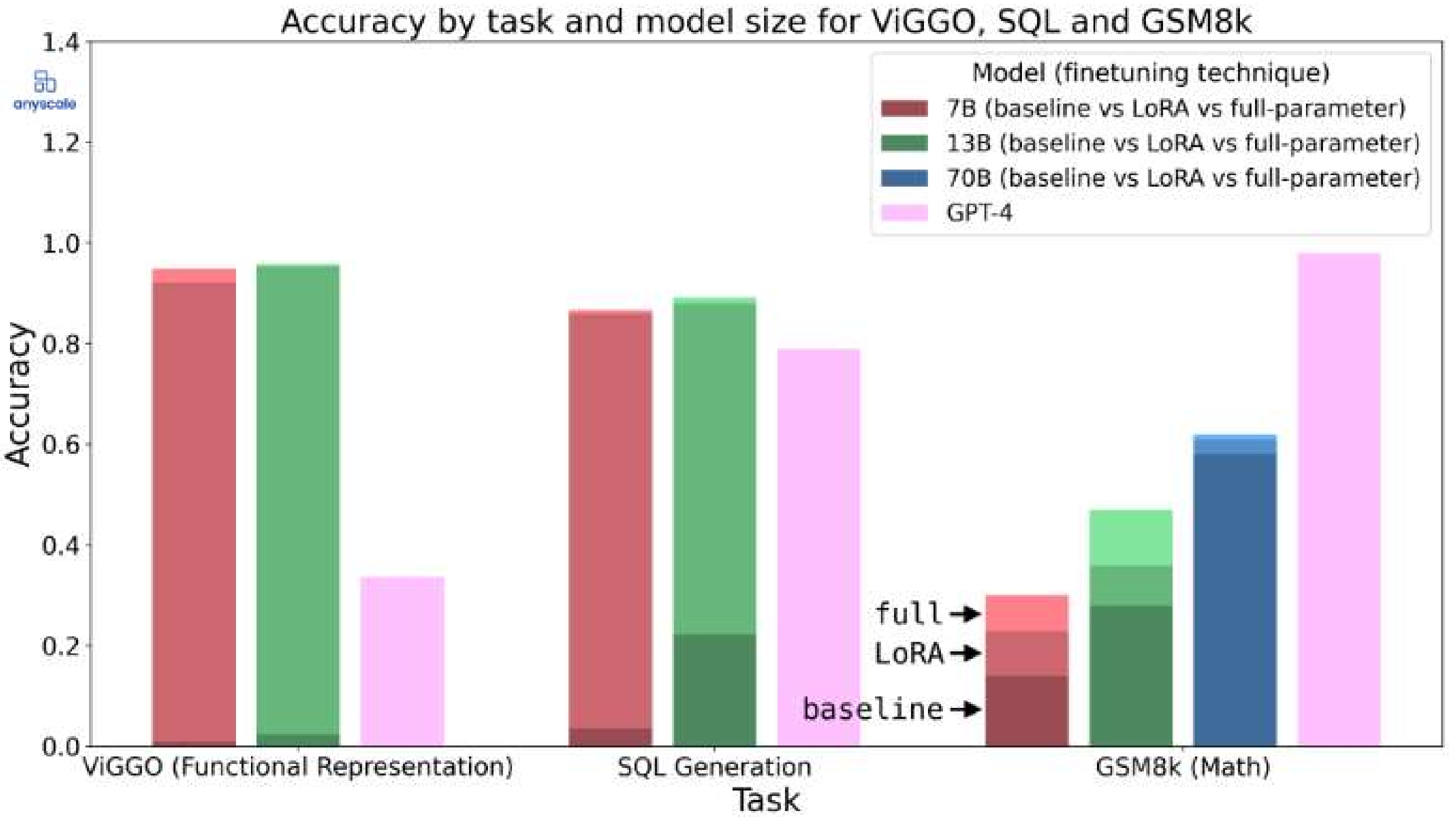

Lab time

https://tinyurl.com/unitbvai

Key Take-aways

1. Utility of SOC
2. Small intro in AI
3. AI & LLMs for SOC and log analysis

Q&A