Democratizing AI Training and Access Statewide

A NAIRR Pilot Expansion Project

Take a Moment, Please Fill Out Our Pre-Workshop Survey:

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Democratizing AI Training and Access Statewide

A NAIRR Pilot Expansion Project

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Disclaimer

This material is based upon work supported in part by the National Science Foundation Award *#2528161*. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Who We Are: Center for Applied AI (CAAI)

Accelerating Discovery with Applied AI & Cross-Disciplinary Collaboration



NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Who We Are: Center for Applied AI (CAAI)

Accelerating Discovery with Applied AI & Cross-Disciplinary Collaboration

What We Do:

- Build secure, interoperable AI toolsets that can be adapted across domains
- Identify and overcome bottlenecks in data, modeling, and analytical workflows
- Work on cross-disciplinary teams to combine AI expertise with domain expertise

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

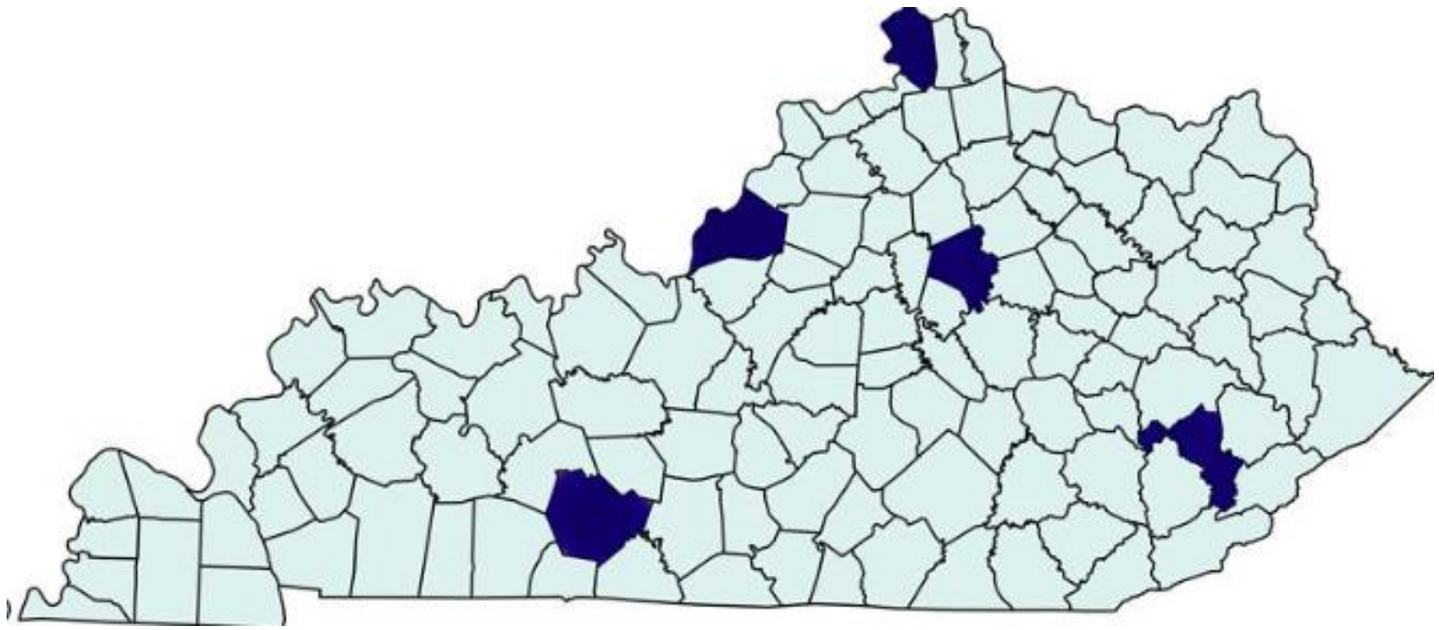
Why This Matters for Kentucky

Empowering researchers and innovators across the Commonwealth

Reducing barriers to AI (knowledge, access, collaboration)

Leveraging statewide partnerships for a connected AI ecosystem

Positioning Kentucky to shape the future of AI



NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Workshop Goals

- Understand core AI concepts and how AI is used in real-world settings
- Explore how AI is shaping research and the workforce in Kentucky
- Gain confidence in foundational ideas like machine learning, large language models, and multimodal AI
- Experience applications of AI through the lens of research workflows
- Learn how to access local and national AI computing resources through the NAIRR Pilot
- Spark new ideas and collaborations across
- Connect Kentucky researchers and innovators to grow a shared AI ecosystem

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

NAIRR Pilot National Artificial Intelligence Research Resource Pilot

Advancing US Innovation in Artificial Intelligence

The NAIRR Pilot aims to connect U.S. researchers and educators to computational, data, and training resources needed to advance AI research and research that employs AI.

nairrpilot.org

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky



CAROLINE LEACH

DATA MANAGEMENT SPECIALIST



AARON MULLEN

SOFTWARE ENGINEER



VAIDEN LOGAN

SOFTWARE ENGINEER



MITCHELL KLUSTY

SOFTWARE ENGINEER

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky



TRAVIS MUNCIE

ASSISTANT VICE PRESIDENT, DATA AND ADVANCED ANALYTICS
COUNCIL FOR POST SECONDARY EDUCATION



DANNA BARNETT

DIRECTOR OF ENTREPRENEURSHIP
SHAPING OUR APPALACHIAN REGION, INC. (SOAR)



RUSTY JUSTICE

CO-FOUNDER, MANAGING MEMBER
BIT SOURCE

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

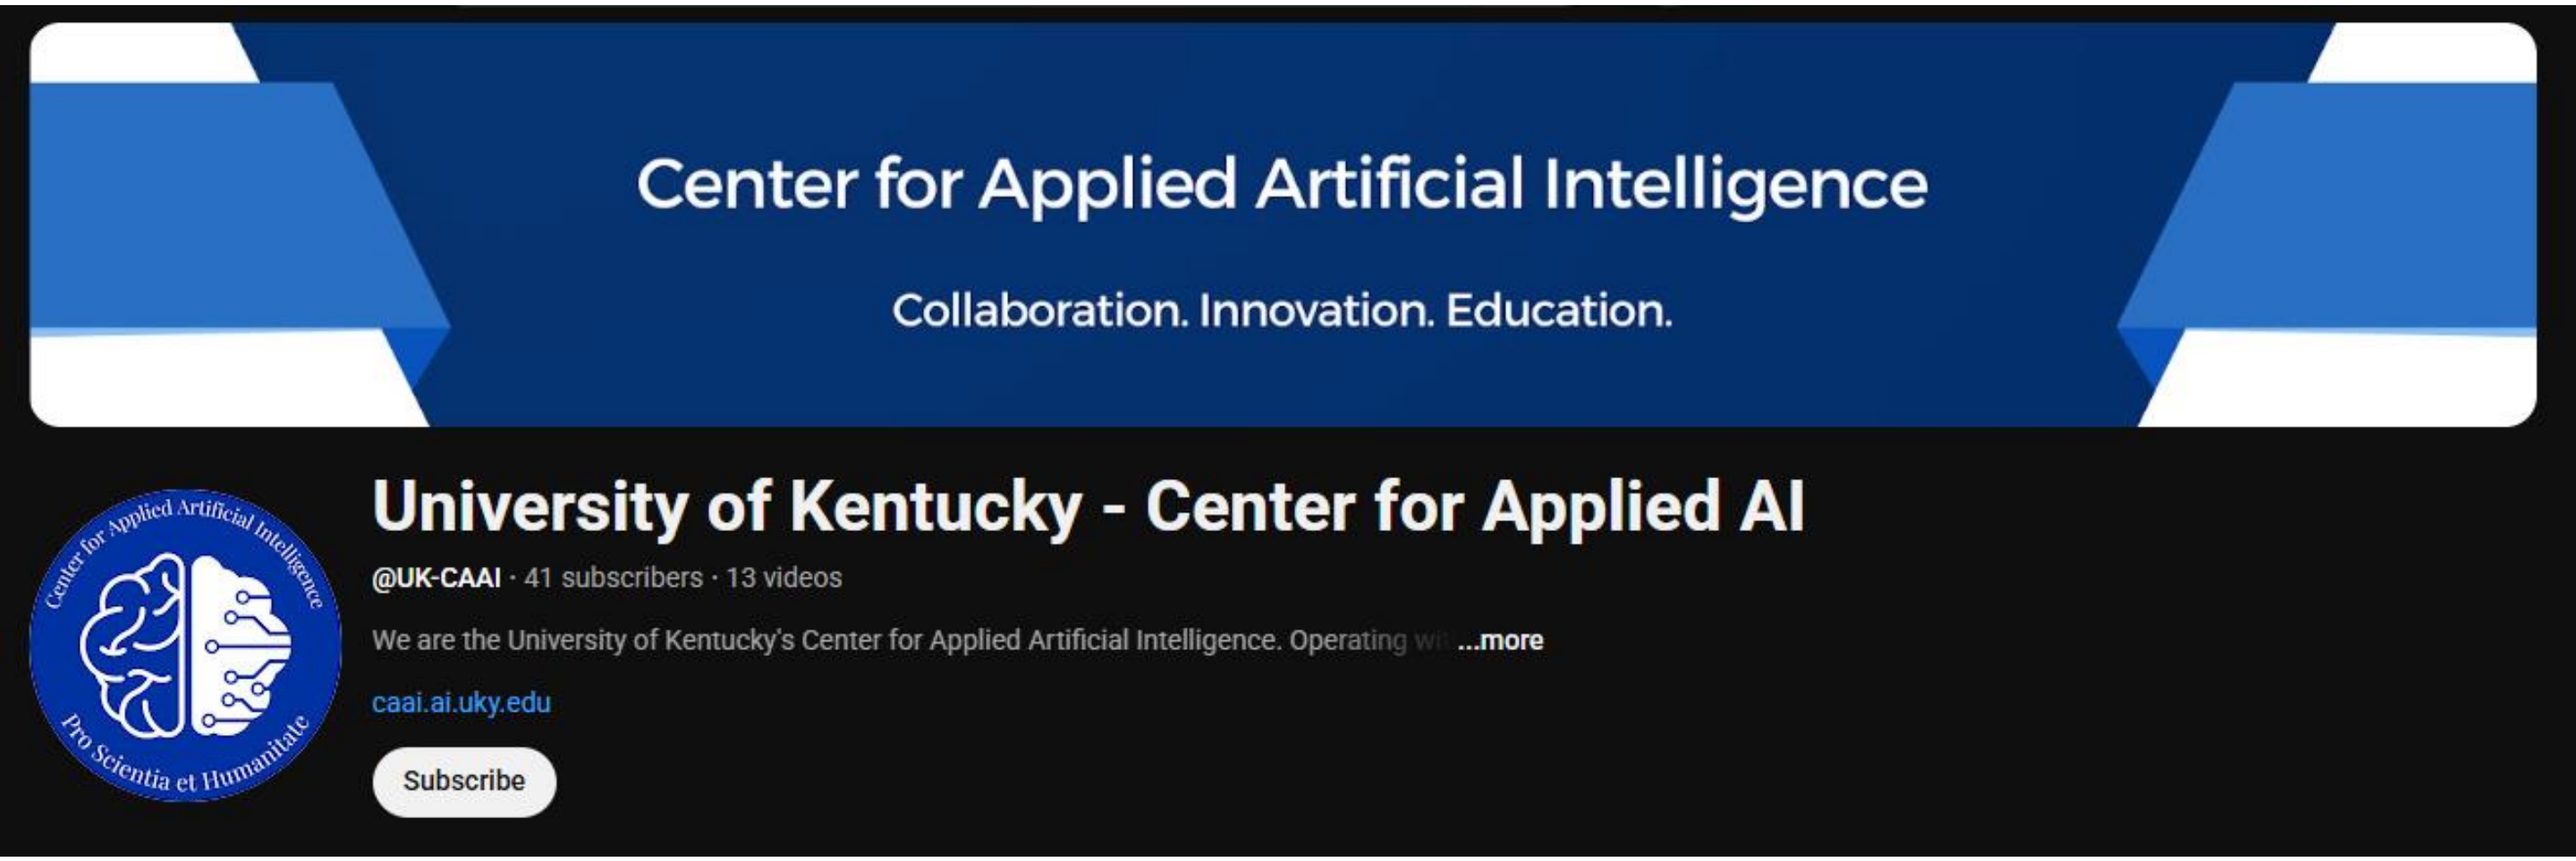
Ongoing resources and support for Kentucky researchers and innovators

To support continued learning and collaboration, we're offering:

- Videos
- Online Forum
- NAIRR Virtual Assistant
- Technical Demos
- Office Hours

Videos and Workshop Recording

https://www.youtube.com/@UK-CAAI



NAIRR Pilot

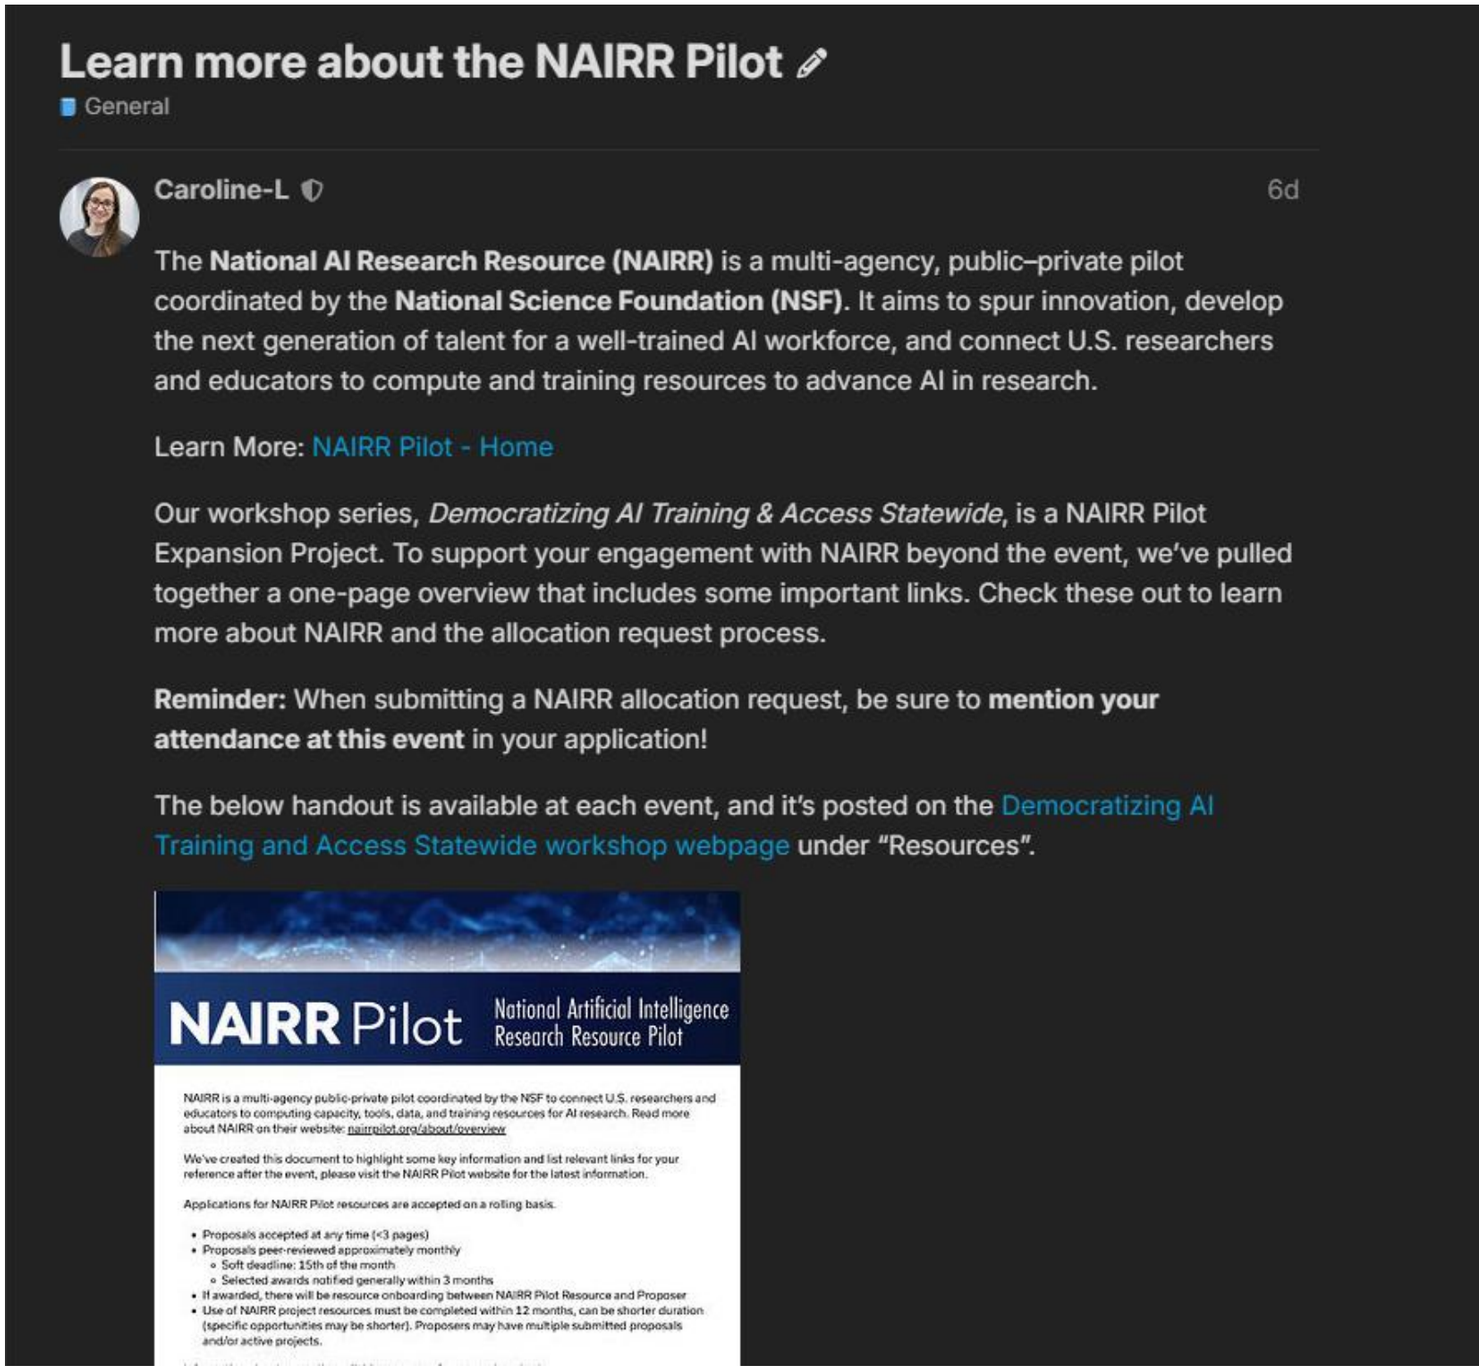
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Online Forum

Scan to create your account on the forums website:

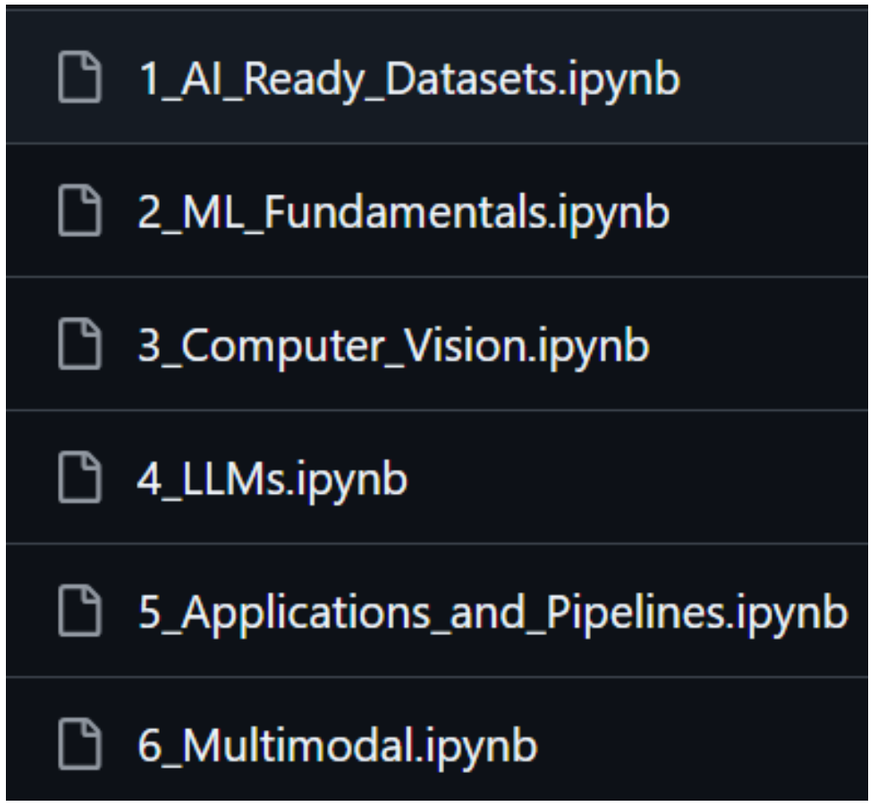


Hazard Community and Technical College

NAIRR Pilot | KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM | CPE Higher Education Matters | University of Kentucky

Demo Repo

GitHub



github.com/Kentucky-Open-Science/CAAI-NAIRR-Pilot-Jupyter-Notebooks

NAIRR PILOT EXPANSION PROJECT:
DEMOCRATIZING AI TRAINING AND
ACCESS STATEWIDE

OFFICE HOURS

SECOND MONDAY EACH MONTH

11:00 AM – 12:30 PM

FEBRUARY 9, 2026

MARCH 9, 2026

APRIL 13, 2026

MAY 11, 2026

Registration required

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

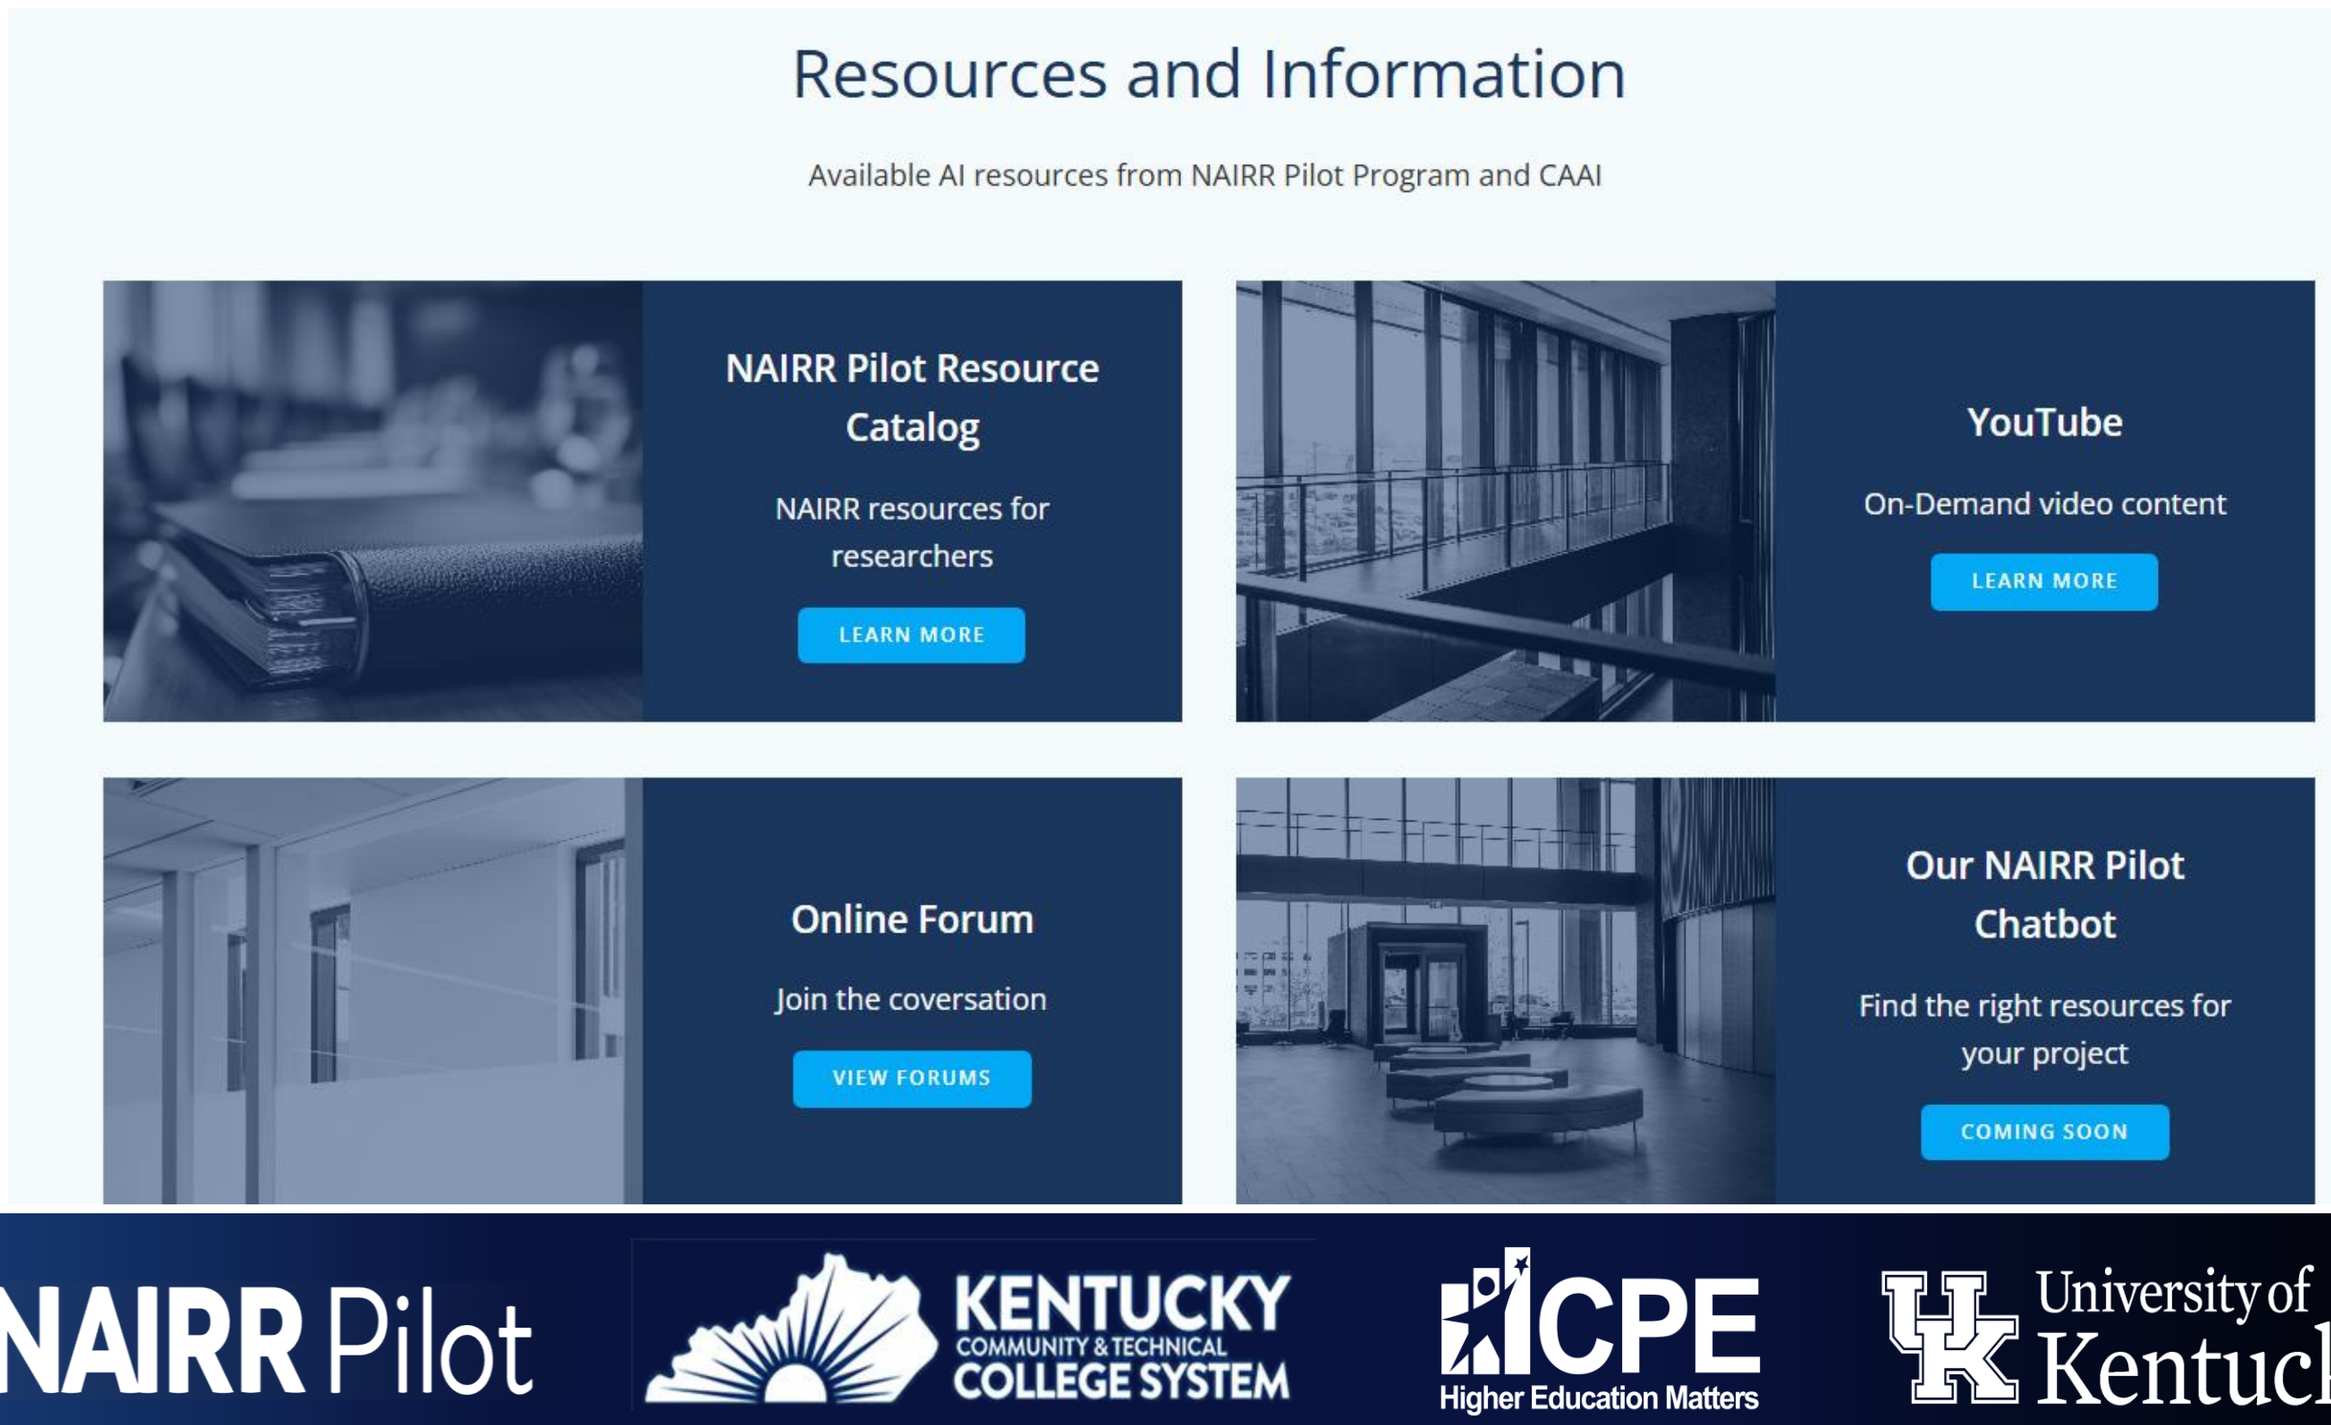
Open-Source



- **CLASSify** – Making Machine Learning Accessible
- **LLM Factory** – Institutionally Controlled Access to LLMs
- **CAT-Talk** – Automated Transcriptions with Timestamps and Speaker-Labels
- **Forecaster** – Evaluate and Explore Timeseries Data with Machine Learning
- **Vision Foundry** – Pretrain, Adapt and Deploy Foundational Vision Models

Democratizing AI Training & Access Statewide Webpage

https://caai.ai.uky.edu/about-us/news-and-events/nairr-pilot-expansion-project/



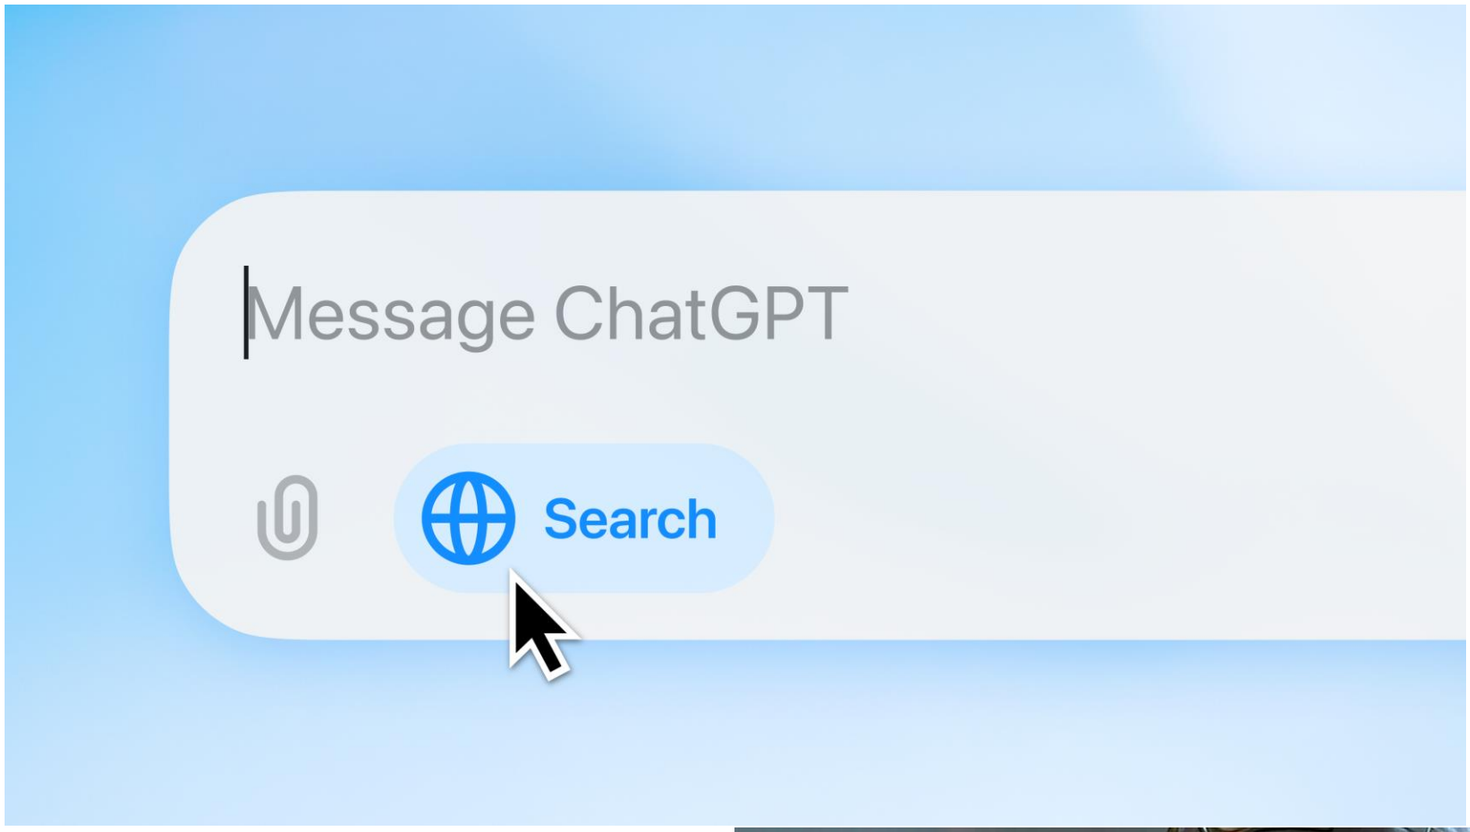
Time	Topic	Speaker
9:15 – 9:20	Welcome	**Dr. Jennifer Lindon**
9:20 – 9:30	Workshop Goals	**Caroline Leach**
9:30 – 10:15	Introduction to AI & AI-Ready Data	**Aaron Mullen**
10:15 – 10:45	AI, The Workforce & Higher Education	**Travis Muncie (CPE)**
10:45 – 11:15	Machine Learning	**Aaron Mullen**
11:15 – 11:30	Accessing NAIRR Resources	**Caroline Leach**
11:30 – 12:15	**Lunch Break**	
12:15 – 1:00	Host Region Presentation	**Danna Barnett & Rusty Justice**
1:00 –1:30	Generative AI & Vibe Coding	**Vaiden Logan**
1:30 – 2:00	Foundational Models & LLMs	**Mitchell Klusty**
2:00 - 2:30	Computer Vision	**Vaiden Logan**
2:30 – 2:45	**Coffee & Peer Networking**	
2:45 – 3:30	AI Pipelines	**Mitchell Klusty**
3:30 - 4:00	Applied AI Panel	**Subject Matter Experts**

Introduction to AI & AI-Ready Data

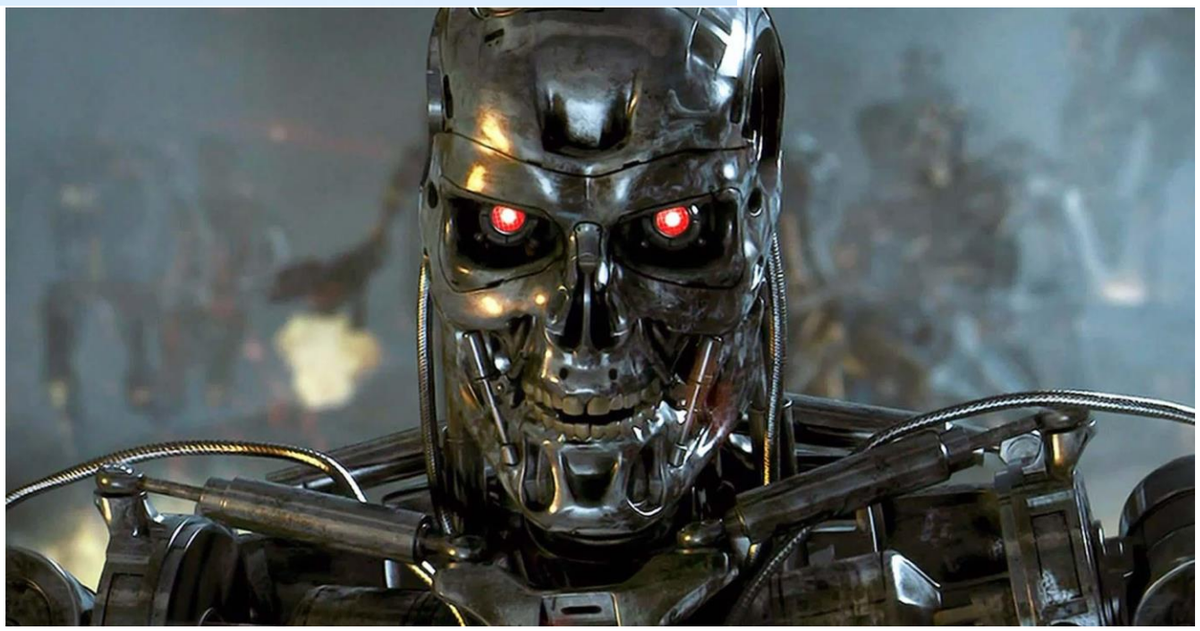
Aaron Mullen, MS

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

What do you think of when you think of AI?







AI is a Moving Target

- AI is more of a broad category than a specific goal
- Optical Character Recognition (OCR) can convert images of text/tables into raw information
  - Example: scanning a check
  - When OCR was developed, it seemed like “magic AI”
  - Now, it’s more or less just a software algorithm
- Thus, AI can have many different meanings and applications

Traffic Predictions

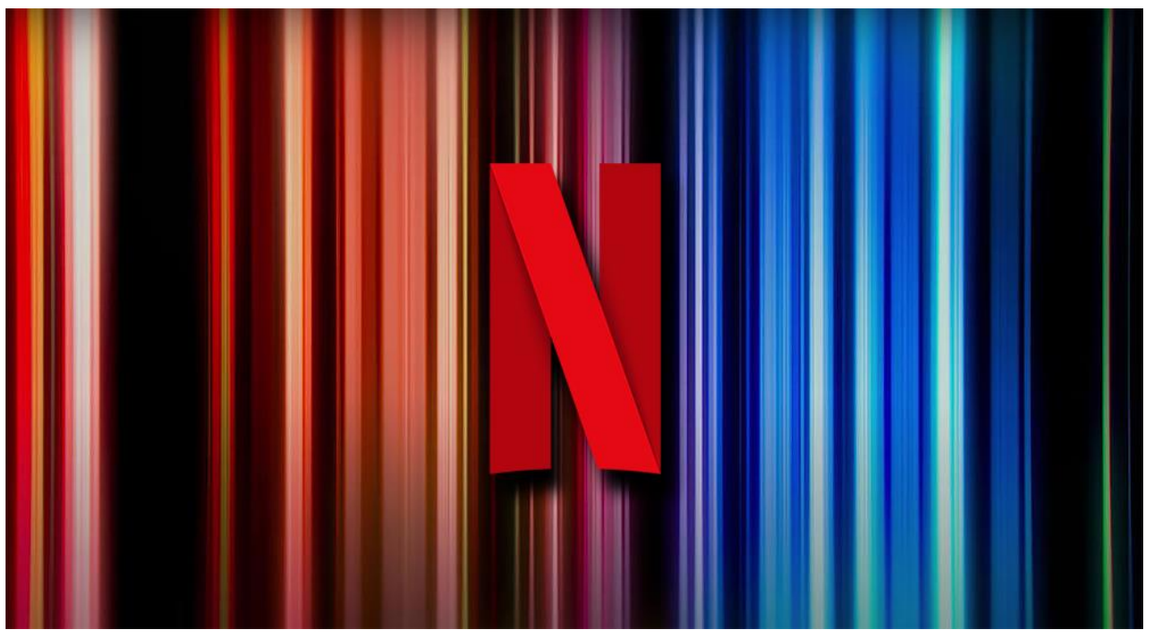
- Route planning apps don't just use algorithms to find the shortest route between two places
  - This could be inaccurate because traffic patterns can change during your drive
- Historical traffic patterns and real time user updates can make predicted ETAs more accurate by forecasting future conditions



https://blog.google/products-and-platforms/products/maps/google-maps-101-how-ai-helps-predict-traffic-and-determine-routes/

Content Personalization

- Streaming services use AI to personalize recommendations to users based on their interests
  - –But this goes beyond just modifying what shows are displayed on the front page
- To promote the last season of House of Cards, Netflix created 10 different trailers that emphasized different aspects of the show
  - –The trailer that was shown to a user on the interface depended on that user's interests



https://research.netflix.com/research-area/recommendations

Improved Crop Management



- Modern John Deere tractors are equipped with cameras that utilize advanced computer vision systems
- These systems have been trained to identify weeds and filter out regular crop plants
  - Herbicide can be intelligently sprayed only where its needed, protecting crops and reducing chemical usage

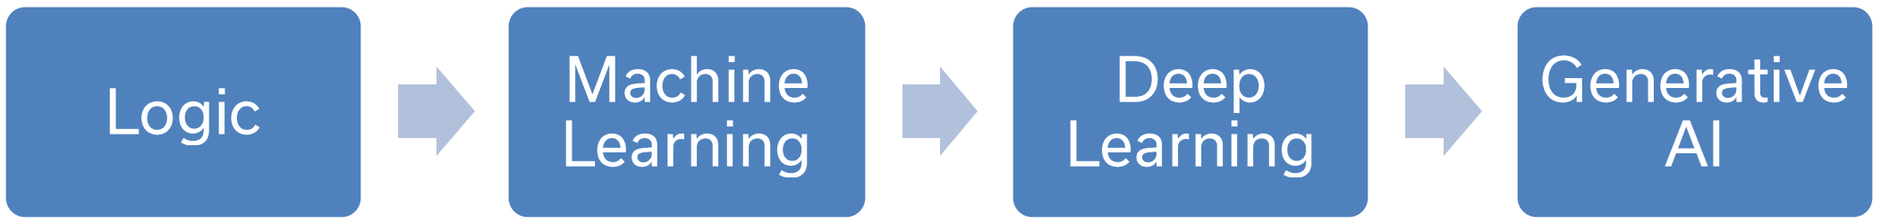
https://about.deere.com/en-us/our-company-and-purpose/technology-and-innovation/sense-and-act

So overall, what is AI used for?

- AI is common in both everyday uses (home voice assistants) and groundbreaking technical advances (predicting protein structures to accelerate drug discovery)
- What do all of these uses have in common?
  - Analyzing and making sense of data
  - Responding to unique and varied situations
  - Automating manual burdens to improve outcomes

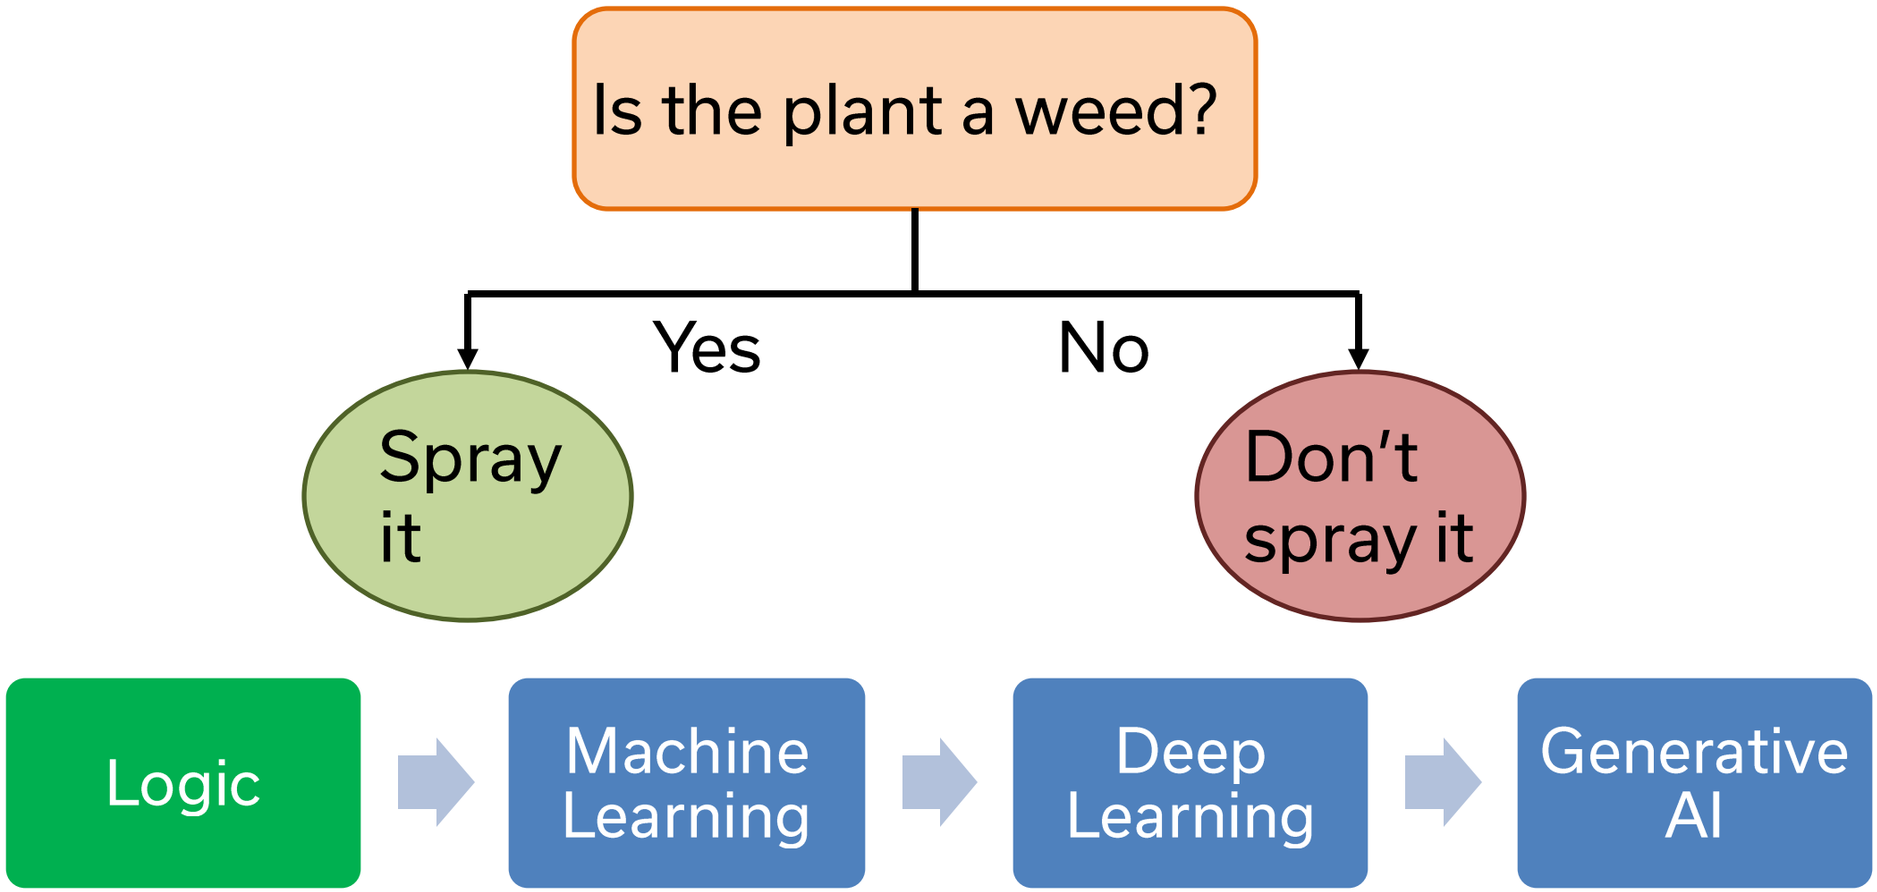
How did we get here?

- Let’s look at how decision-making systems were developed over time to lead to modern AI

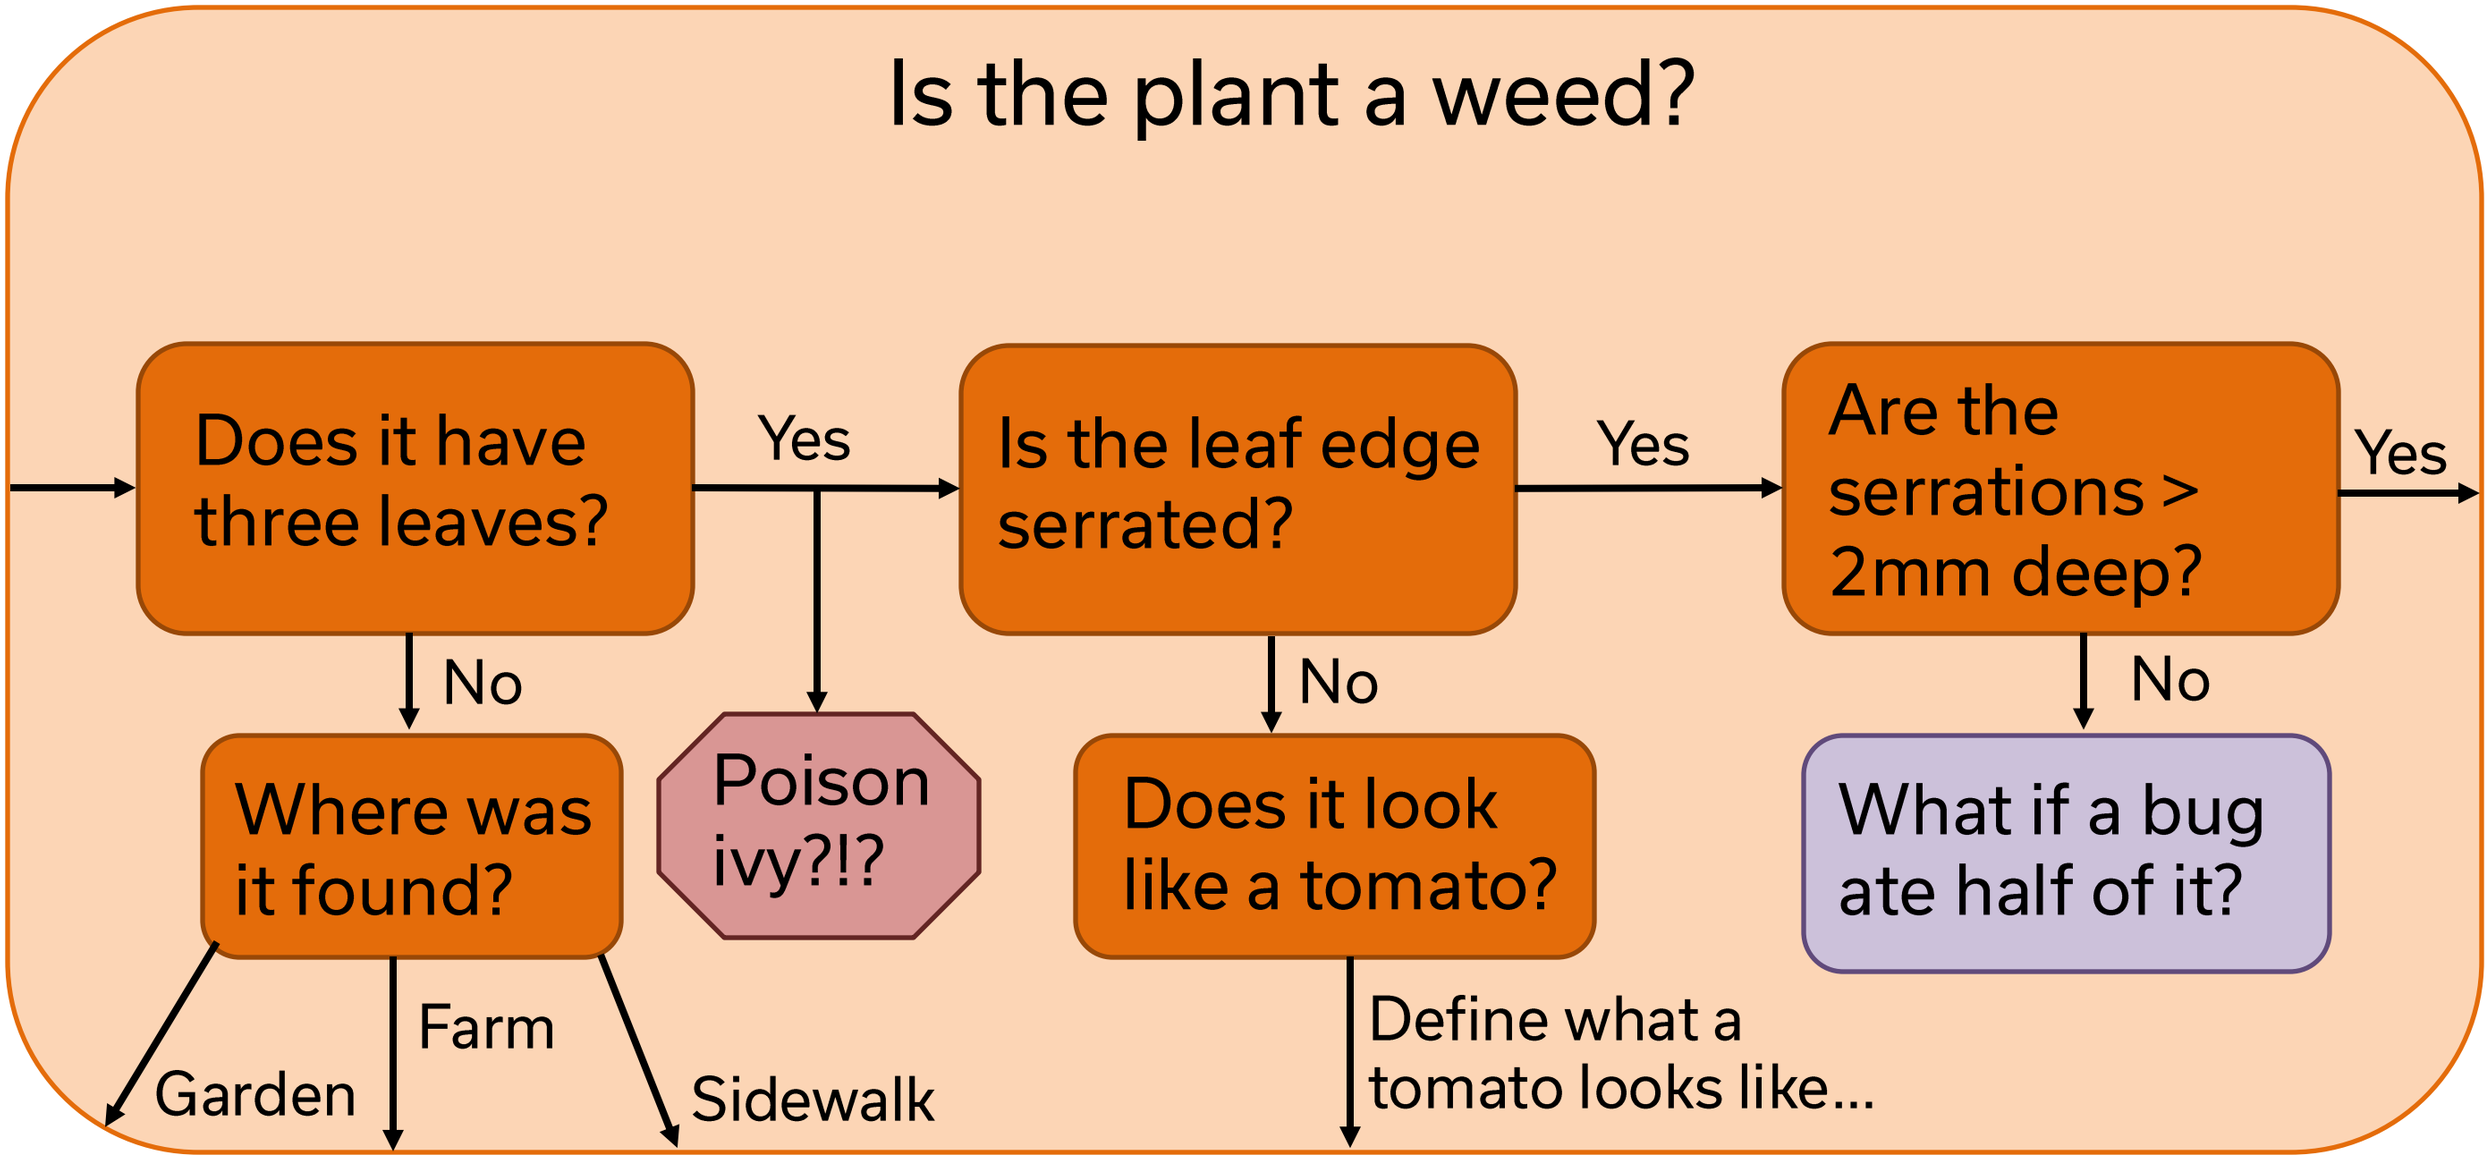


The Age of Logic (1950s-1980s)

- The first computers were being developed, and with them, techniques for computer-based decision-making
- Systems were meticulously coded in an “if-then” format of explicit rules



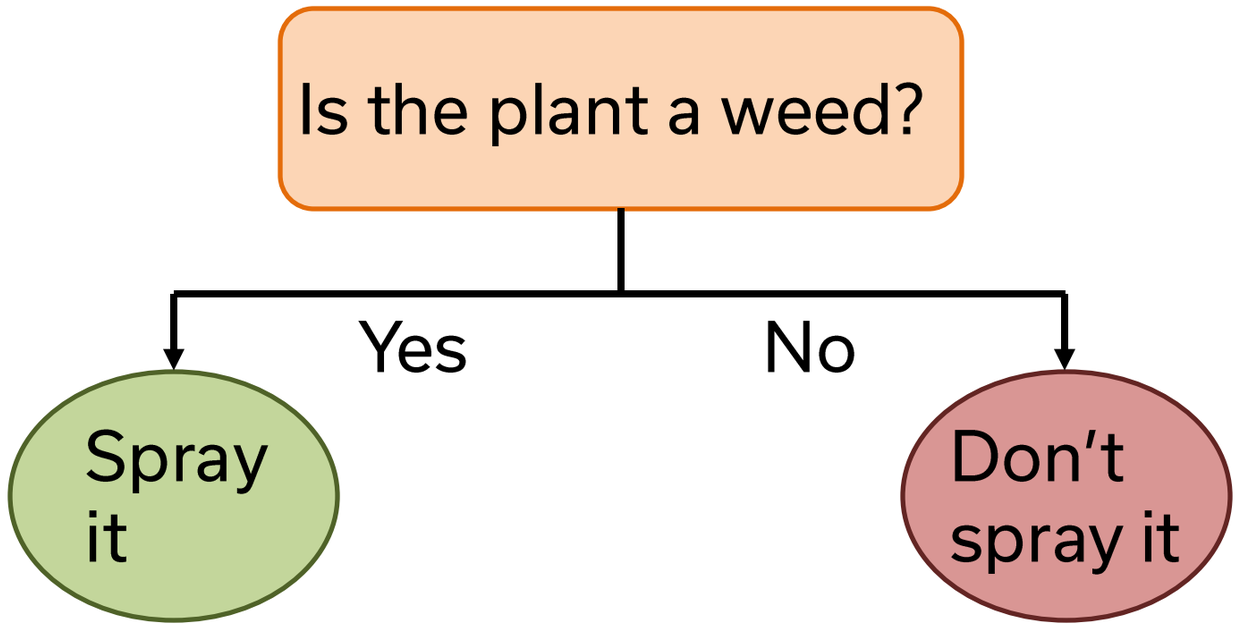
The Age of Logic (1950s-1980s)



Logic → Machine Learning → Deep Learning → Generative AI

The Age of Logic (1950s-1980s)

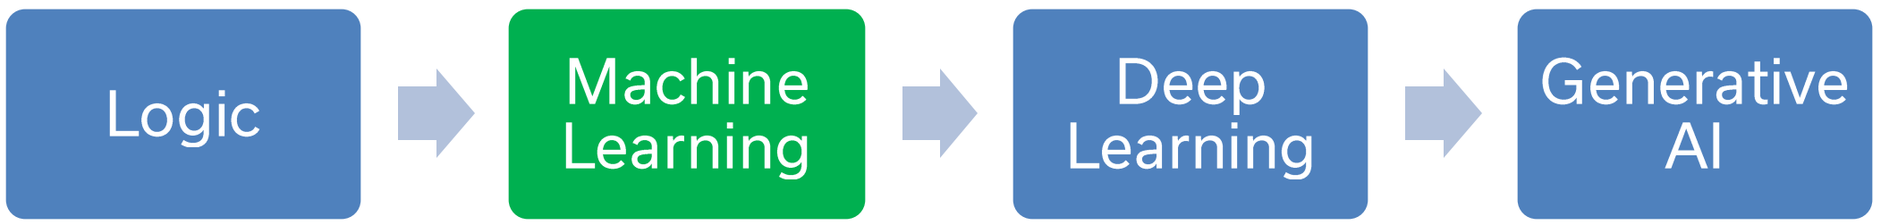
- These systems are very rigid and require definitions ahead of time for every possible scenario that could be encountered
- But they proved that computers can make logical decisions in narrow fields using human knowledge



Logic → Machine Learning → Deep Learning → Generative AI

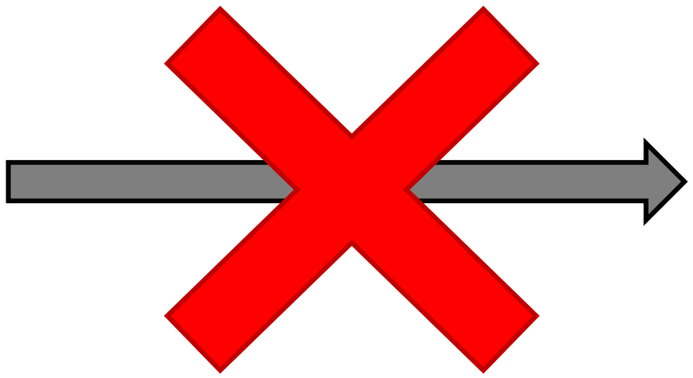
The Age of Machine Learning (1980s-2010s)

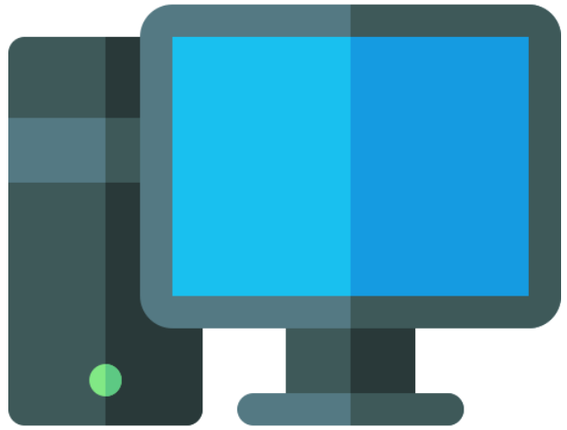
- Computers have progressed more and can make determinations without definitions of specific rulesets
- However, computers are too slow to process much raw data
- So, feature engineering is still required to limit the data the computer has to analyze
  - Must define what aspects of the problem are important for the computer to consider



The Age of Machine Learning (1980s-2010s)



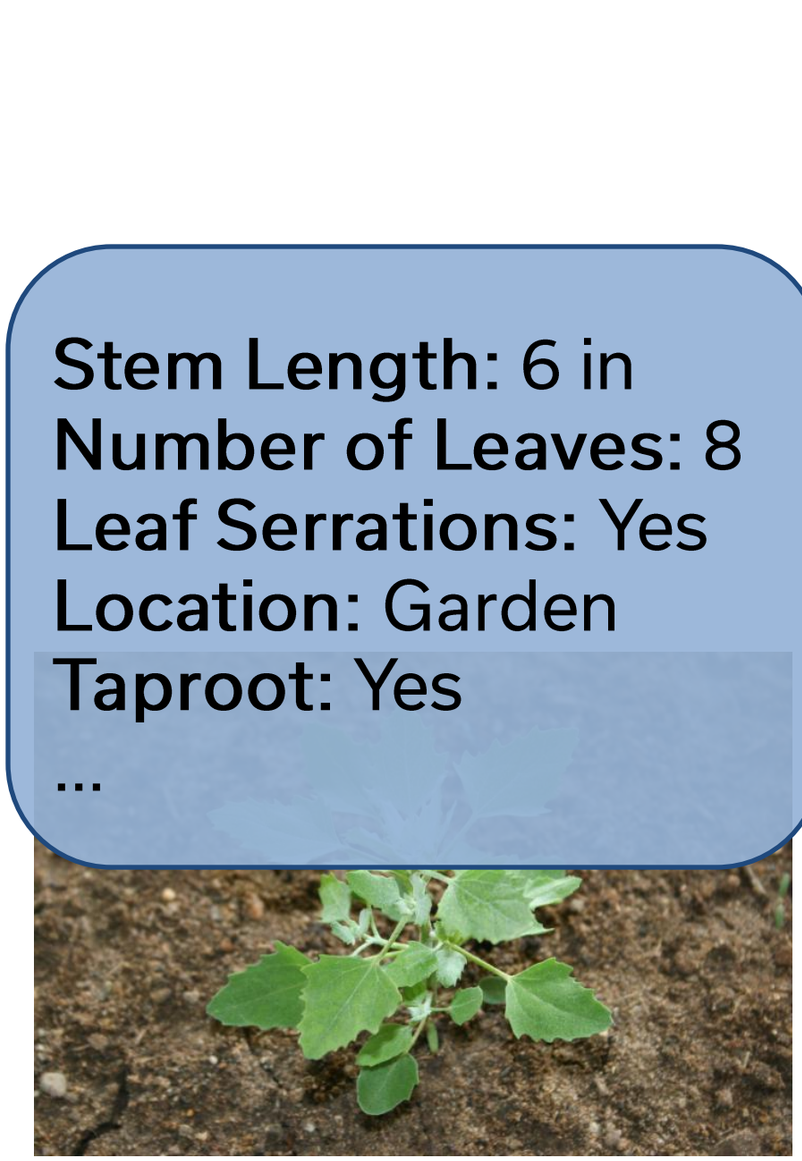


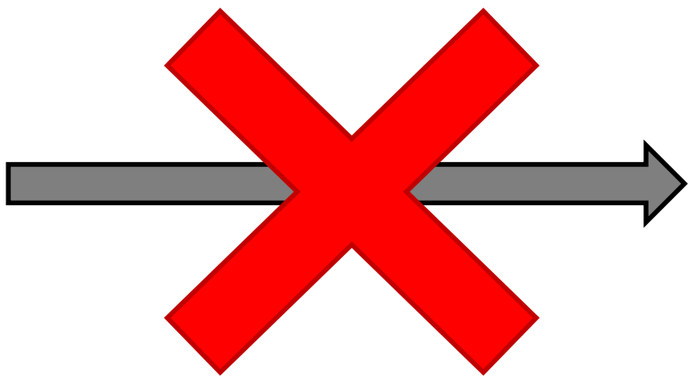


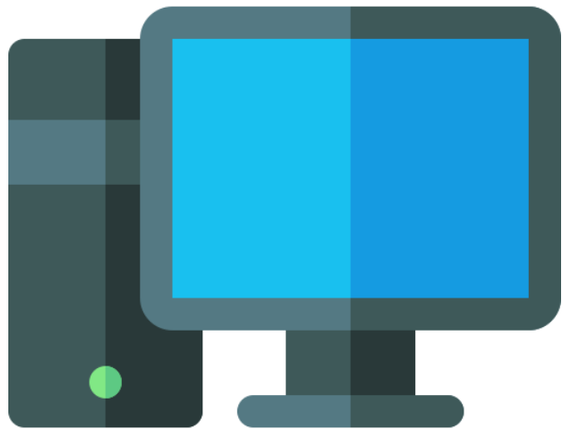
Logic → **Machine Learning** → Deep Learning → Generative AI

The Age of Machine Learning (1980s-2010s)

Stem Length: 6 in
Number of Leaves: 8
Leaf Serrations: Yes
Location: Garden
Taproot: Yes
...



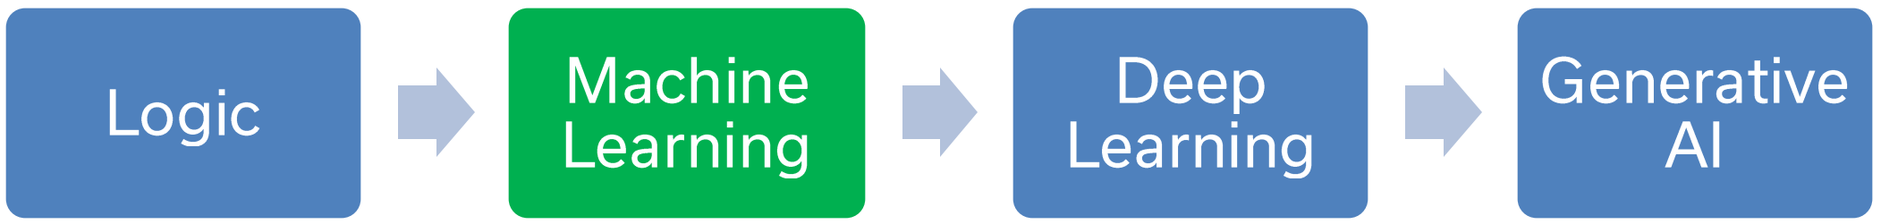




Logic → Machine Learning → Deep Learning → Generative AI

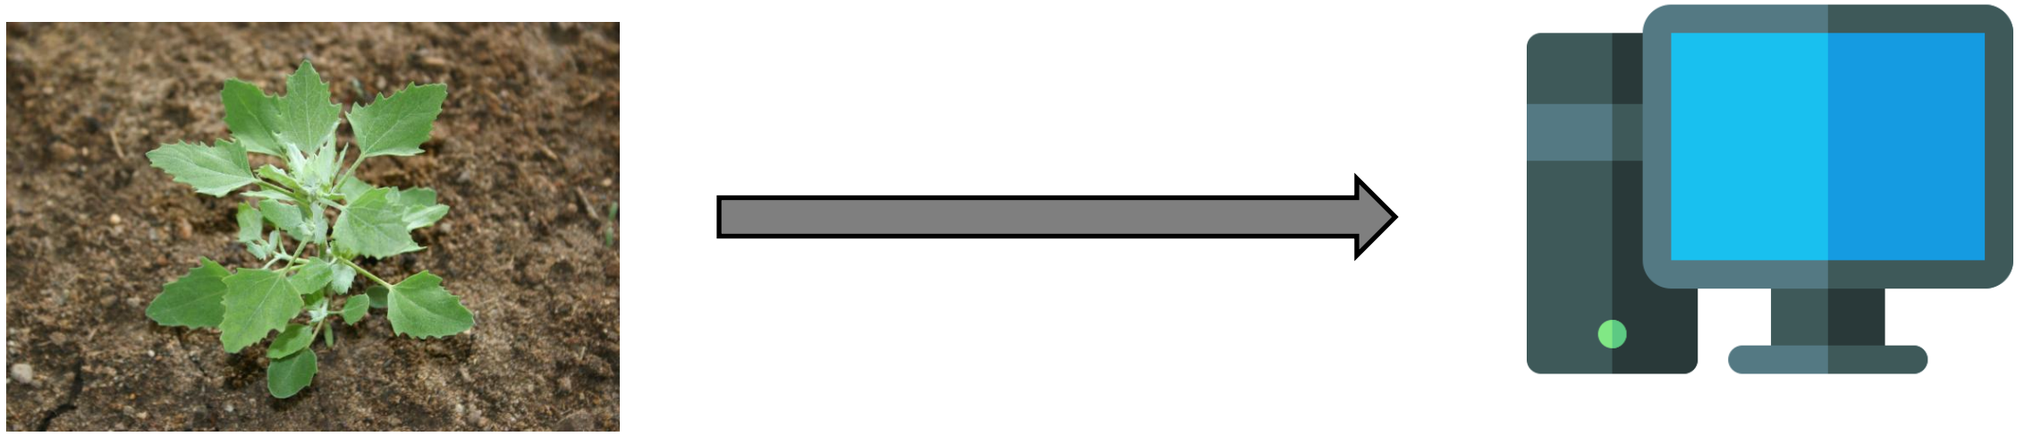
The Age of Machine Learning (1980s-2010s)

- More flexibility than rule-based systems because computers can determine which features are important for the task
  - –Conditions for every possible scenario no longer need to be defined
  - –New models like support vector machines and simple neural networks
- Human involvement is still required to define features of data



The Age of Deep Learning (2010s-2020s)

- Computers and GPUs have advanced to allow for much more complex interpretations of data
- Now, raw data can be fed directly into a model
  - Deep neural networks and transformers can determine their own sets of features to maximize performance



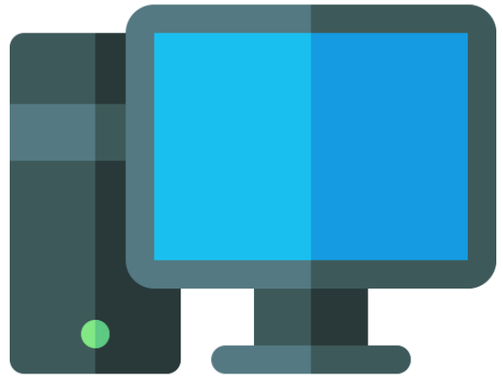
Logic → Machine Learning → **Deep Learning** → Generative AI

The Age of Deep Learning (2010s-2020s)

- The ability of computers to now process and understand these huge quantities of data prompted a new idea:
  - If our model understands everything about what defines a plant as a weed based on images, would it be able to generate its own image of a weed?



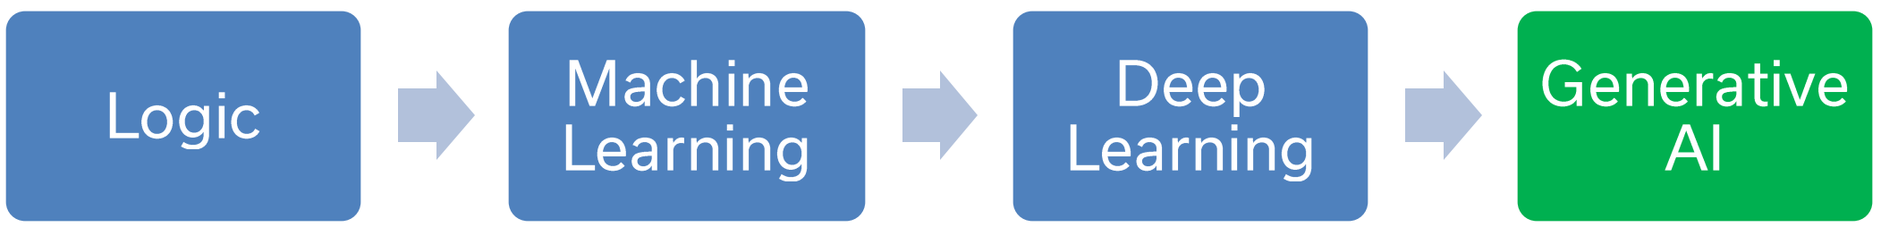




Logic → Machine Learning → **Deep Learning** → Generative AI

The Age of Generative AI (2020s)

- New types of models could be trained using Generative Adversarial Networks (GANs)
  - Two models: a generator (creates fake images/text) and a discriminator (tries to classify which objects are real or fake) pitted against each other
- Transformers developed with "attention" mechanisms
  - A model can look at an entire sentence or image and understand how every aspect relates to each other at the same time
  - Models can be trained on huge amounts of data and learn from all of it



Where do we go from here?

- The meaning of AI and machine learning has changed significantly over time
  - Techniques possible today seemed like science fiction 10 years ago
- The development of new AI techniques has only gotten faster
- How will AI continue to change?

AI Today is Narrow

- Narrow AI: Designed for one specific task
- Can sometimes outperform human expertise/ability at those tasks
  - AI systems tasked with detecting breast cancer from X-rays have been able to reduce both false positives and false negatives compared to human doctors
- However, these systems have no intelligence outside of their own domains
  - AI that detects breast cancer cannot be used to identify weeds

General AI

- Artificial General Intelligence (AGI): AI that can accomplish new tasks in completely unfamiliar contexts without human intervention
  - Able to reason, plan, and solve problems with same flexibility as a human
- AGI is purely theoretical and does not exist yet
- However, multimodal systems are expanding the definition of "narrow" AI

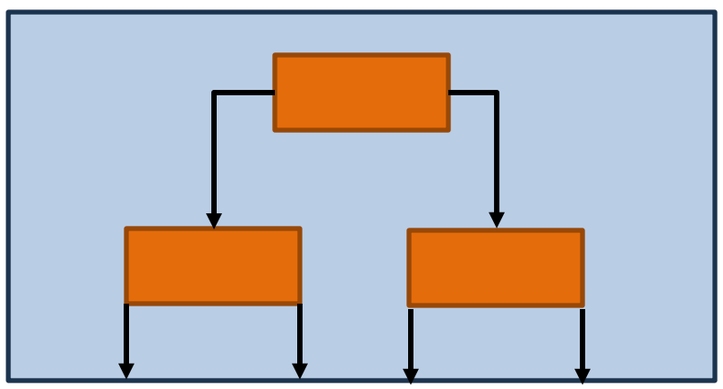
Important Caveats About AI

AI doesn't "know" anything!

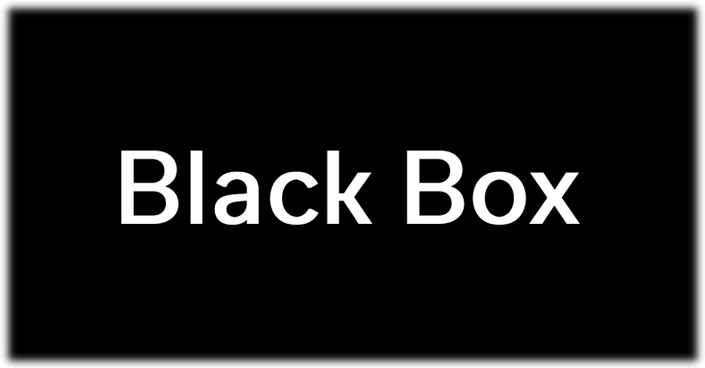
- AI that generates text doesn't "understand" what it's saying in the same way a human does
  - –As it's writing, it's using an algorithm to predict the most likely next word at each step, based on the data it was trained with
- AI that identifies breast cancer doesn't know what breast cancer is
  - –It knows the data features that make a certain response more or less likely to be rewarded as "correct" based on training data
- This doesn't mean that AI can't be useful, just that we must be cautious with interpretations

Transparency and Trust with AI

- Ethical use of AI is crucial as this technology develops
- Training data:
  - When AI is trained, it "copies" the data it is developed with
  - Thus, human biases present in the data will be "copied" as well
  - This will be discussed in more detail later
- AI tends to lack explainability (why does it predict something?)
  - This is a problem if AI is to be used with high stakes decision-making



vs.



Hallucinations

- A chatbot on the Air Canada website told a passenger that he could apply for refunds for a bereavement fare after traveling
- After attempting to do so, he was told bereavement rates could not retroactively apply to completed travel by the company
- Air Canada was forced to pay the customer the difference because the chatbot had misled him
- Raises interesting questions: who is responsible for the content AI produces?

https://www.theguardian.com/world/2024/feb/16/air-canada-chatbot-lawsuit

Ethics of Generative AI

- Deepfakes: how can we trust media we view?
  - –Believing something to be real that is fake
  - –Believing something to be fake that is real
- Copyright conflicts
  - –OpenAI was trained on millions of New York Times articles
  - –Is that learning or plagiarizing?

- AI is neither good or bad, but how we use it can be
- Today, we will learn much more about how to use AI effectively and ethically

Ethics of Generative AI

- AI is powerful which makes it frightening
  - It's important to remember that this power can be and is used for good too

Giving independence to blind people, allowing them to take pictures of anything and listen to conversational descriptions

Allowing patients with ALS to record their voice before losing it to preserve their identity with generative speech

Providing useful tutors to students who can't afford private help

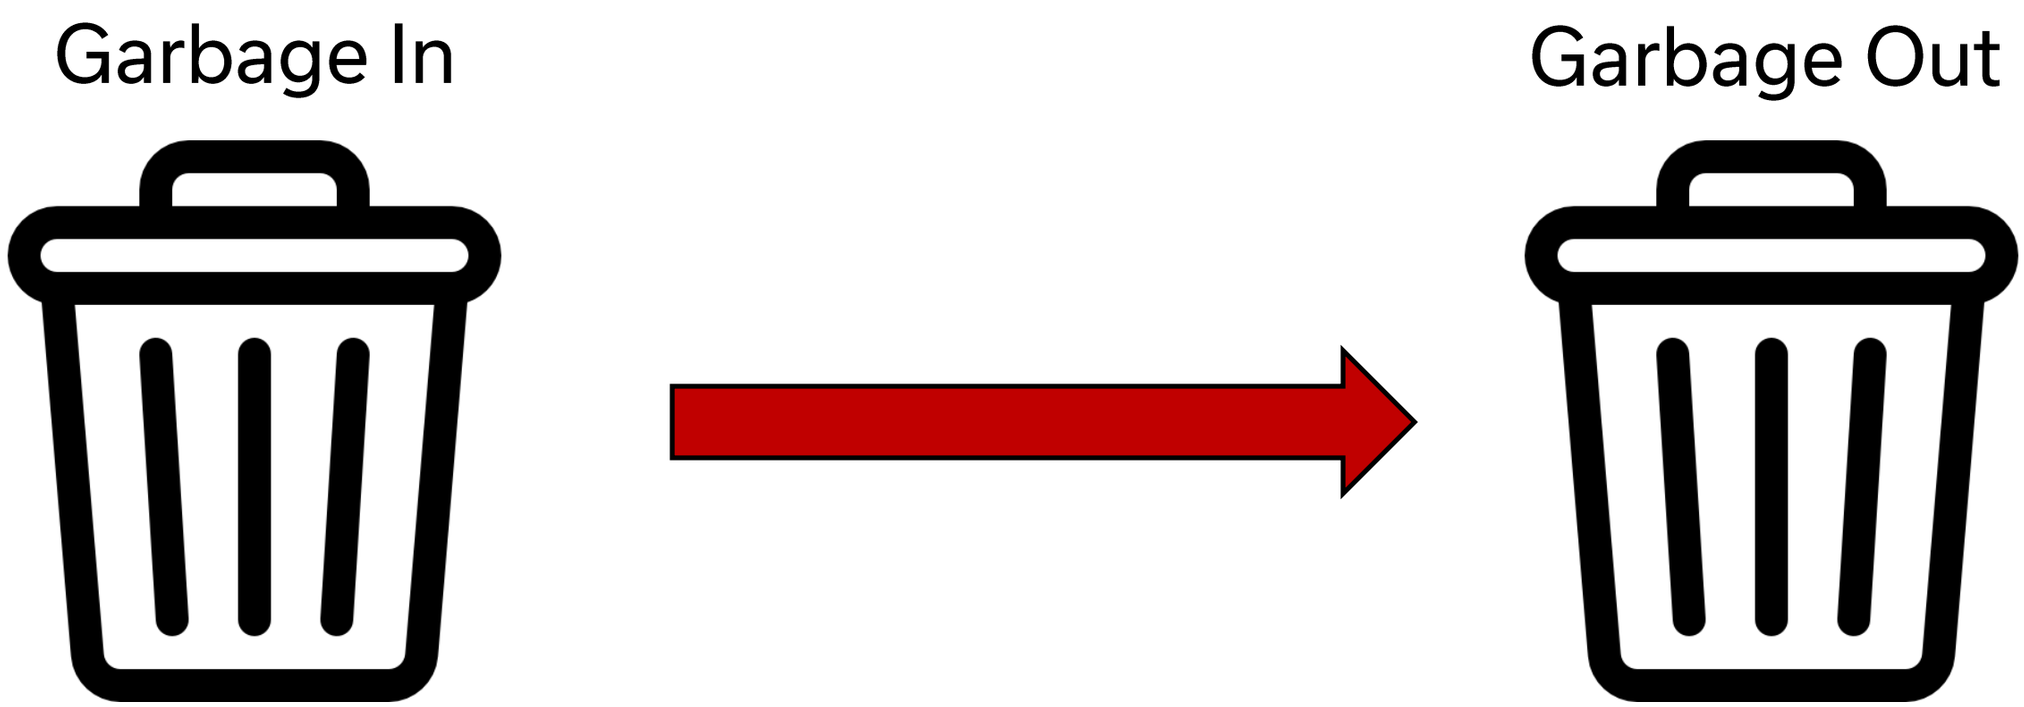
Forecasting floods and natural disasters in developing countries to provide equity in safety

How are these AI systems created?

Data!

Data is important!

- All AI systems require data for training and evaluation
  - –AI learns from data rather than pre-defined rules
- Bad data will lead to a bad AI model



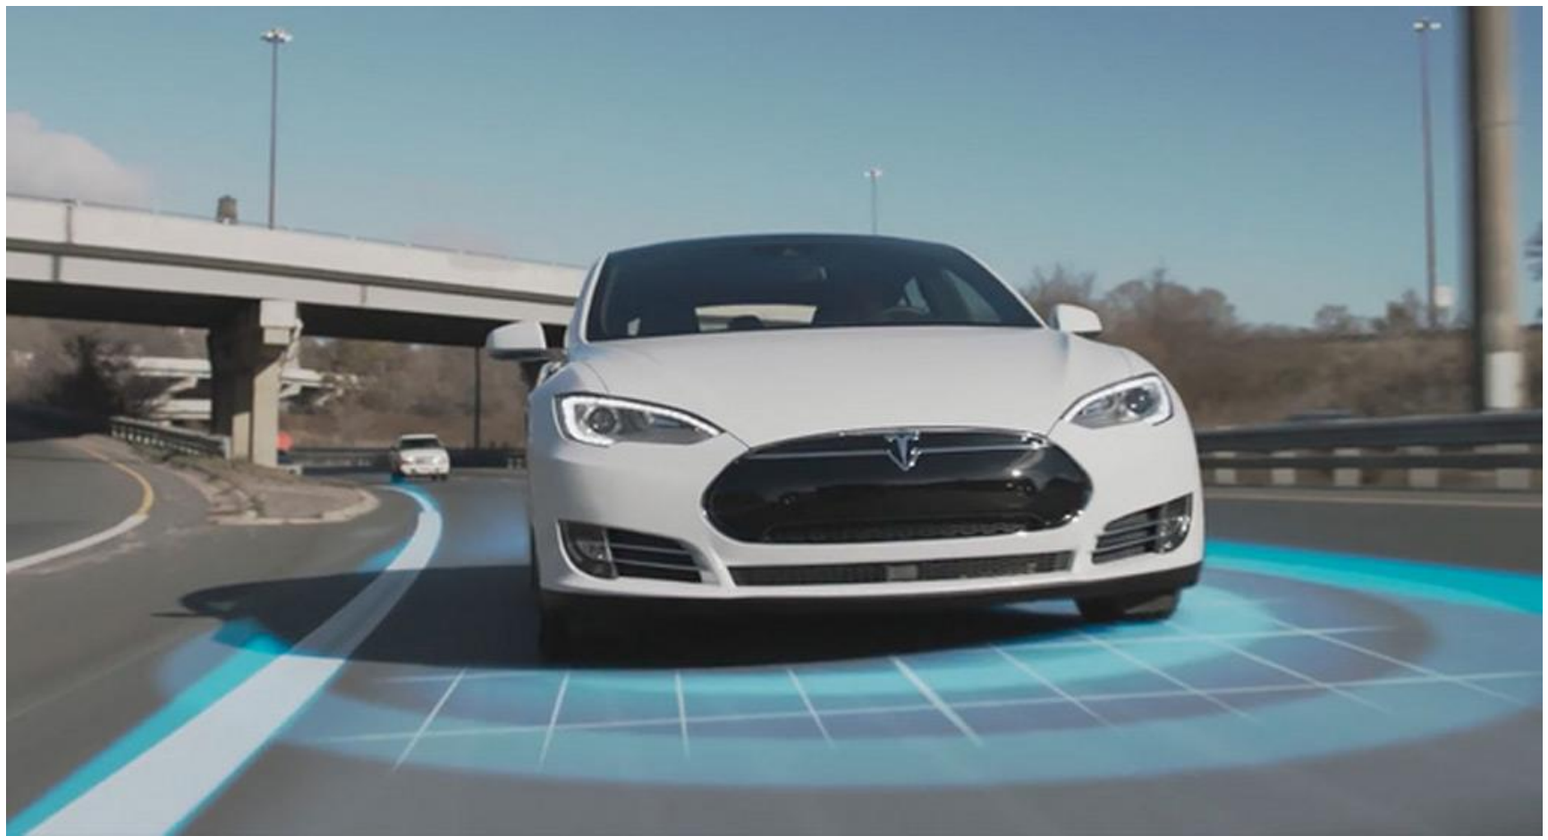
- So what makes data bad?

Rule 1: Data must be sufficient

- If you don’t have enough (good) data, any AI model trained with that data will underperform
  - –Won’t generalize well to unexpected scenarios
  - –Systems only have knowledge of what they were trained with

Rule 1: Data must be sufficient

- In 2016, Tesla' autopilot caused a fatal accident when it failed to detect a tractor trailer turning in front of the car
  - –The white side of the trailer blended in with the brightly lit sky
  - –This rare condition was not accounted for in the autopilot's training data



https://www.theguardian.com/technology/2016/jun/30/tesla-autopilot-death-self-driving-car-elon-musk

Rule 1: Data must be sufficient

- So how much data is sufficient?
  - –It depends!

How narrow is the AI's task?	• Self driving car models must account for many more possibilities than models that scan X-rays
What's the domain of the AI?	• Images vs text vs numbers
How complex is the problem?	• More data is required to account for uncommon/difficult tasks

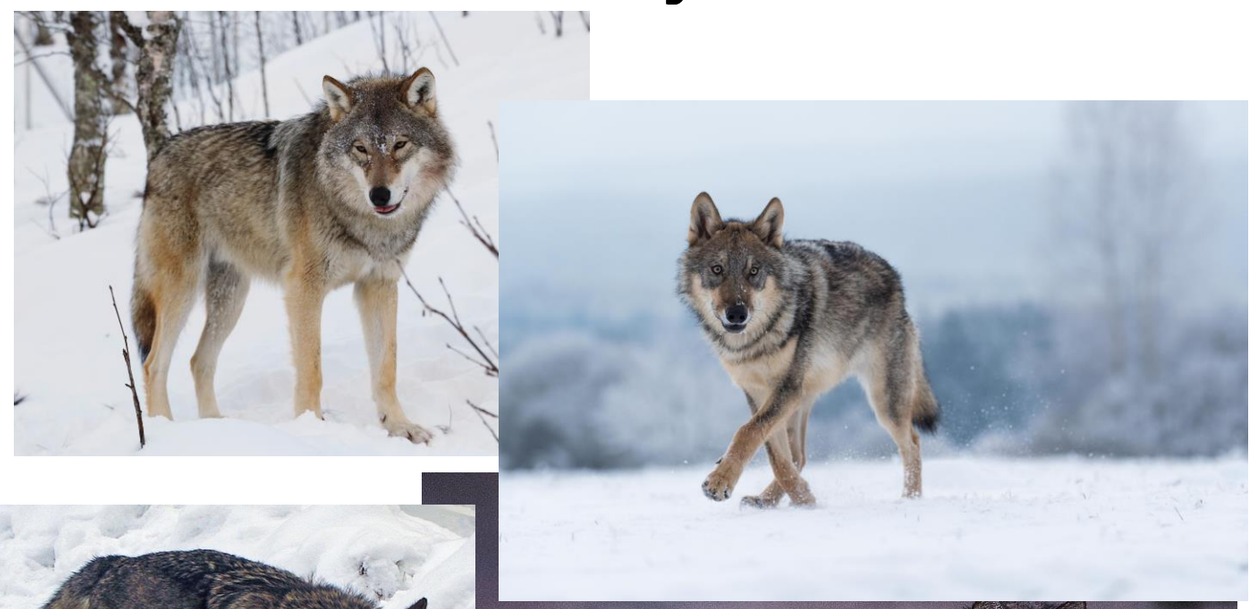
Rule 2: Data must be free of confounding variables

- Researchers built a model to classify images of huskies and wolves

Model says: HUSKY



Model says: WOLF





"Why Should I Trust You?" Explaining the Predictions of Any Classifier. Ribeiro M, Singh S, Guestrin C.

Rule 2: Data must be free of confounding variables

- They found that in testing, the model did not work great

Model says: WOLF



Model says: HUSKY



Rule 2: Data must be free of confounding variables

- Could the model be learning something unintended?

Model says: HUSKY



All green/grassy!

Model says: WOLF



All white/snowy!

Rule 2: Data must be free of confounding variables

- Rather than being able to tell huskies from wolves, the model could just tell snowy landscapes from grassy ones
- **Lesson**: closely examine all aspects of your data
  - –Consider factors aside from your primary focus that could affect results

Model says: WOLF



Model says: HUSKY



Rule 3: Data must be unbiased

- AI will reflect any biases present in the training data
- These biases must be considered and removed to create fair and equitable AI systems
- AI requires human data, and unfortunately, human data is very rarely completely unbiased

Rule 3: Data must be unbiased

- Amazon wanted to automate recruiting new employees
  - Candidates would be ranked based on their resumes
  - Training data consisted of past employee performance, as well as who got hired during previous recruitments vs who did not
- They discovered this recruitment tool immediately favored male candidates over female candidates
  - Some resumes would be discarded simply because the candidates went to women's colleges or participated in women's clubs
- The AI system itself was not inherently discriminatory
  - But the data that the system was trained on (Amazon's past hiring practices) contained these biases, and the system simply reflected those

https://www.hubert.ai/insights/why-amazons-ai-driven-high-volume-hiring-project-failed

Rule 4: Data must be appropriate in scope



Rule 4: Data must be appropriate in scope

- IBM wanted Watson's capabilities to extend to real-world problems
- They developed Watson for Oncology to analyze medical literature, treatment protocols, and emerging research
  - –It could recommend treatment plans and assist oncologists
- It was trained on data from Memorial Sloan Kettering Cancer Center
  - –High quality data, but reflected only patient demographics and treatment practices for that specific area

https://www.henricodolfing.ch/case-study-20-the-4-billion-ai-failure-of-ibm-watson-for-oncology/

Rule 4: Data must be appropriate in scope

- **The problem**: Watson for Oncology was heavily marketed in traditionally under-served regions like Thailand, India, and South Korea
- It would give unsafe or irrelevant recommendations because important demographics and treatment practices were different
  - –Certain drugs may not be available
  - –A different patient population will react differently to different treatments
- Watson for Oncology was discontinued in 2023

Rule 4: Data must be appropriate in scope

- **The lesson**: the scope of your data should match the scope of your AI system's intended usage
- If Watson had only been used for MSKCC patients, it would have likely been seen as a success
  - –But expansion to new patient populations required new data
- A model trained on local data only is not inherently bad!
  - –But the scope of the system's application must remain local

So good data is…

Large enough for your domain/problem

Free from additional confounding variables

Without human biases

Appropriate to the scope of your usage

Good Data Example

- Google developed an AI-powered diabetic retinopathy detection model
- It was trained on data from underserved countries like Thailand and India with the hope of providing care where healthcare resources are more scarce
- The model achieves 95% accuracy, slightly higher than retina specialists
- Over 600,000 screenings have already been performed with this system
  - Millions more are planned over the next decade

https://blog.google/company-news/inside-google/around-the-globe/google-asia/arda-diabetic-retinopathy-india-thailand/

What made this dataset good?

It was large- initial training data consisted of over 128,000 images

It was high quality- standardized labeling determined by expert teams to ensure accurate ground truth

It was appropriate to its scope- the models have only been applied in countries that data was used from for training

What does AI-ready data actually look like?

- Data can take many different forms depending on the domain and intended goal
- Three main types:
  - Structured
  - Unstructured
  - Semi-structured

Structured Data

- Typically takes the form of rows and columns in a spreadsheet or relational database
- Quantitative rather than qualitative

Structured Data Use Cases

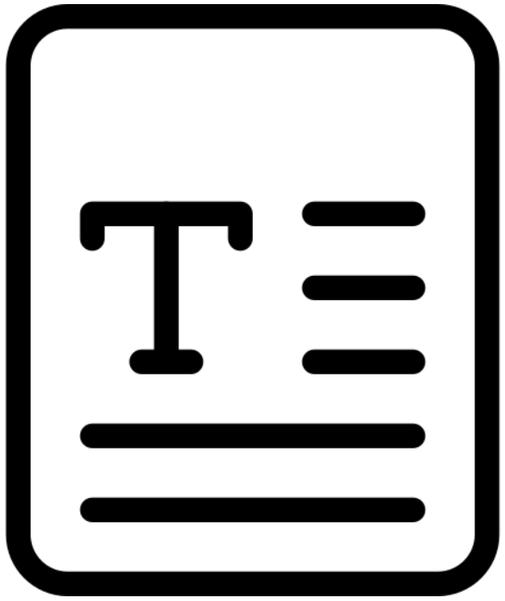
Type of Application	Type of Data	Example
Forecasting	Time series data- measurements gathered at regular intervals in a specific order	A hospital wants to forecast how many patients will come to the ER on a given day to determine staffing and transfer decisions
Classification	Tabular, cross-sectional data- snapshot of entities at a single point in time	An educator wants to predict which students will pass/fail their final exam, based on prior performances, so assistance can be given
Regression	Similar to classification, but target is a number rather than a category	Ride-sharing apps use regression algorithms to determine dynamic pricing based on a variety of factors
Anomaly Detection	Highly granular data generated by actions or events, where a small proportion of rows/events need to be flagged	Banks analyze credit card purchases to determine when transactions are likely to be fraud

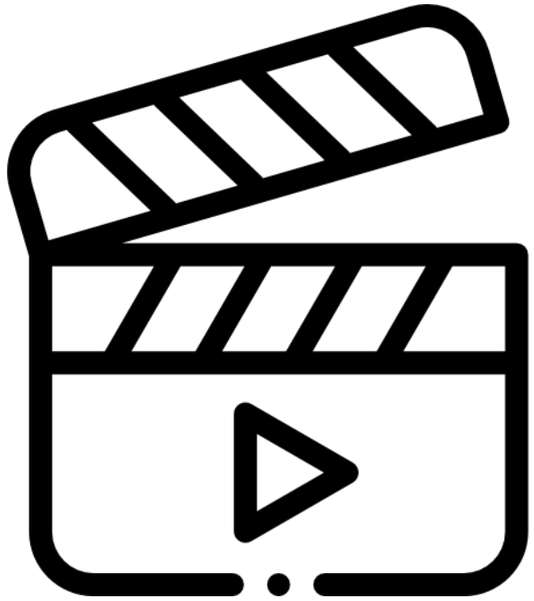
Structured Data Pros/Cons

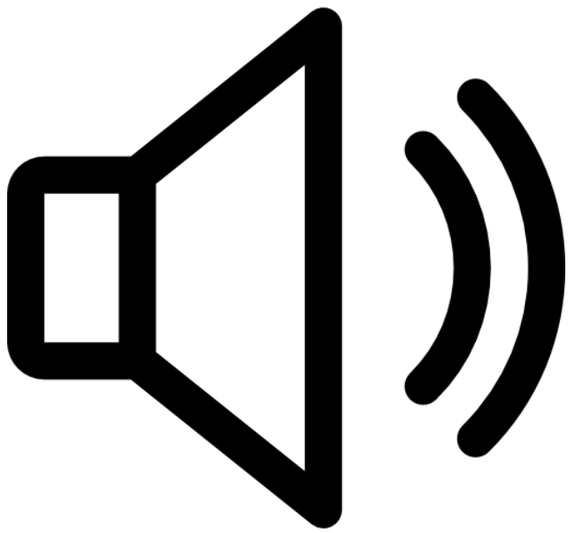
Advantage	Disadvantage
Easy to use due to rigid structure	Much of the world's real data is not gathered in a structured way
Easy to view/interpret as a human	Lots of pre-processing can be required to format data correctly and consistently
Applied to a huge variety of use cases and types of problems	Can fail due to small imperfections, like missing data
Many machine learning techniques are built around structured data	Doesn't work well with most kinds of data aside from numbers/booleans

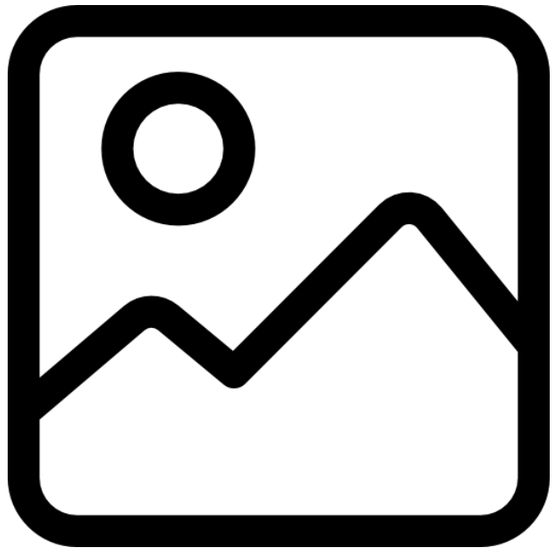
Unstructured Data

- Data that cannot be stored in a traditional database structure

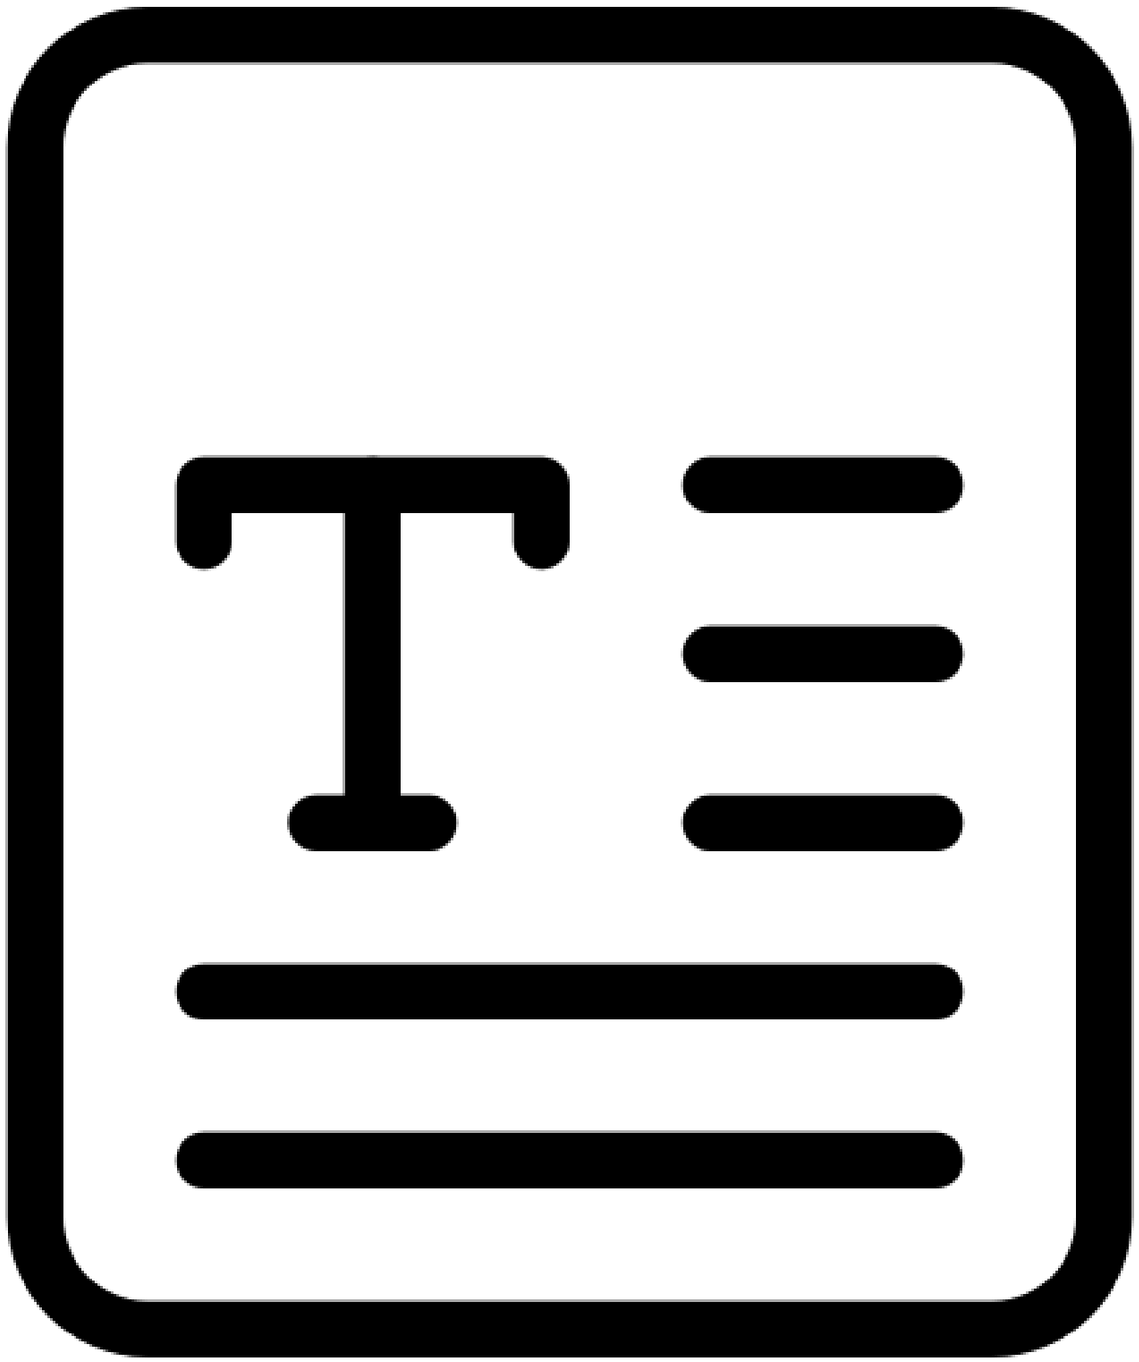






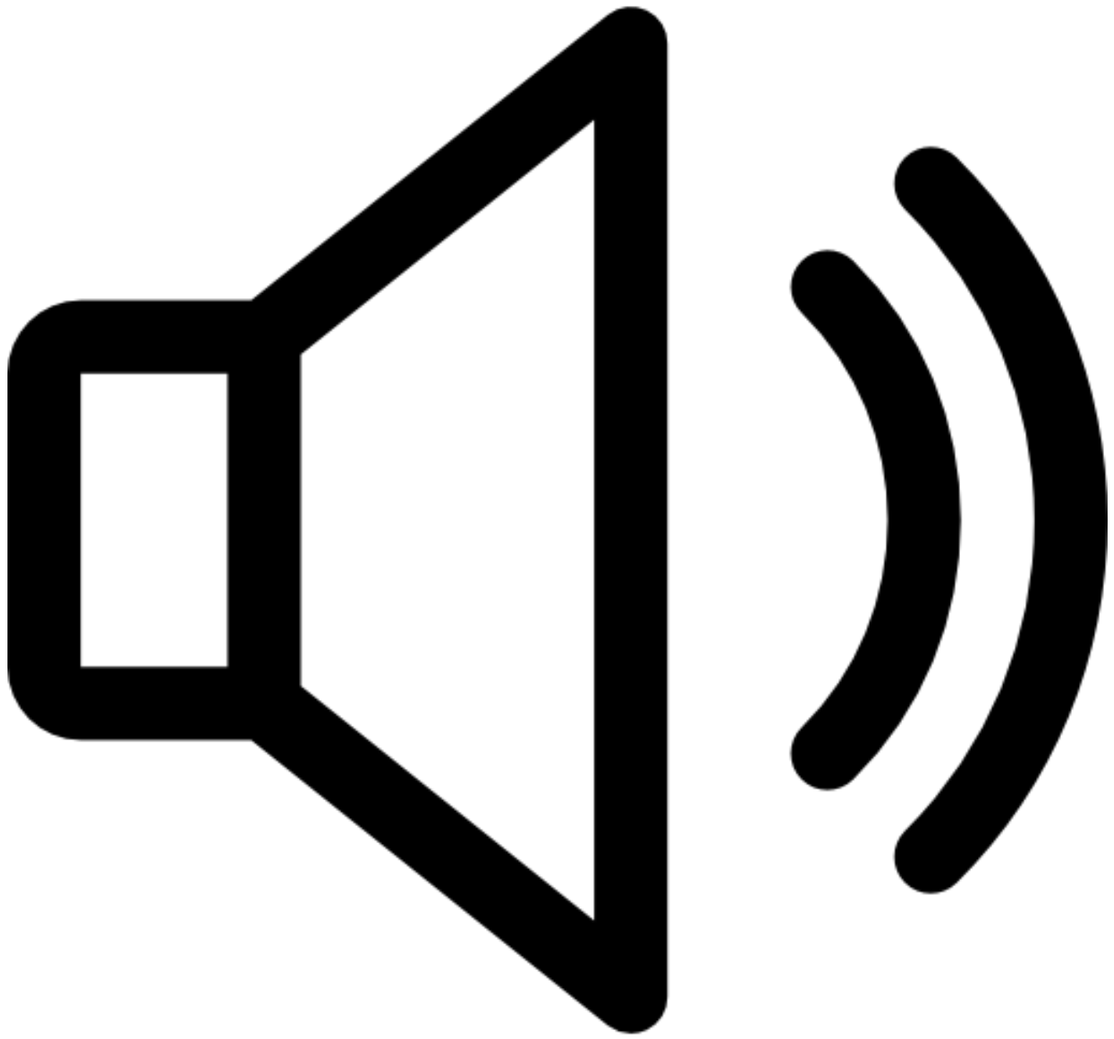


Unstructured Data



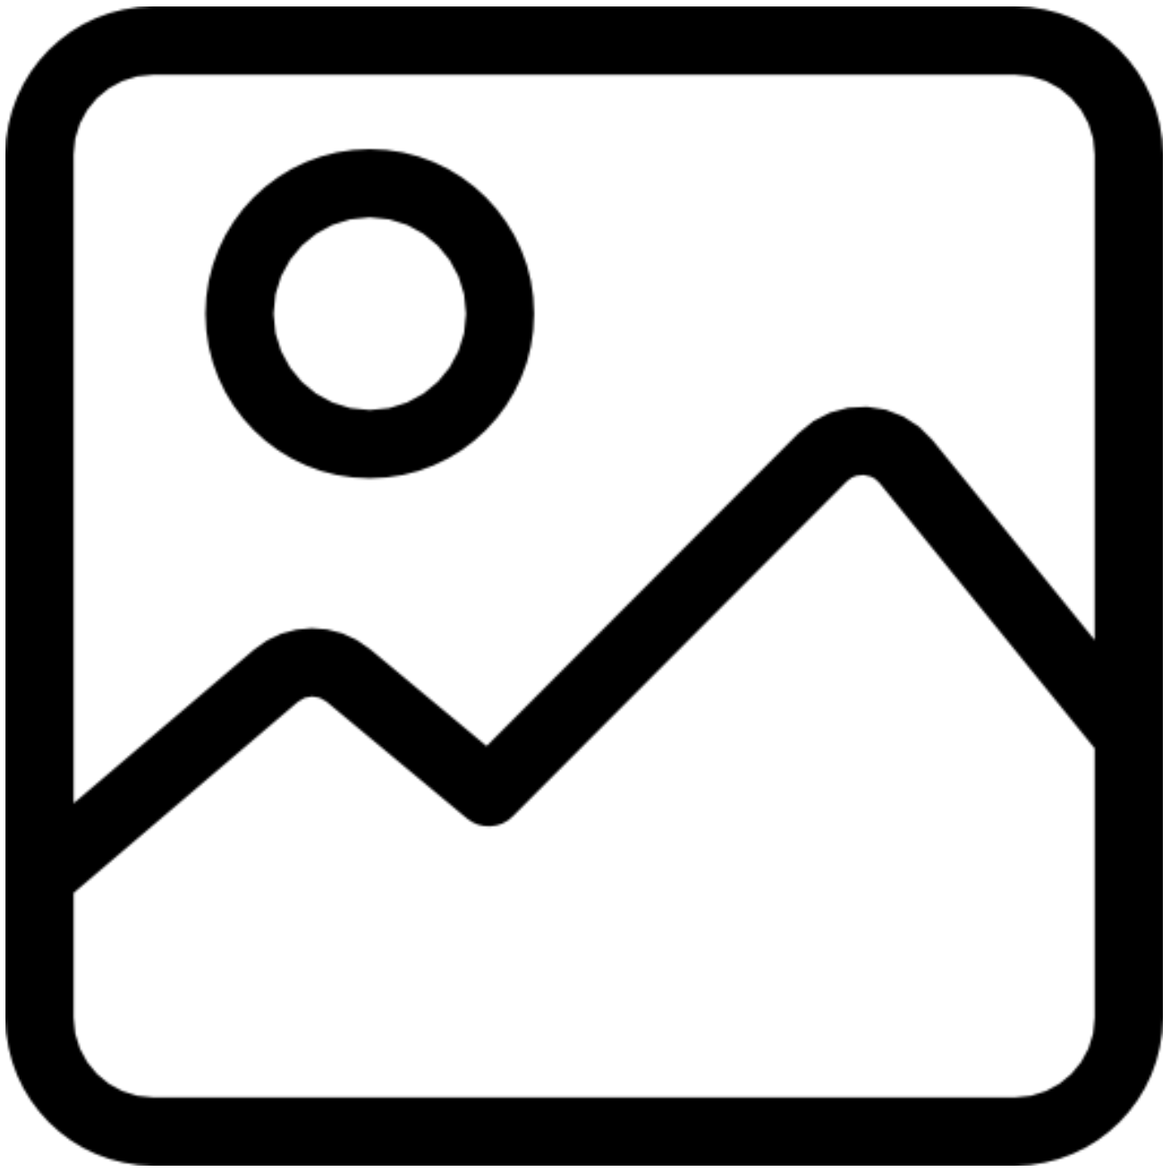
- **Text data**
- Interpretable to humans but viewed by machines as sequence of tokens (words)
- Words/phrases can be embedded into numbers
- Example: a hotel chain analyzes thousands of user reviews to flag complaints without a human reading them

Unstructured Data



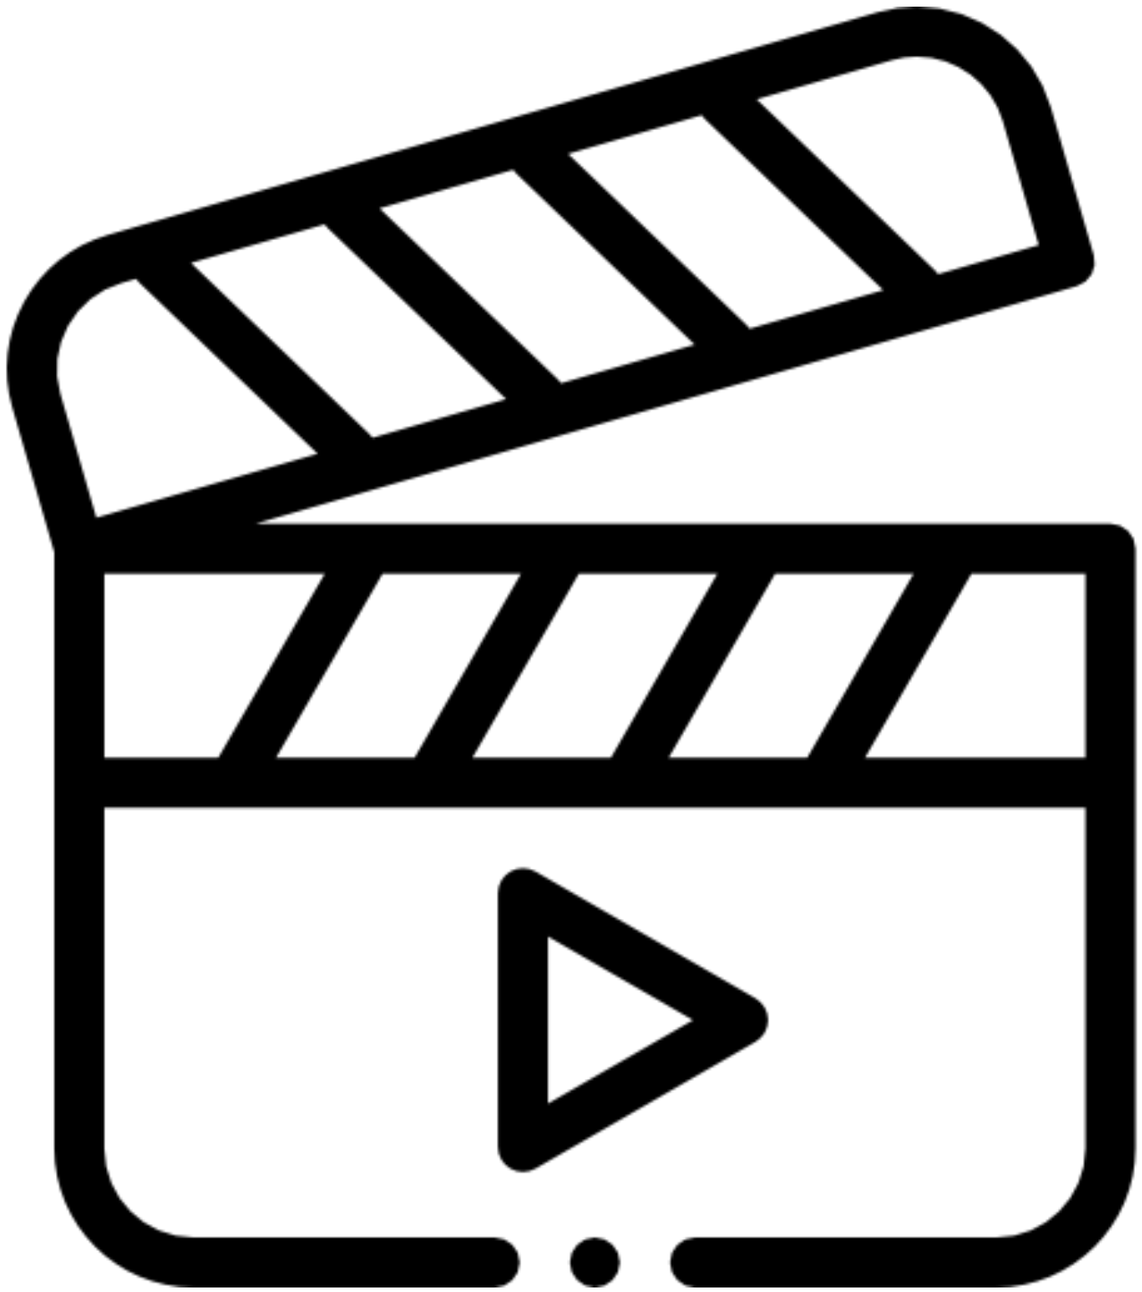
- **Audio data**
- Continuous signal representing air pressure changes
- Measured with attributes like frequency and amplitude
- Example: software that automatically transcribes meetings by listening to dialogue

Unstructured Data



- **Image data**
- To a computer, just a massive grid of numbers representing colors at each pixel
- The goal may be to classify an entire image, or to segment specific parts of it
- Example: X-ray images can be used to identify potential tumors

Unstructured Data



- **Video data**
- Each frame of the video can be analyzed the same as an image
- The combination of these frames, and changes over time, can be used to identify motion
- Example: video monitoring of laboratory rats can be used to identify seizures

Unstructured Data

Pros:

Very common and easy to obtain

Covers a huge range of data types and domains

Understandable and clear to humans

Cons:

Some limited use cases due to lack of structure

Burdensome to manually review/analyze

Frequently very computationally expensive to use

Semi-Structured Data

- Doesn’t follow a rigid table format, but it does follow a looser structure that can still be analyzed and parsed
- Can take on a variety of forms

Semi-Structured Data

<h1> **Websites are built using this HTML format** </h1>

 It's flexible and allows for a lot of variation, but there are still key structured elements, such as nesting and tags

<h2> **Example:** </h2>

 An AI that scrapes the web for ticket prices needs to be able to parse elements from browser HTML code

Semi-Structured Data

{

“title”: “JavaScript Object Notation (JSON) is a very common semi-structured data format”,

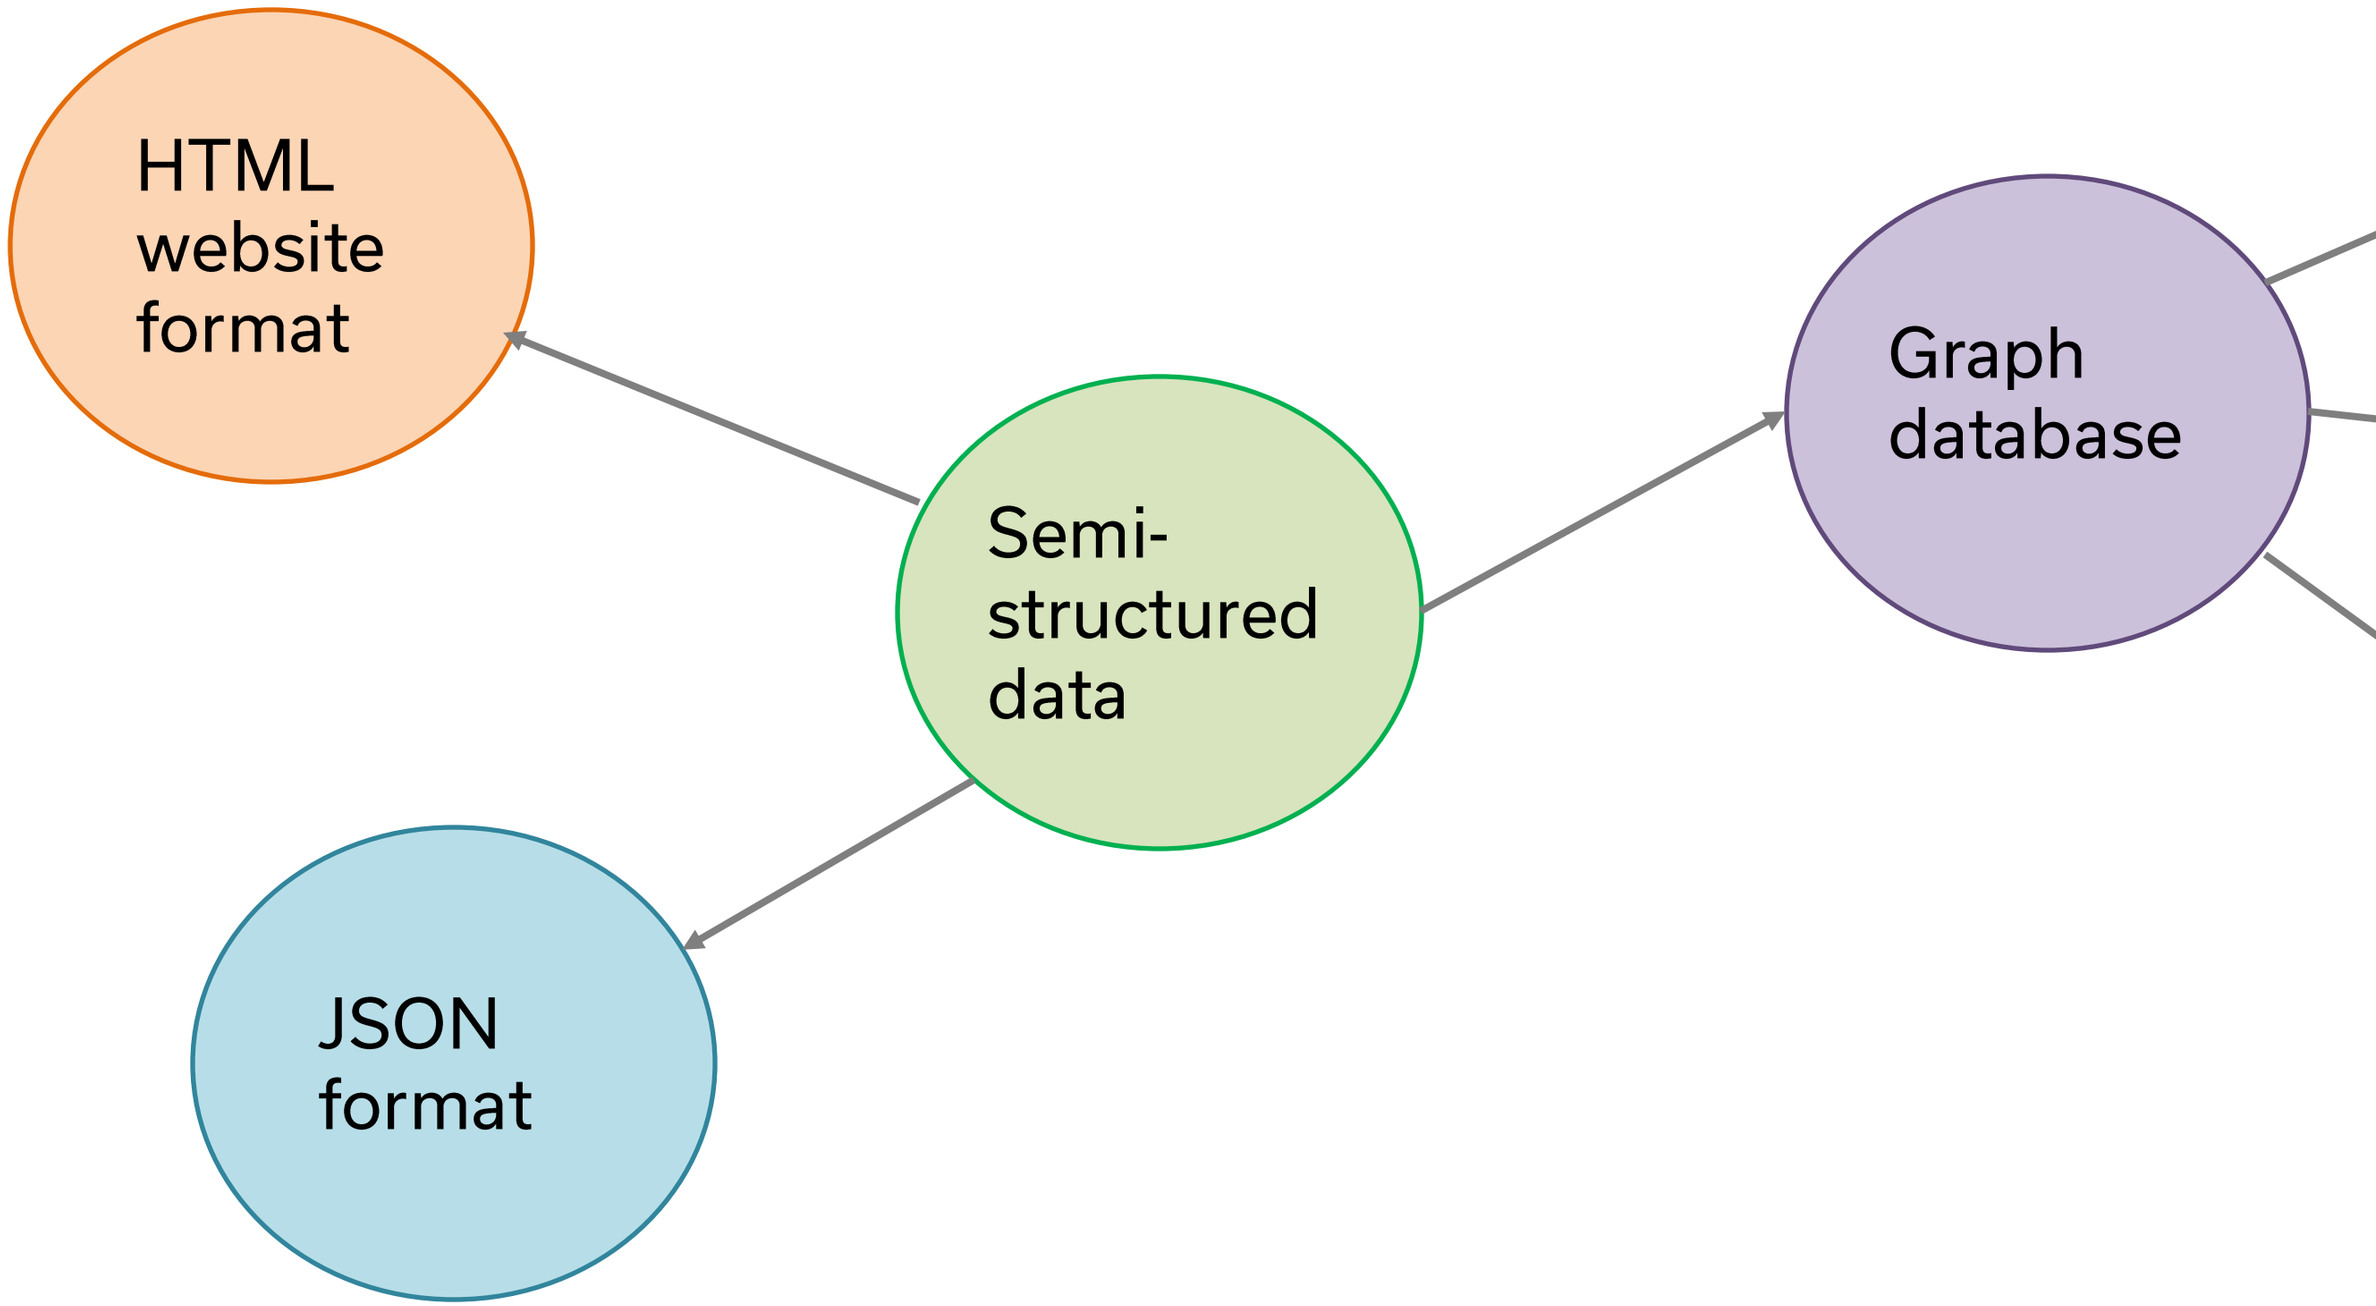
“definition”: “Data is organized in groups of ‘keys’ and ‘values’”,

“sub-definition”: “Values may follow no consistent structure, but they can always be indexed using the keys, which are reliable”,

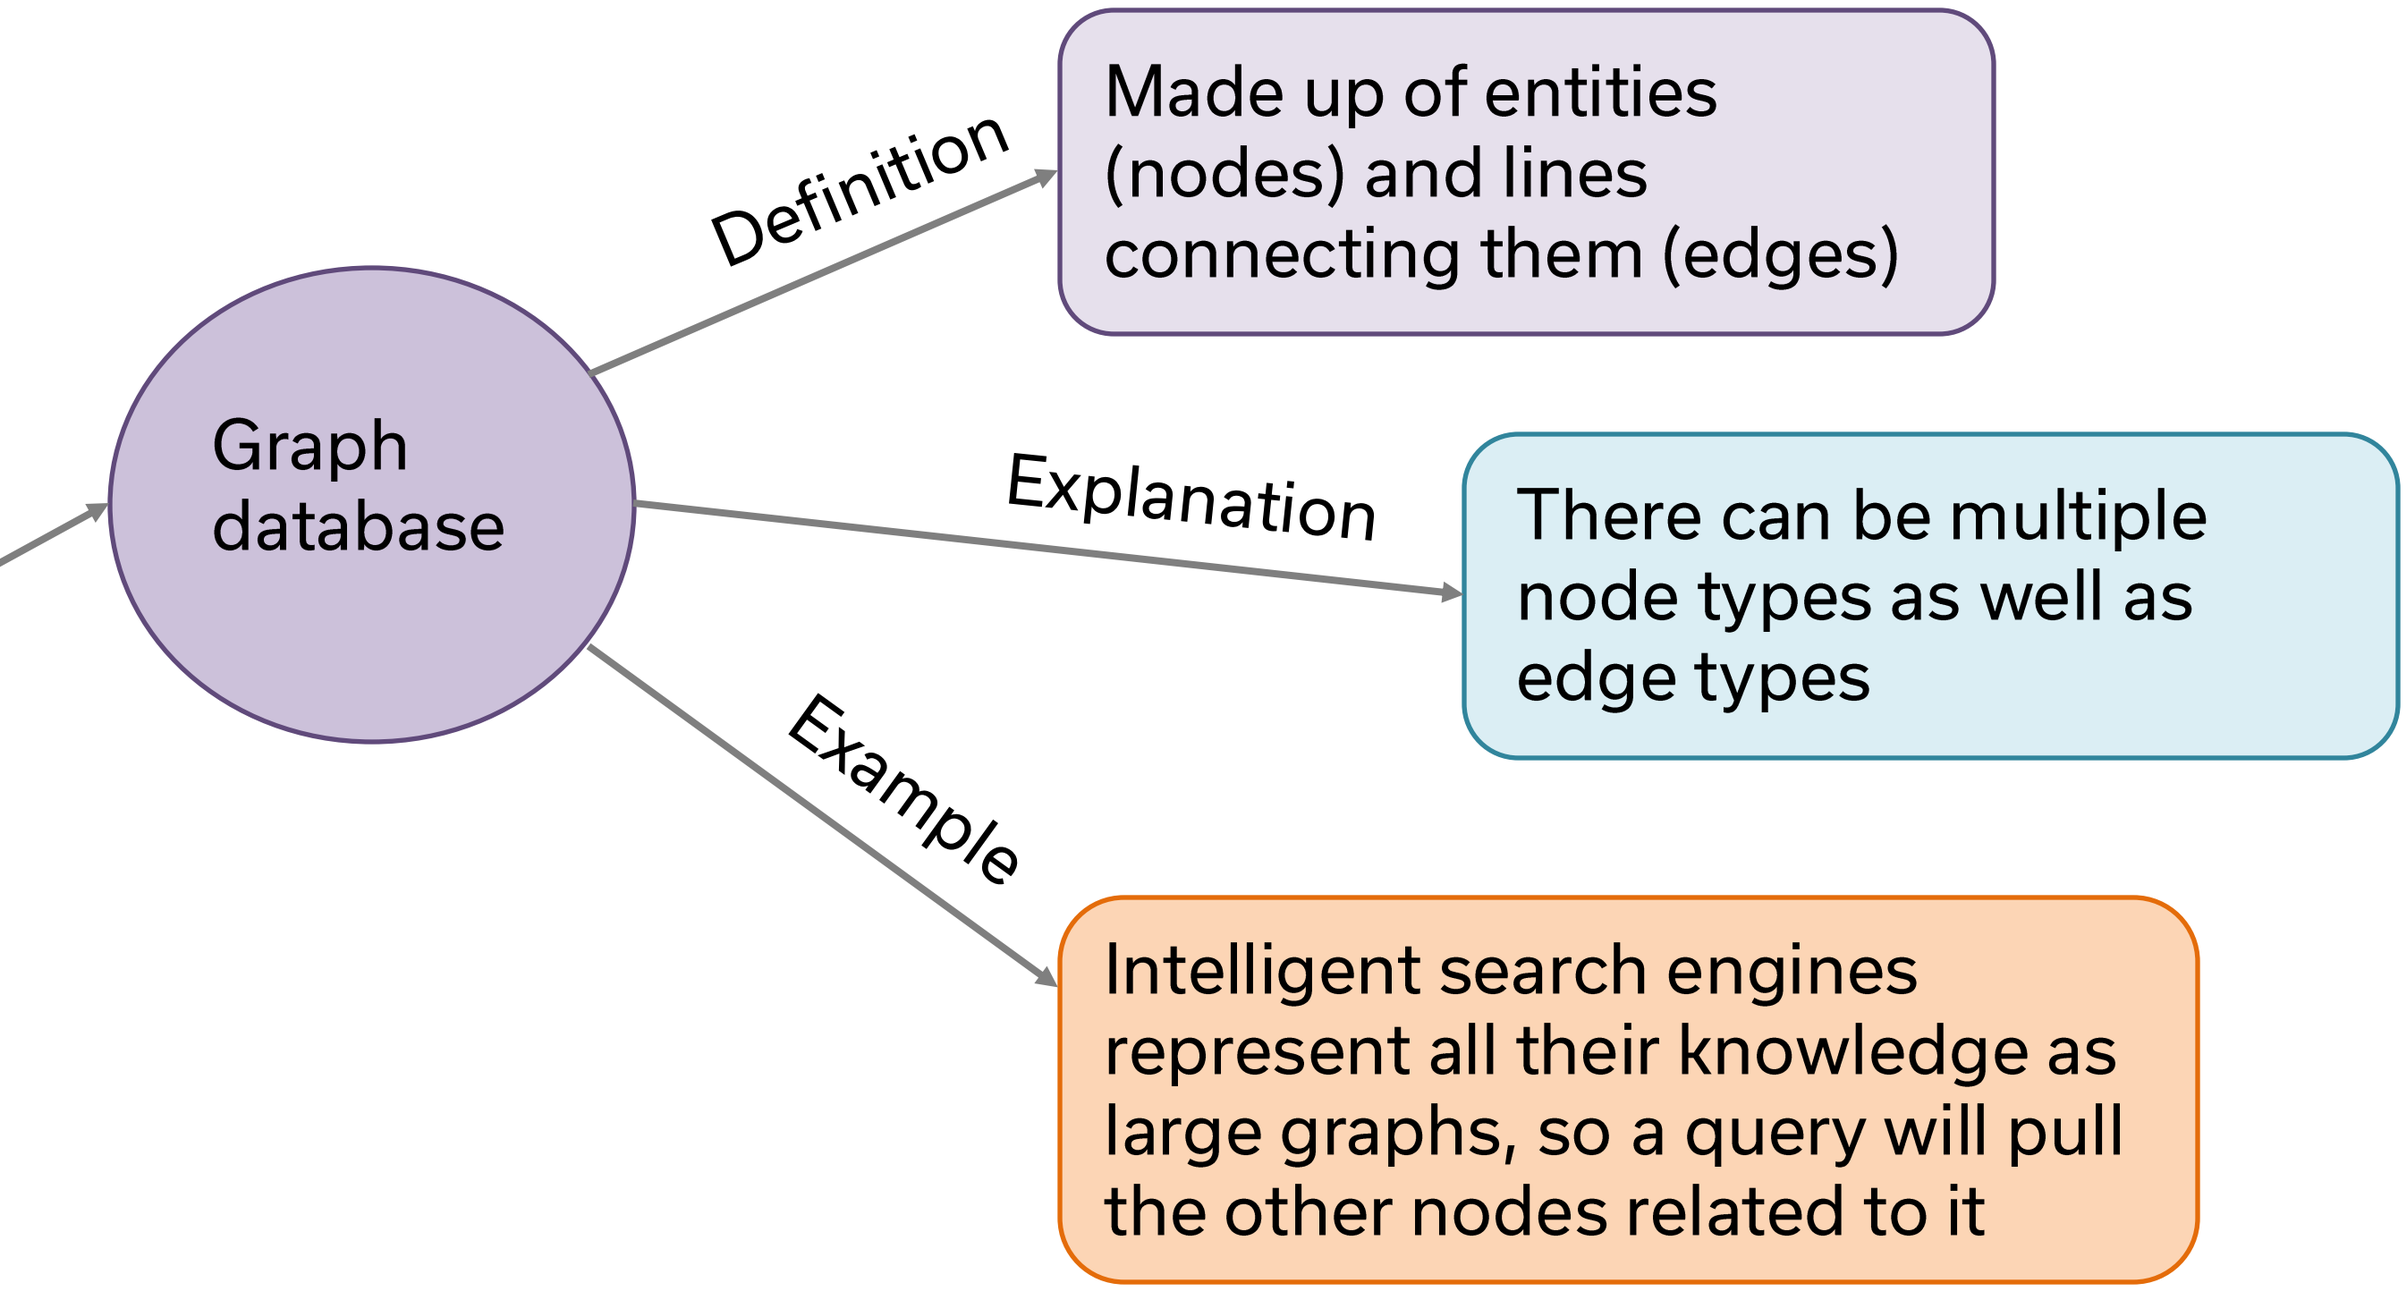
“example”: “Smart devices, such as thermostats, frequently send data in JSON format, so an AI system can use this to intelligently manage heating in a building”

}

Semi-Structured Data



Semi-Structured Data



Semi-Structured Data

- Pros:
  - Easier to produce than fully structured data
  - Rigid enough to still be effectively analyzed
  - Very common for certain use cases
- Cons:
  - Can still require pre-processing and data cleaning
  - Not always interpretable for humans
  - May require custom ML solutions

How is data stored?

- Lots of different database management systems (DBMS)

Structured	Unstructured	Semi-structured
MySQL	amazon S3	MongoDB.
ORACLE	Microsoft Azure Blob Storage	neo4j
PostgreSQL	Google Cloud	redis

NAIRR Data Resources

- NAIRR provides multiple resources for storing data:

Amazon Web Services (AWS)	• Cloud platform for data storage, computing, and model training/inference
DataBricks	• Cloud-based data lake for reporting and analytics
Google Cloud	• Access to datasets and computing
Microsoft Azure	• Data storage and model services

Data Concerns

Bias and fairness

Quality (accurate labels, complete records)

Privacy and security

Data Concerns

Privacy and security

- Data anonymization is important in many fields
- Multiple techniques:
  - –Data masking
  - –Suppression to reduce details/group values
  - –Replacing identifiers with codes
  - –Synthetic data generation
    - **NAIRR Resource**: Lexset Seahaven

Conclusion

- AI is as diverse as the data required to create it
- Good data needs both quantity and quality
- Watch for bias, confounding factors, scale
- Tailor your use case to your available data
- Data can take many different forms and machine learning techniques exist for just about all of them
  - Think through how your data can be best represented

AI, the Workforce, and Higher Education:

Implications and Opportunities for Kentucky

Travis Muncie, CPE

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

Agenda

- **AI and the Workforce: Implications for Kentucky**
  - Early signals of broad disruption
  - How can national patterns inform our work in Kentucky?
  - What if Kentucky isn’t prepared?
  - How do we get prepared? AI can actually help!

- **AI and Postsecondary Education: Student Success Opportunities**
  - Sustaining success in Kentucky
  - Informing the public

AI and the Workforce:
Implications for Kentucky

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Early Signals of Broad Disruption

AI performs ~16% of U.S. labor tasks today

- AI writes over **a billion lines of code per day**
- AI reads over **150,000 words per second**

Widespread layoffs and hiring freezes

- Major companies are laying off **thousands in IT management and HR**
- **Hundreds of thousands of engineers** were laid off this year with nationwide hiring freezes
- Some data corporations estimate **50% of their customer service and sales work** in 2025 was done by AI

National Insights -> Kentucky

By 2030, thirty percent of the workforce will be unmeasurable

Manufacturing:

There will be a **50%+ reduction in production jobs** by 2030

Customer Service:

Over **60% of customer service** transactions are automated

Transportation and Logistics:

Autonomous vehicles logged **71 million miles** last year

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

What if Kentucky isn't prepared?

Core industries are at risk

- **Transport/Logistics and Manufacturing** account for nearly **¾ million jobs.**
  - Experts estimate that KY could face a **10% reduction in median income** and a **double-digit spike in unemployment**
  - A **quarter million Kentucky jobs** and **$11 billion of revenue** at imminent risk
- Counties along the **I-64 and I-75 corridors** are at severe risk

States lack visibility into AI-driven labor shifts

- For most, it seems like trying to study and understand water while holding it in your hand

How do we get prepared?

State and local gov't must investigate current impact

- **Create AI labor indices** to quantify and understand the problem – specifically measuring **AI job loss** – at a local, granular level.
- **Measure and understand industries at risk** in Kentucky's workforce
  - Map **skills and upskill potential** to other industries
- **Risk Modeling**
  - Test, simulate, predict, and retest – **AI can actually help us prepare!**

How do we get prepared?

Postsecondary education must adapt for future impact

- Are there opportunities to **reprioritize academic programs?**
- Advocacy for and adaptability of **state support around critical programs less threatened by AI**
  - **Healthcare**: HWIF
  - **Skilled Trades:** Kentucky Work Ready Scholarship (KWRS)
  - **Education:** Teacher Program Scholarships
  - **Early Postsecondary Ops:** Dual Credit and WRDC Scholarships

AI and Postsecondary Education: Student Success Opportunities

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

How do we get prepared?

Sustaining student success improvement

- How can we better understand **how key factors (affordability, transitions, etc.) impact student success?**

Increasing public access to PSE data

- How can we **help the public and lawmakers quickly answer simple but important data questions** about higher education?

Student Success

Example: Using ML and AI to target financial aid

- We modeled how out-of-pocket costs (and at what thresholds) impact second-year retention
- Now we're using ML and AI tools to **turn inferential research into predictive power**
- **Providing information back to institutions** to help target institutional aid and resources

Student Success

Implementation on a grand scale?

- **How can we integrate AI tools with the best longitudinal postsecondary data system in the United States?**
  - Estimate massive **exploratory inferential models**: What could we be investigating that's not our radar?
  - **Predict, simulate, and constantly inform** student success interventions
  - How can we do this in a controlled server with **minimal security risks?**

Informing the Public

How can we use AI to increase public awareness?

- **Example:** How have retention rates at Kentucky public institutions changed over the last decade?

- Create interactive AI chat to **rapidly answer simple, important questions**

- We're investigating creating secure servers with aggregated data tables, **creating zero security risk opportunities**

References

- Chopra 2025. “Project Iceberg: Workforce Security for Kentucky in AI Age.”
- Kentucky Cabinet for Economic Development 2026. “Top Industries.”
- Kentucky Center for Statistics (KYSTATS) 2026. “Employment and Wages by Industry.”
- Kentucky Workforce Innovation Board 2026. “Kentucky’s Top Five In-Demand Sectors.”

Machine Learning

Aaron Mullen, MS

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

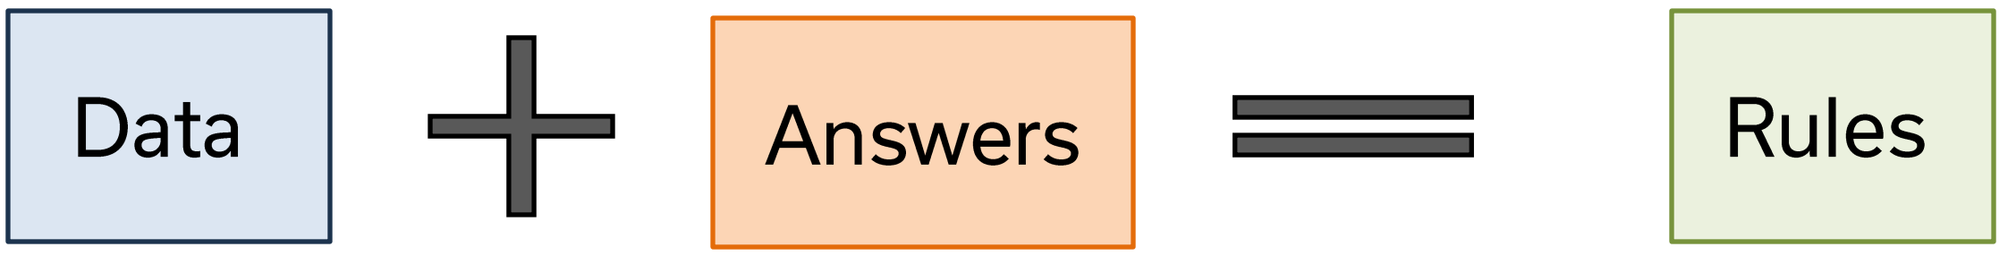
What is Machine Learning?

- Getting computers to perform a task without being explicitly programmed by humans
- The computer learns from data in its own way

Traditional Learning



Machine Learning



Machine Learning Example



Machine Learning Example

- Email providers use classification algorithms to analyze subject line, sender information, and specific phrases in emails
- Models can be trained with many examples of emails, labeled "Spam" or "Safe" so the machine can learn the difference
- It learns that certain phrases ("Congratulations! You won…"), senders, times, or images are associated with spam emails more frequently



Machine Learning Example

Why Machine Learning?



- The first spam filters were rule based and checked for specific keywords:
  - IF email contains "Click Here" THEN mark as Spam
- Scammers could change spelling/punctuation to avoid these filters
  - "Cl1ck Here" or "C!ick He re"
- ML checks for intent and structure rather than specific keywords

Categories of Machine Learning

- Supervised Learning
- Unsupervised Learning
- Reinforcement Learning

Supervised Learning

- Model learns from data that is already labeled with correct answers
  - Spam detection is example: emails used to train models are already labeled “Spam” or “Safe” so the model can learn the difference
- Multiple types:
  - Classification (malignant vs benign tumor)
  - Regression/forecasting (how many sales made in the next quarter)
  - Segmentation (which part of this image contains a dog)
- But how does the machine actually “learn” from these examples?

How does a machine learn?

- Most models iterate through epochs, learning from the data at each step
- Dataset is split into different groups:

Training set
- ~70% of your data
- Used to train the model

Validation set
- ~15% of your data
- Used in training to ensure the model is on the right track

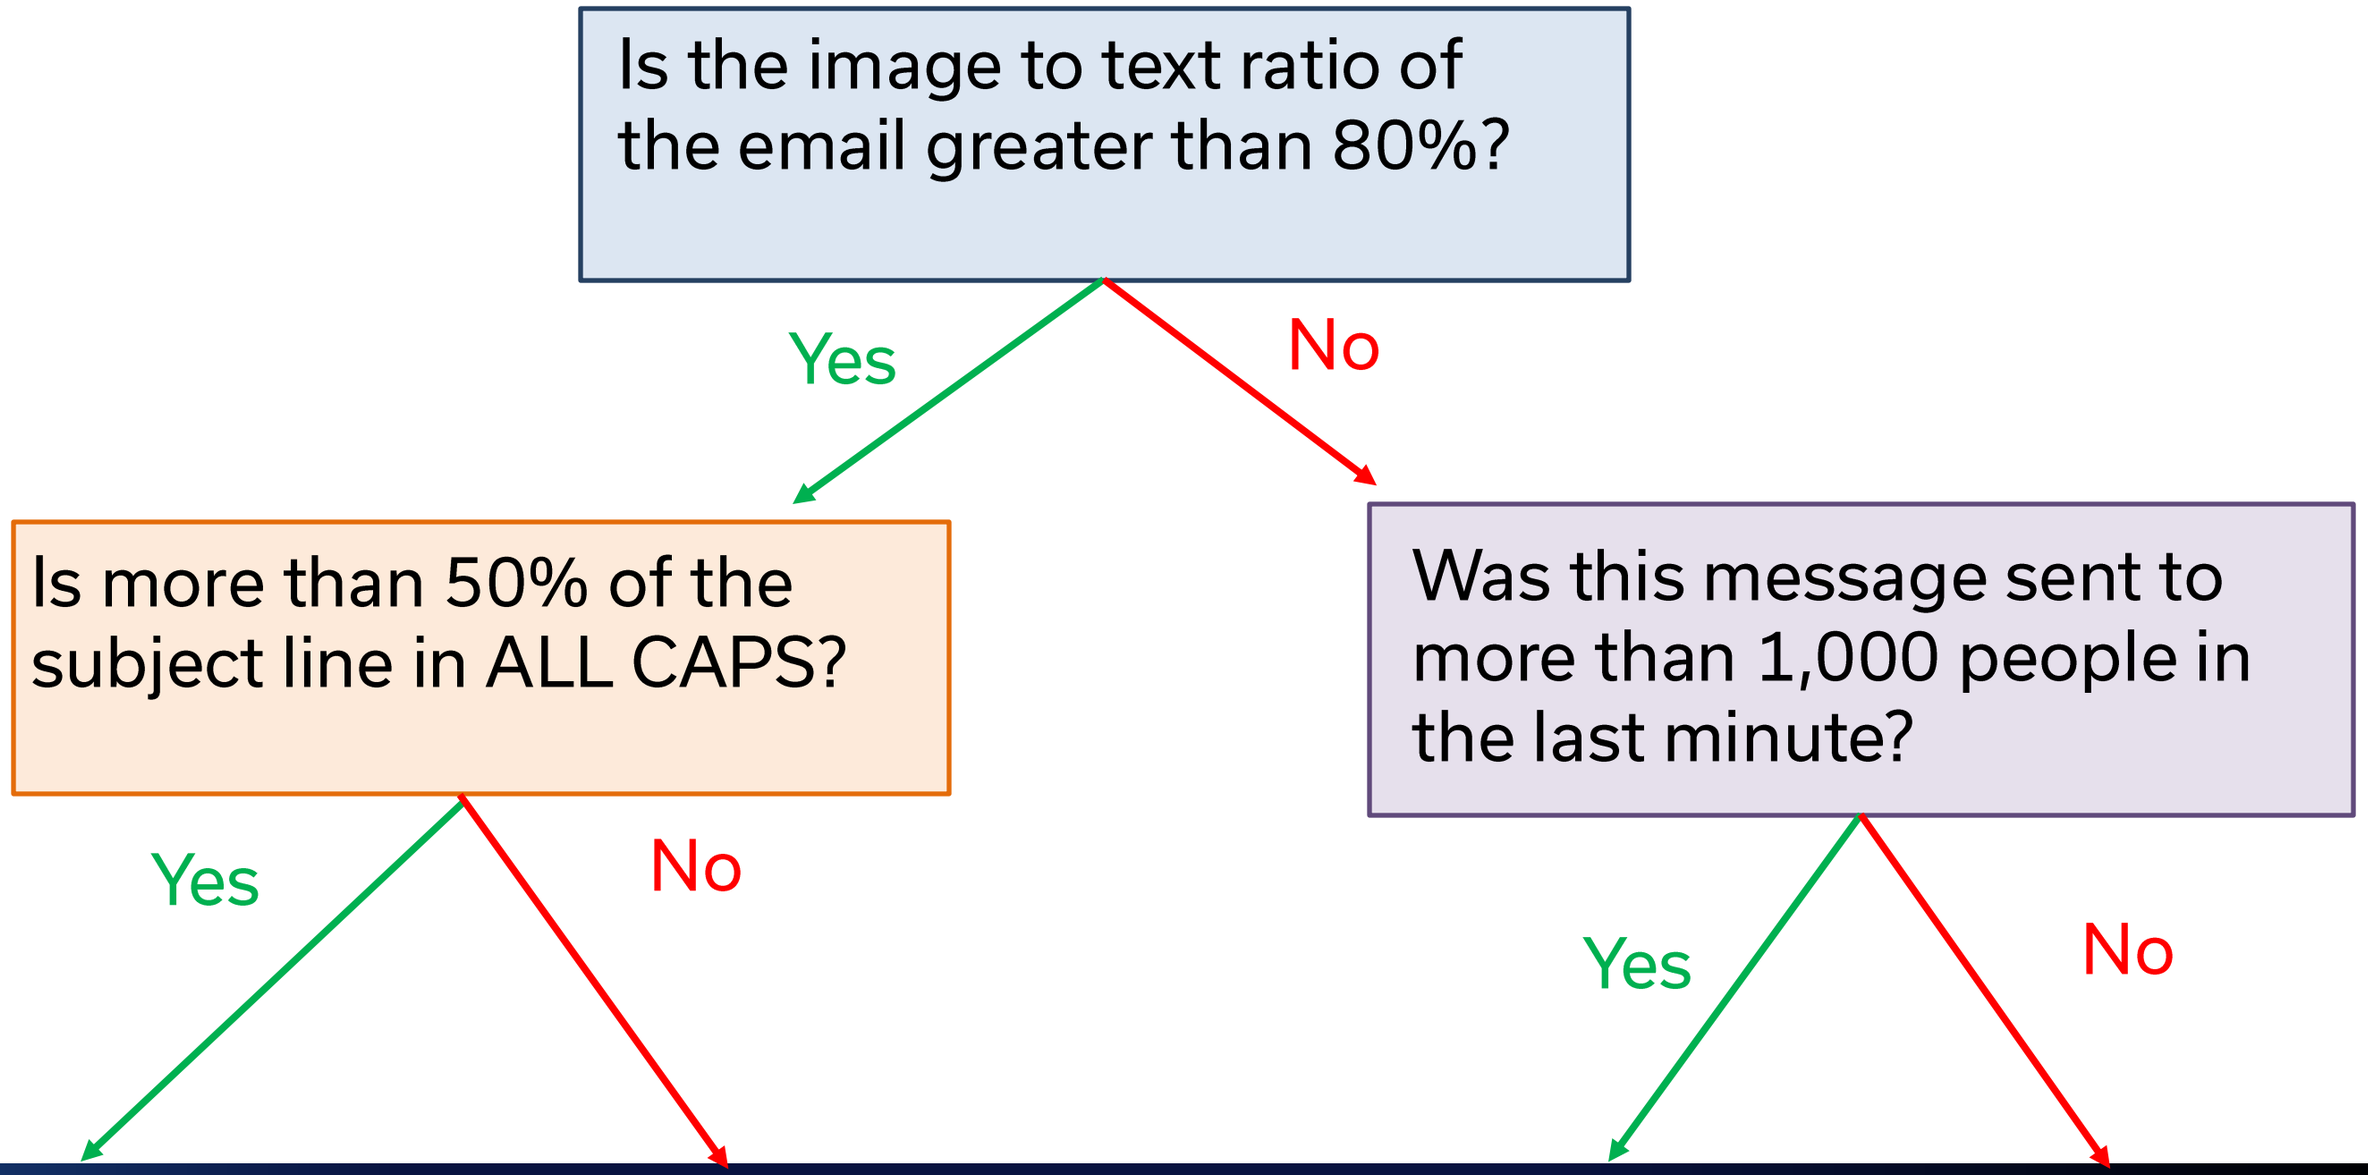
Test set
- ~15% of your data
- Completely unseen during training and used at the end to determine applicability of the model

How does a model learn features?

- During the training step, how does the model determine what features it should use to distinguish classes or make predictions?
- It depends on the type of machine learning model!

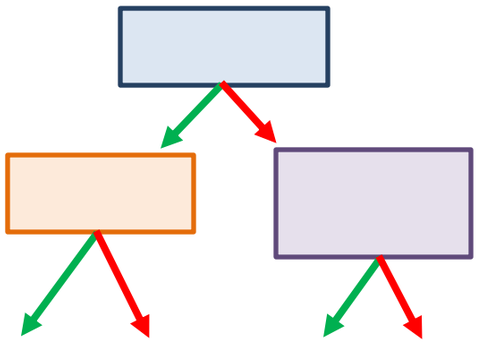
Decision Tree

- The model essentially builds a flowchart that maximizes accuracy when followed for classification



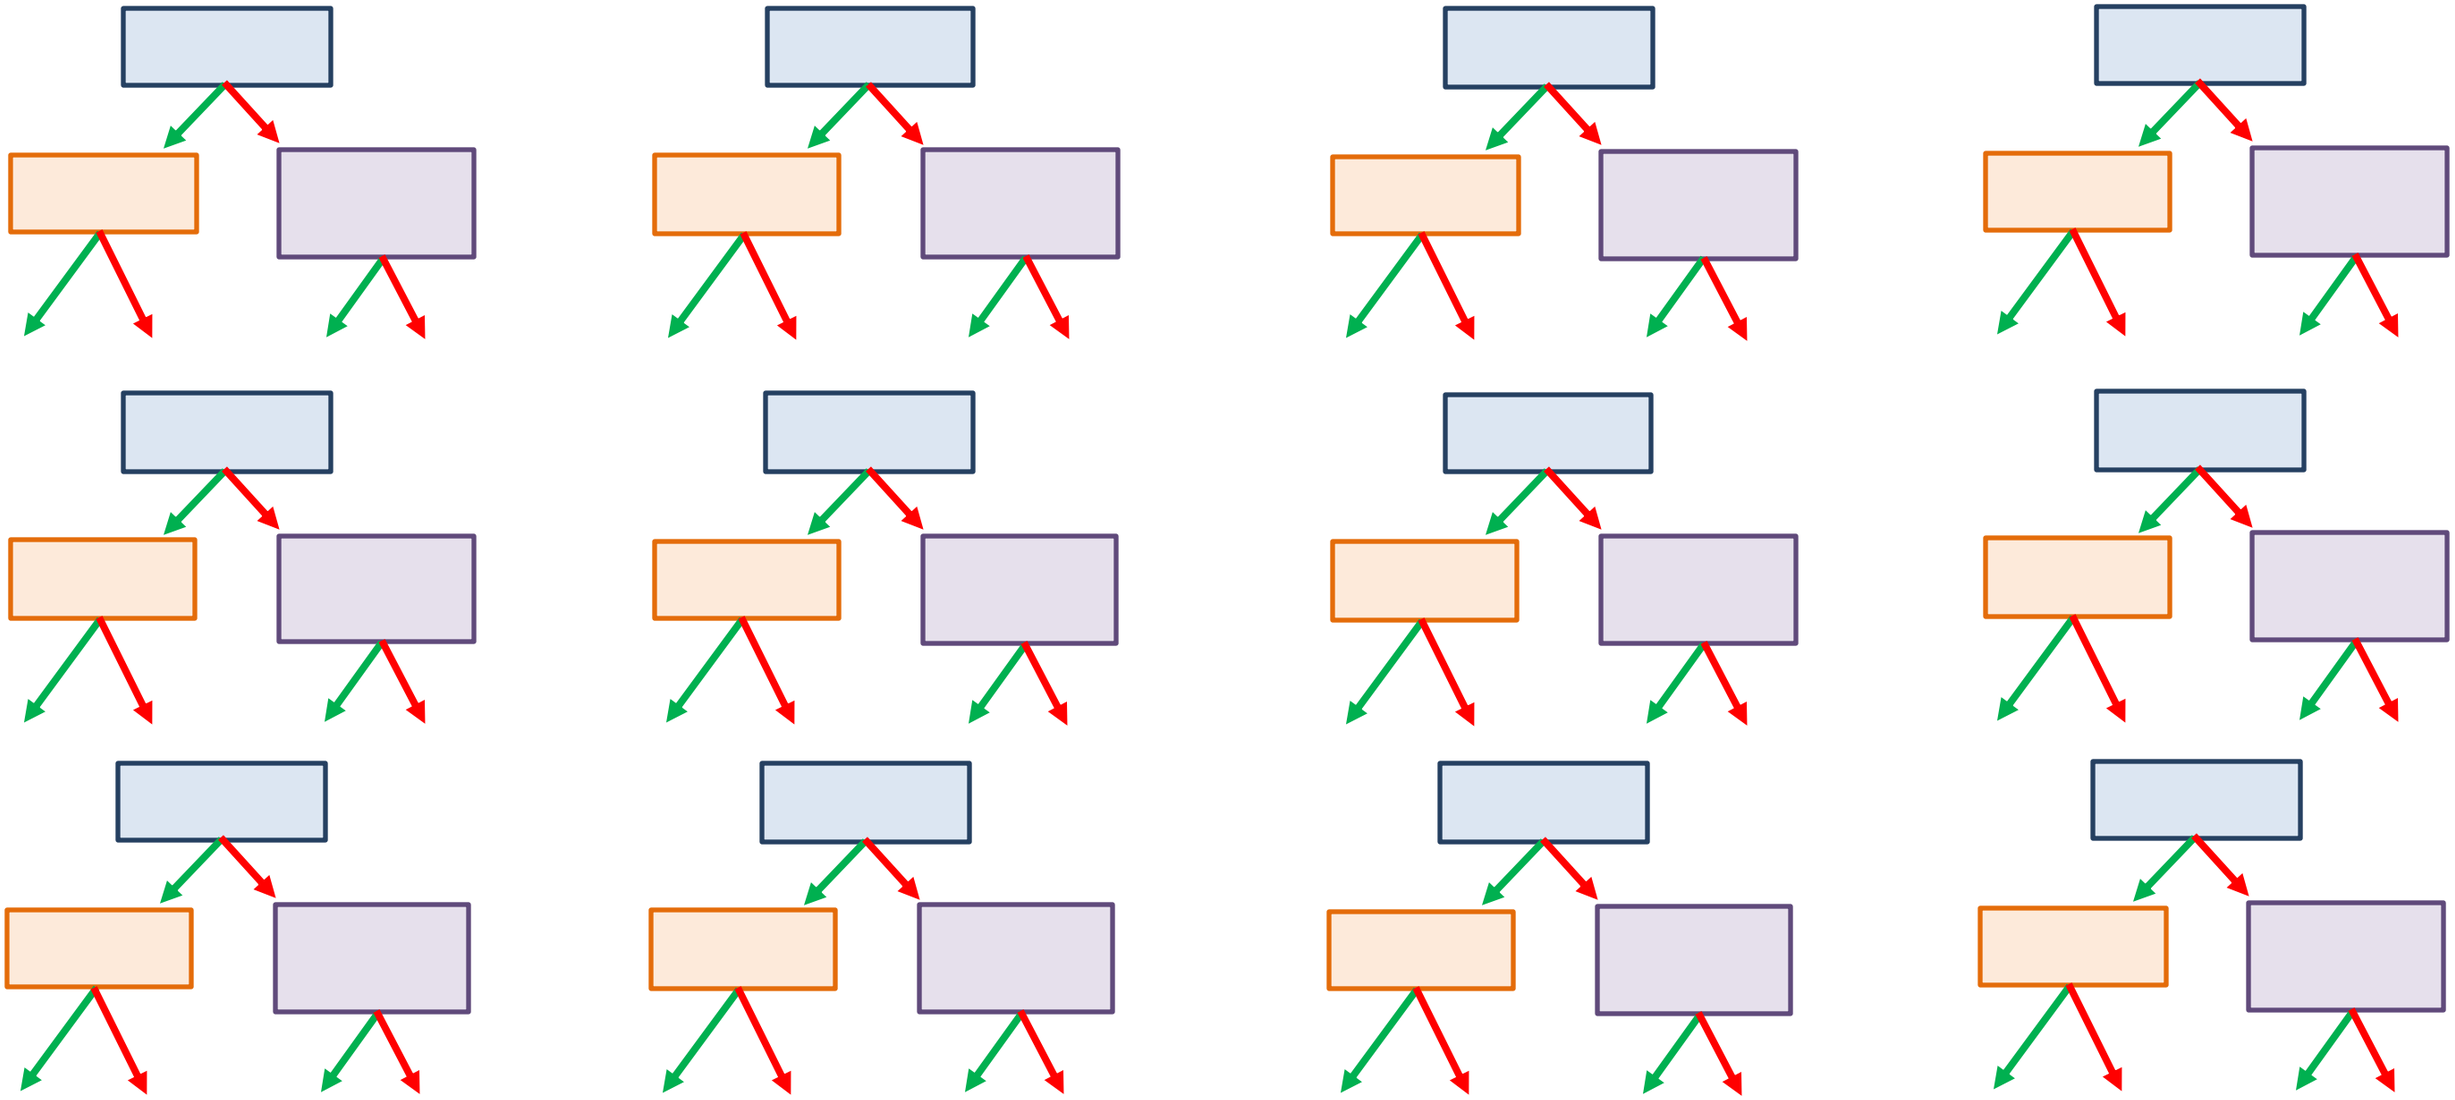
Decision Tree

- The model essentially builds a flowchart that maximizes accuracy when followed for classification

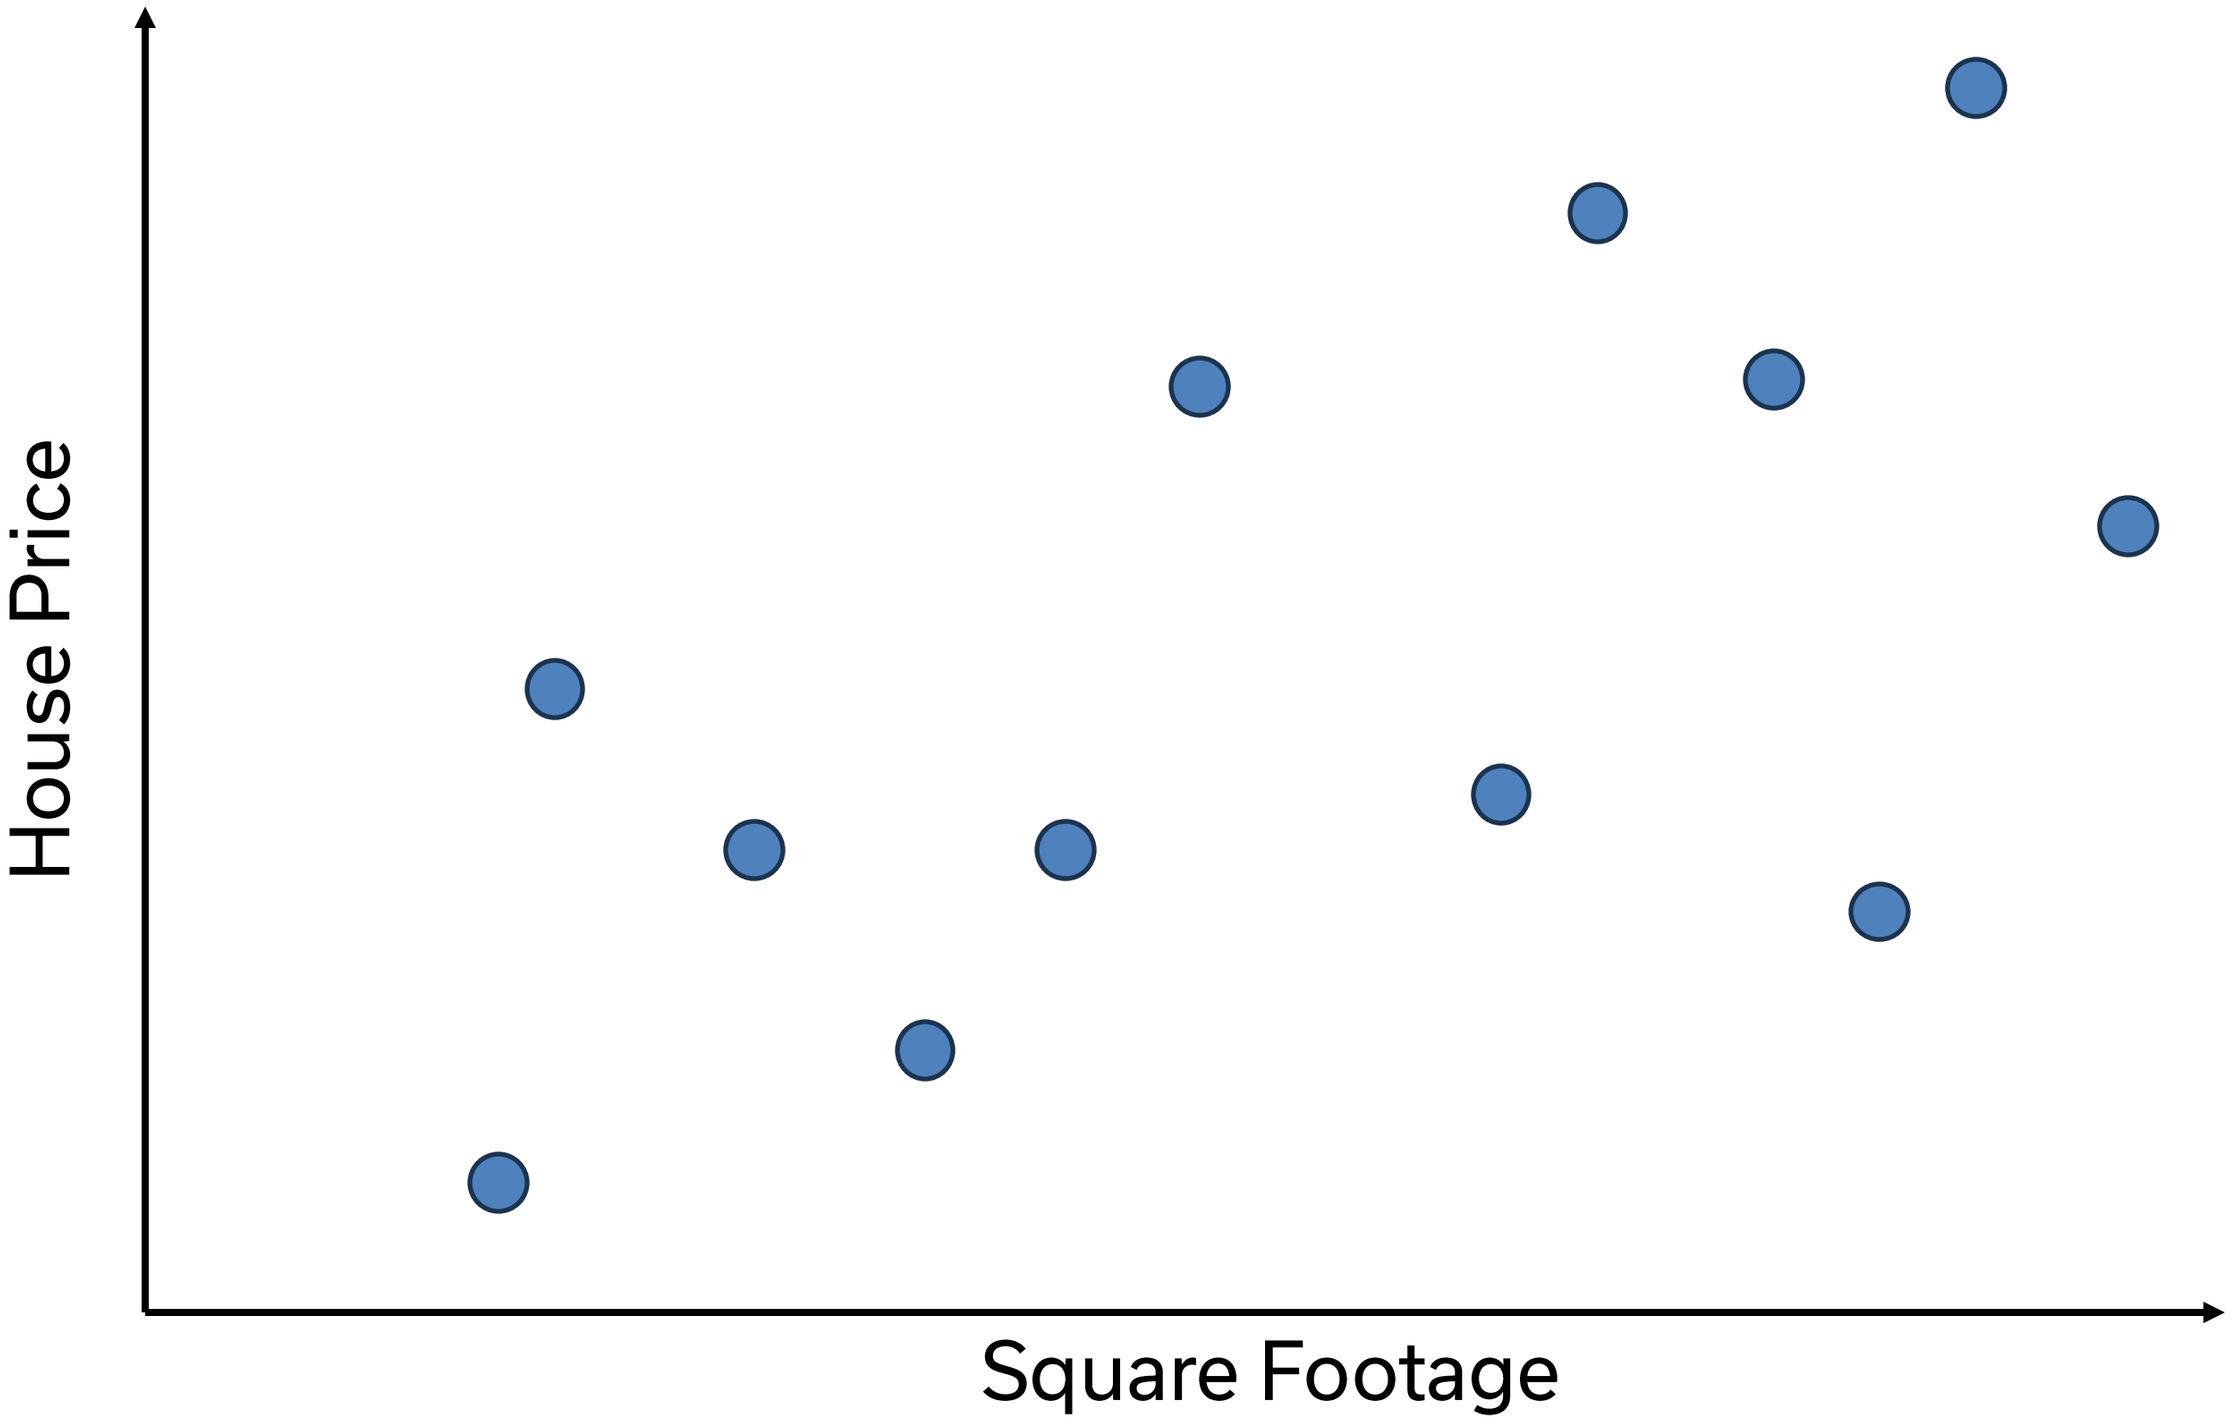


Random Forest

- Building many decision trees based off different subsets of data and features improves the model significantly

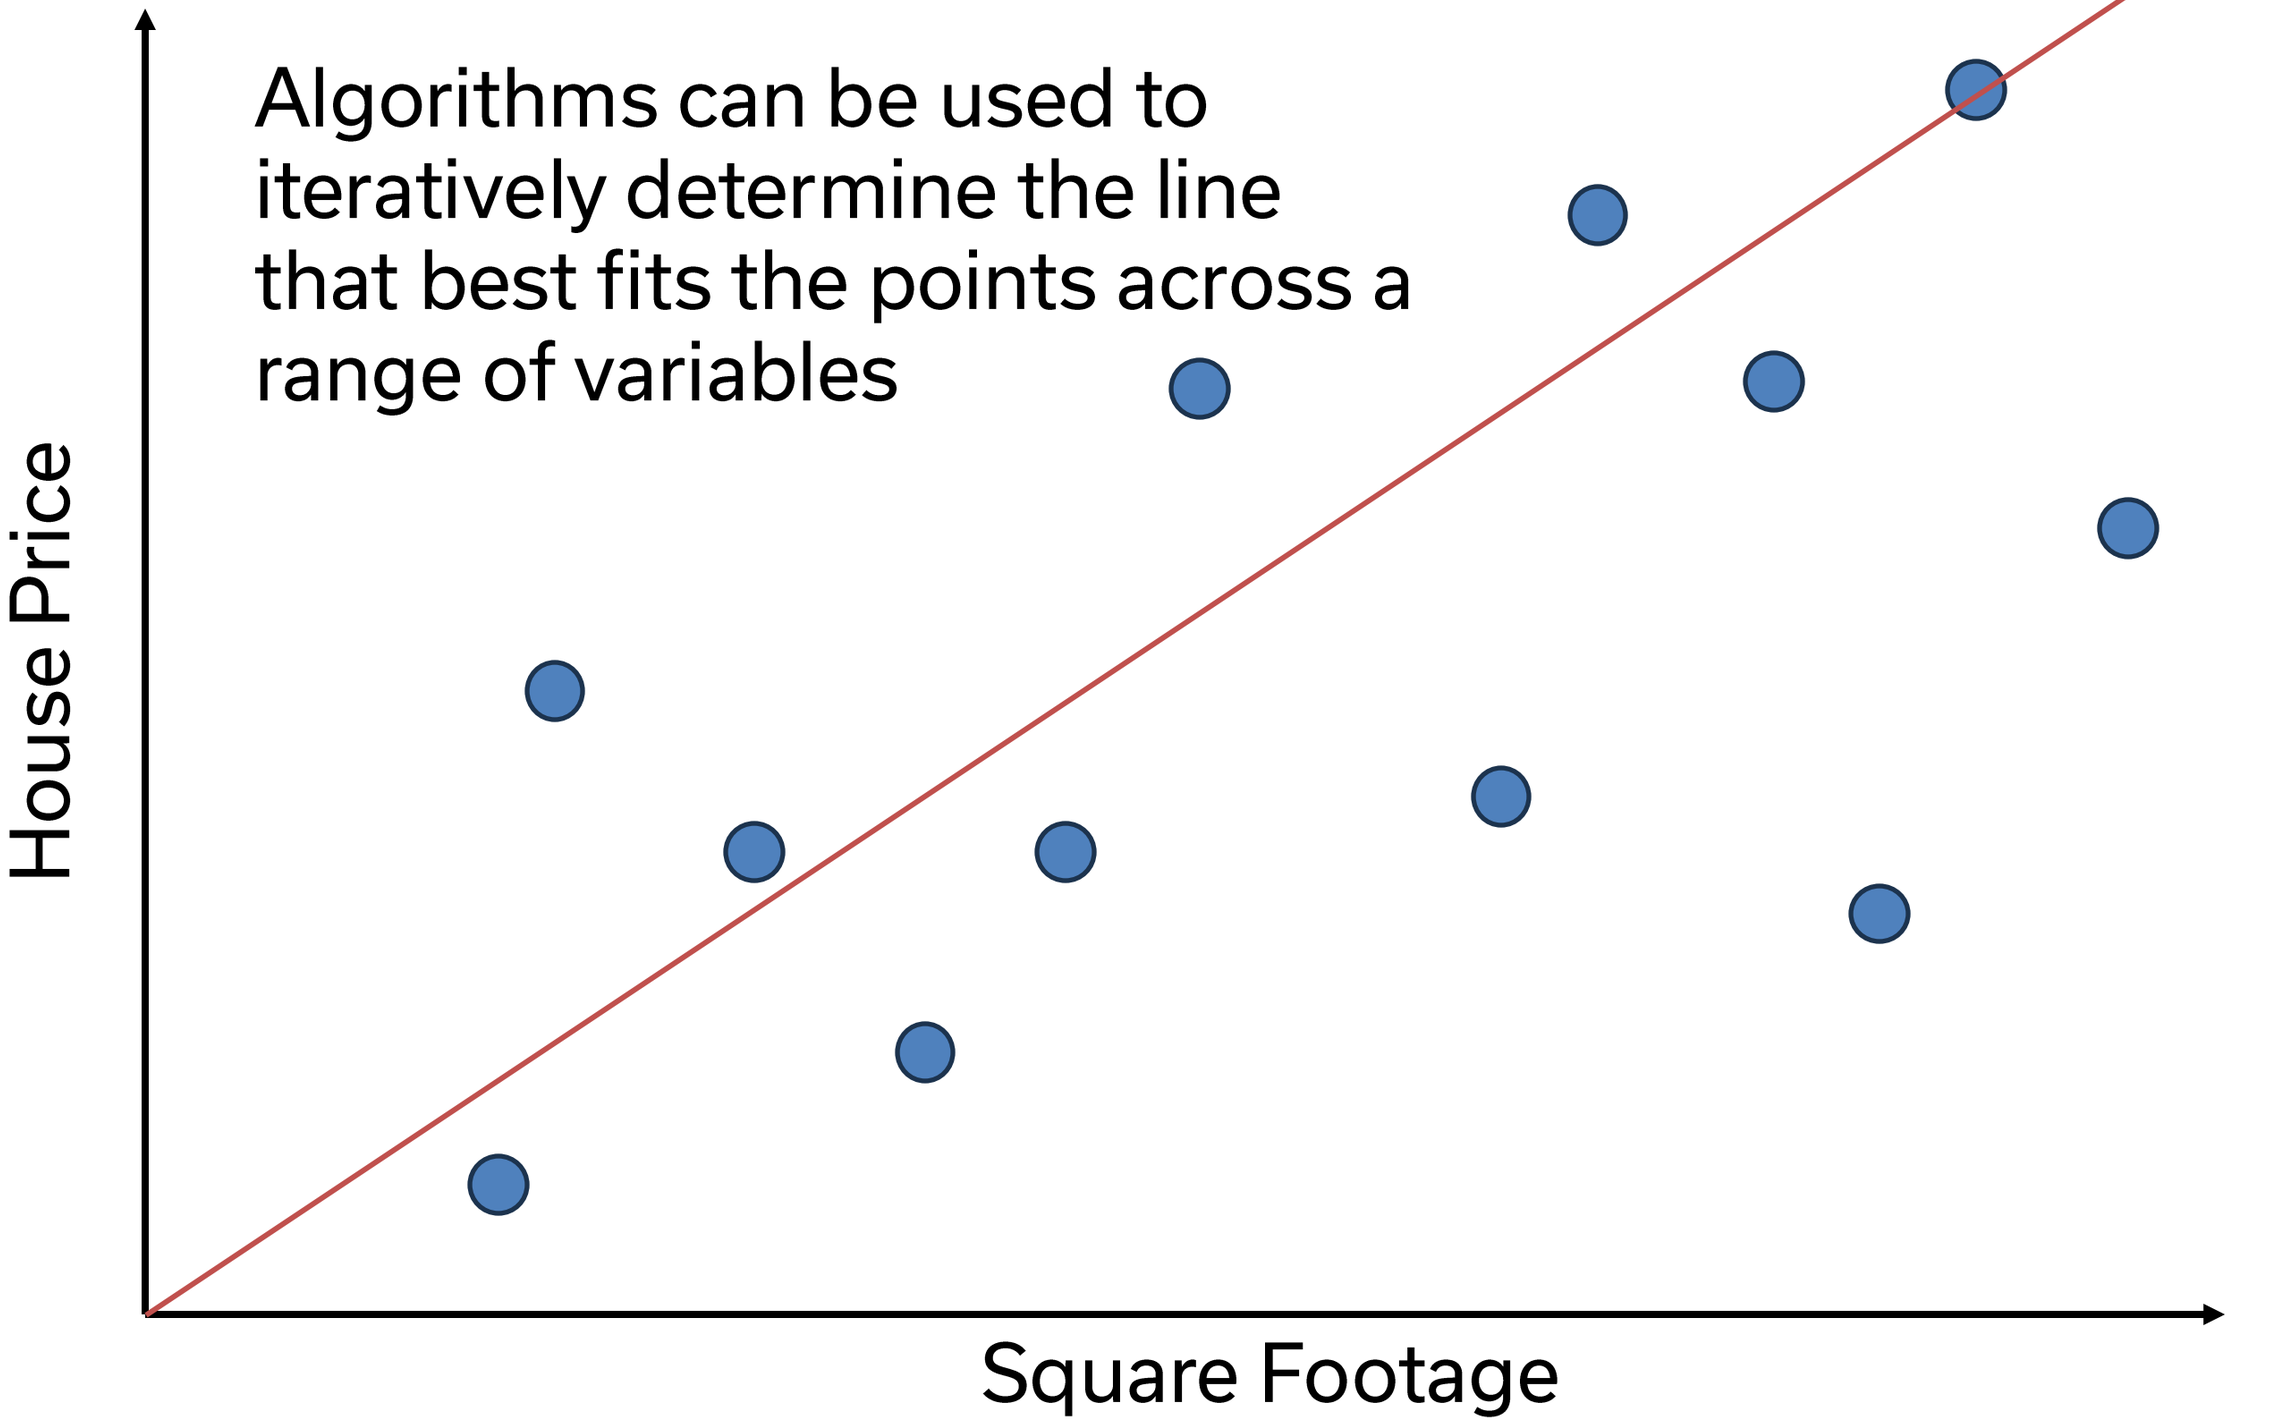


Linear Regression

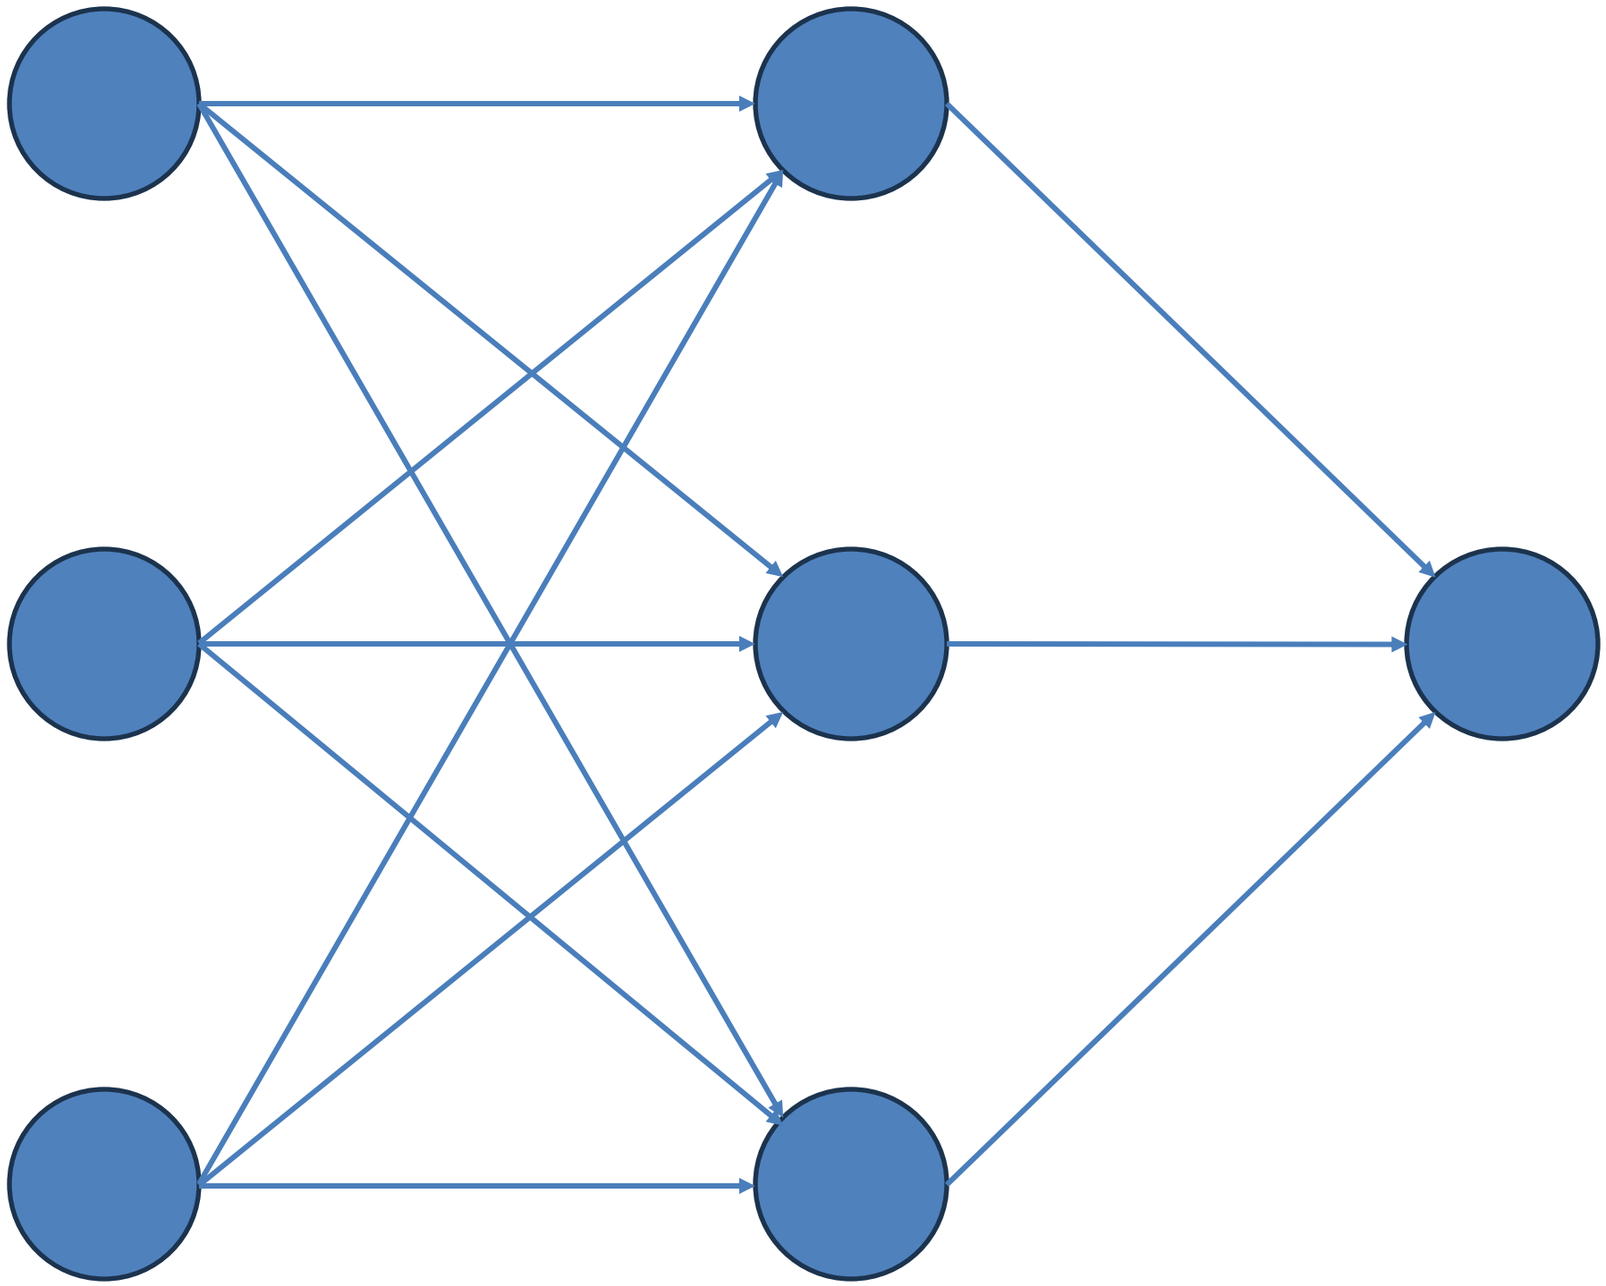


Linear Regression

Algorithms can be used to iteratively determine the line that best fits the points across a range of variables

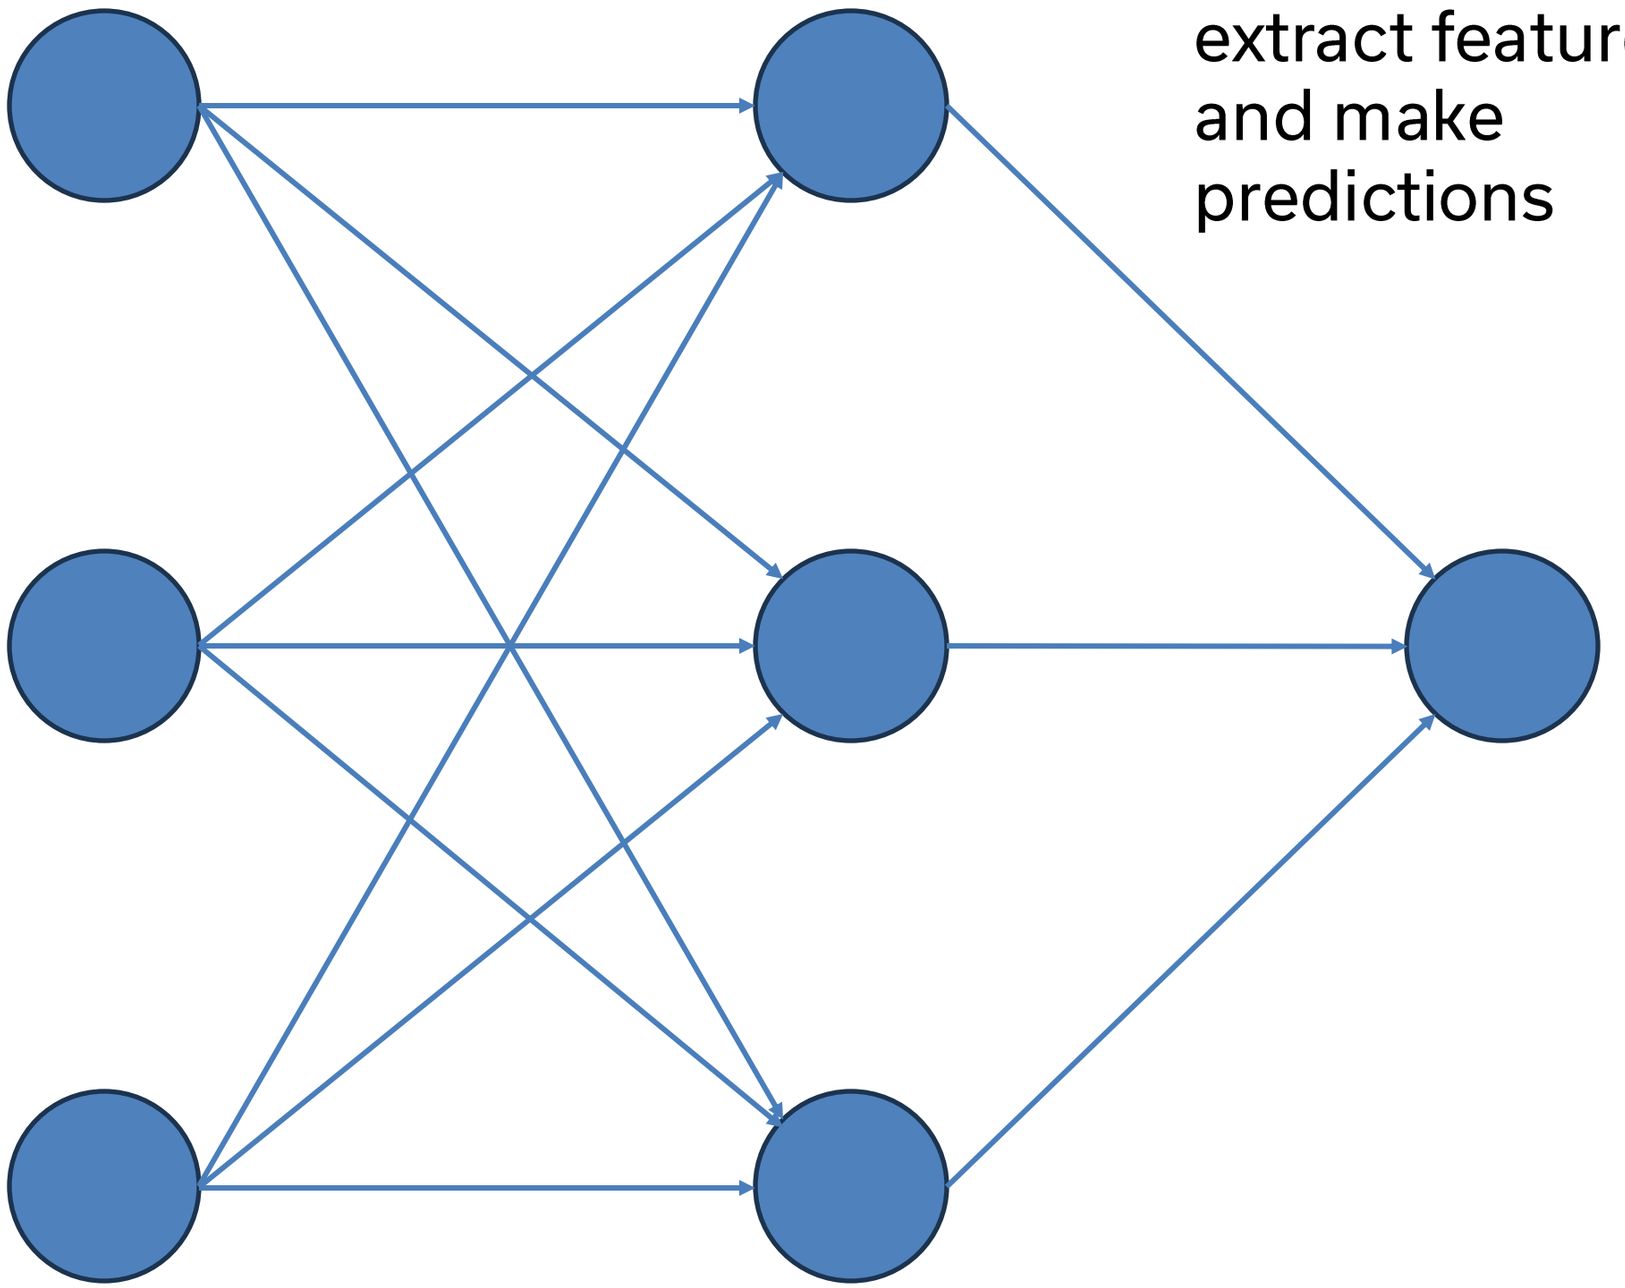


Neural Network



Neural Network

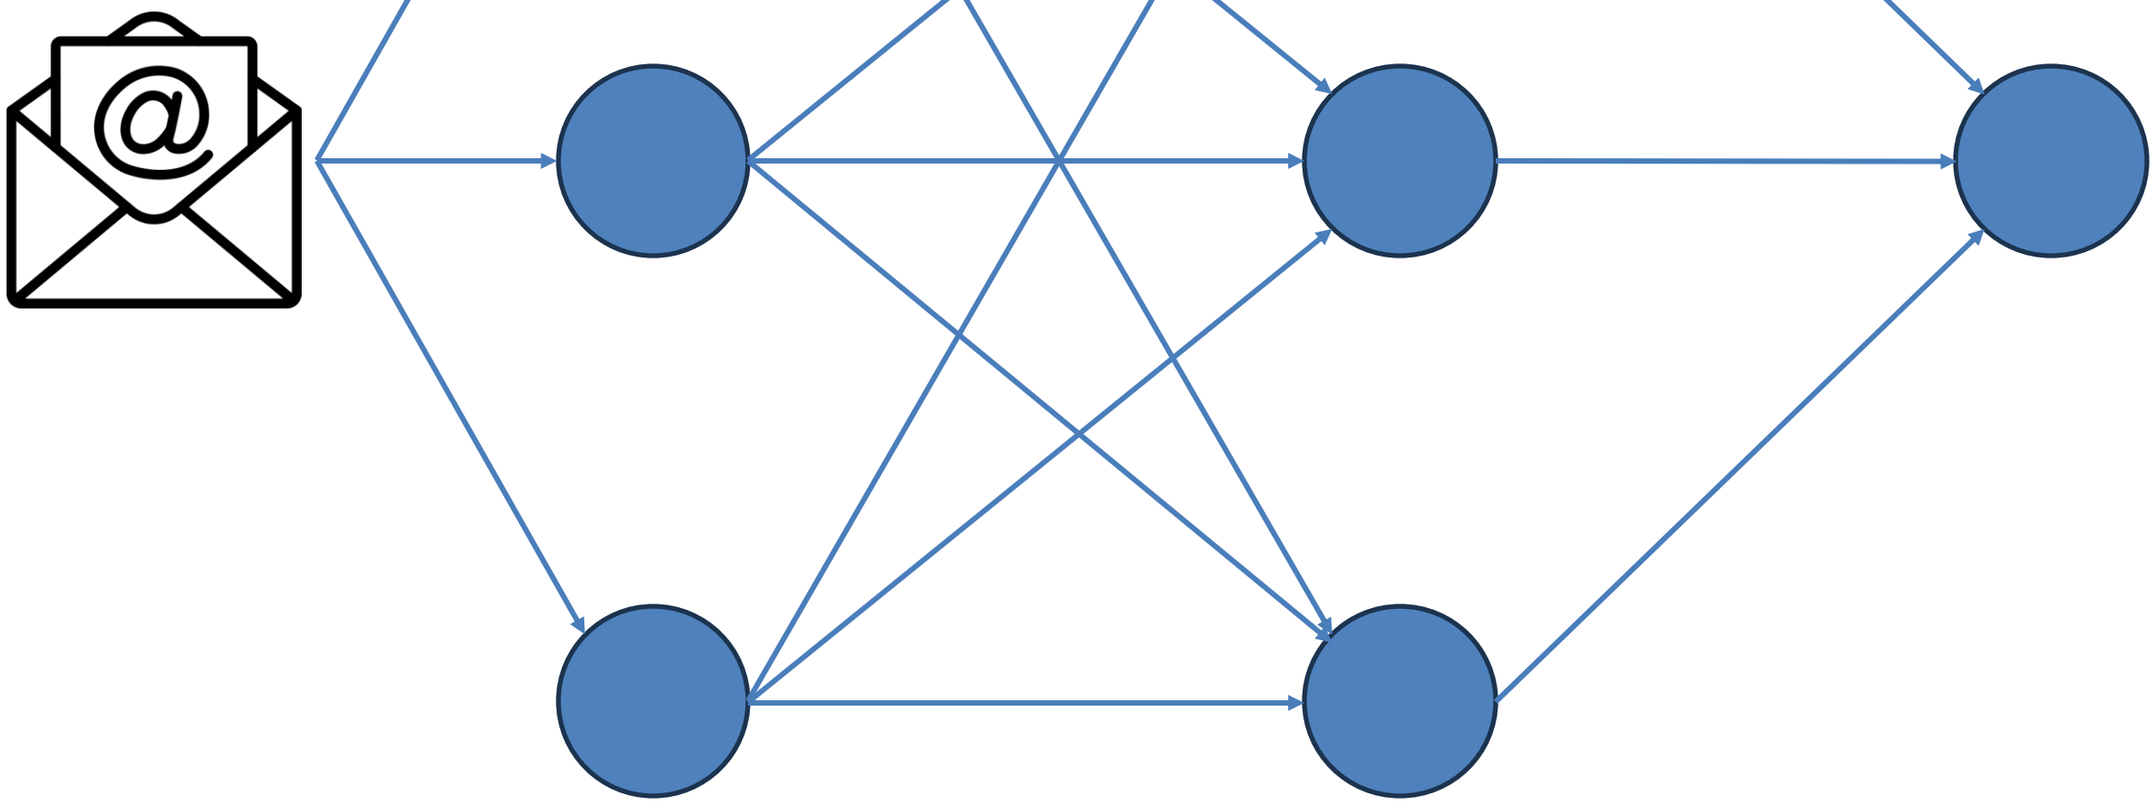




Each node of the network is iteratively tuned to extract features and make predictions

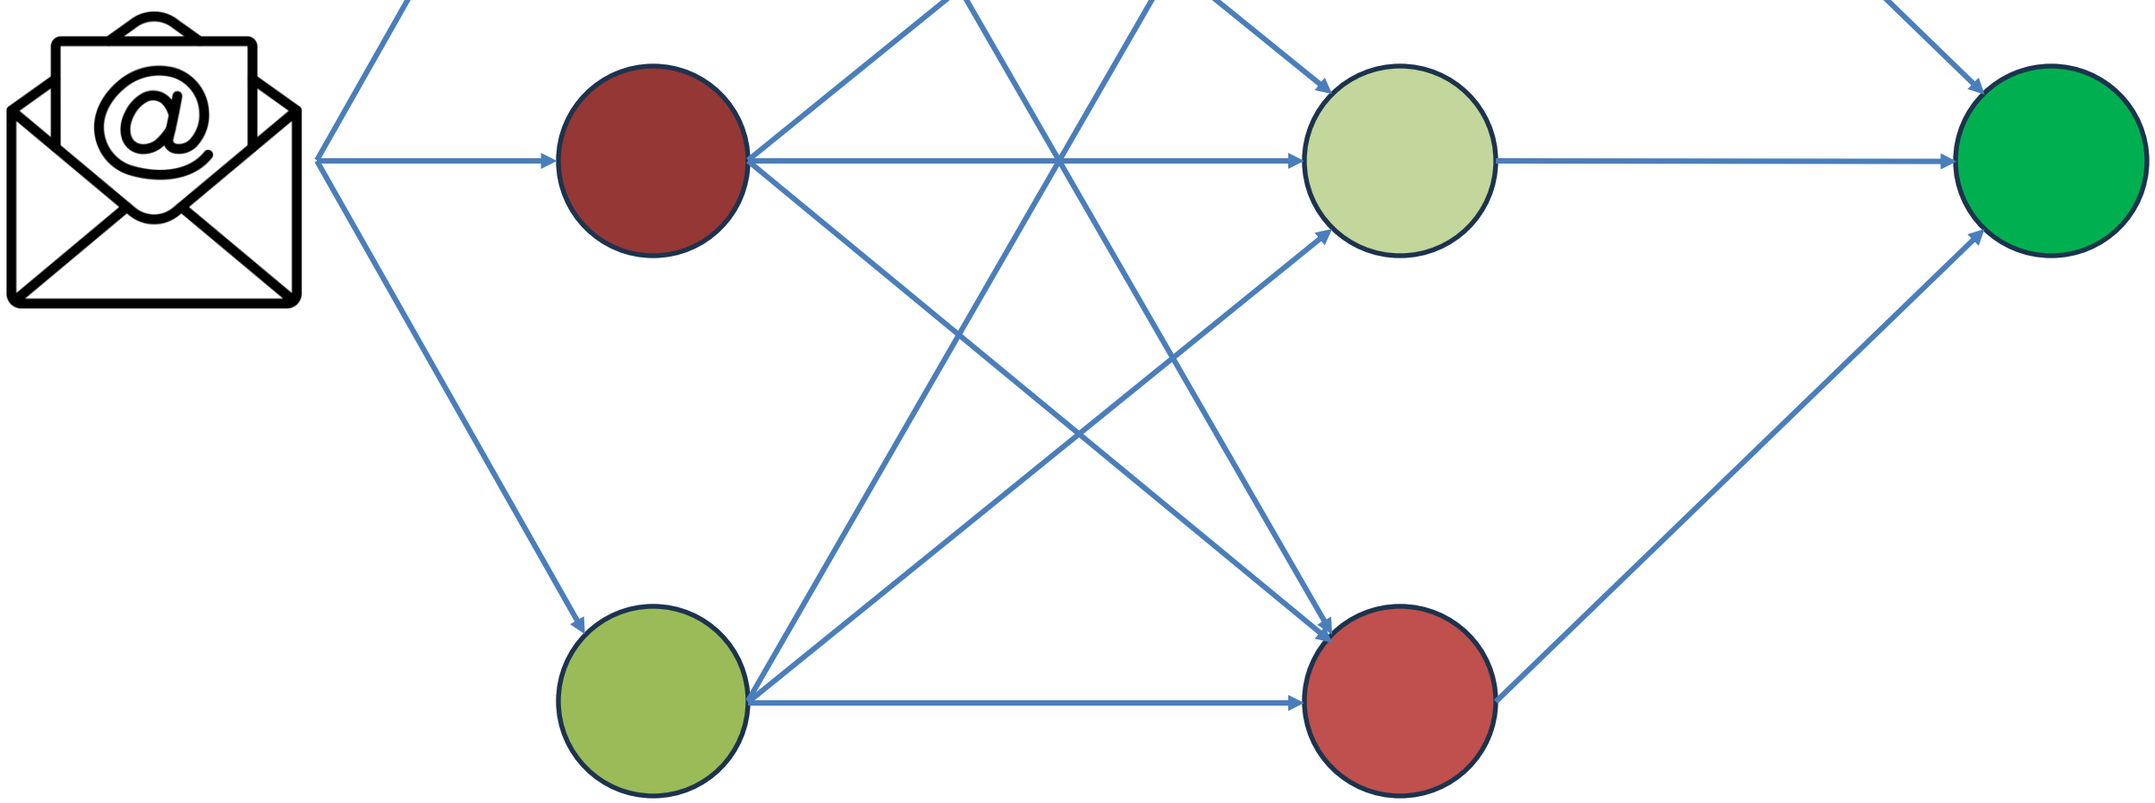
Neural Network

Each node of the network is iteratively tuned to extract features and make predictions



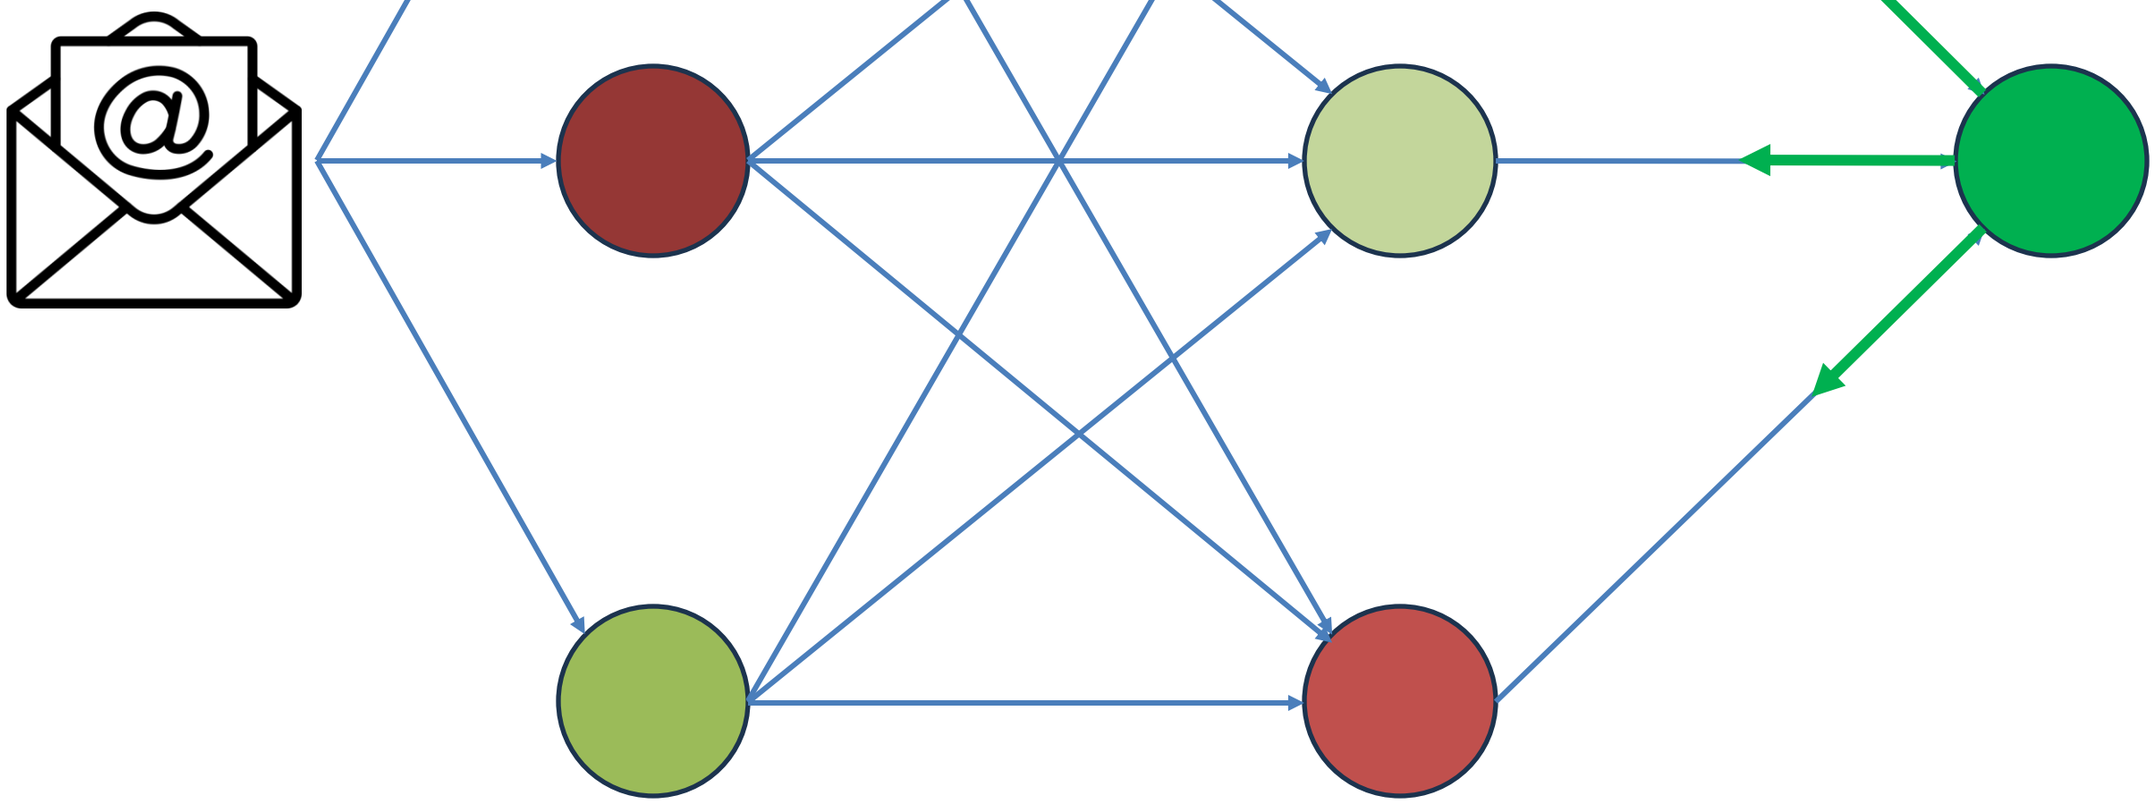
Neural Network

The features are turned into numbers which propagate through the network, with some nodes getting activated and others not



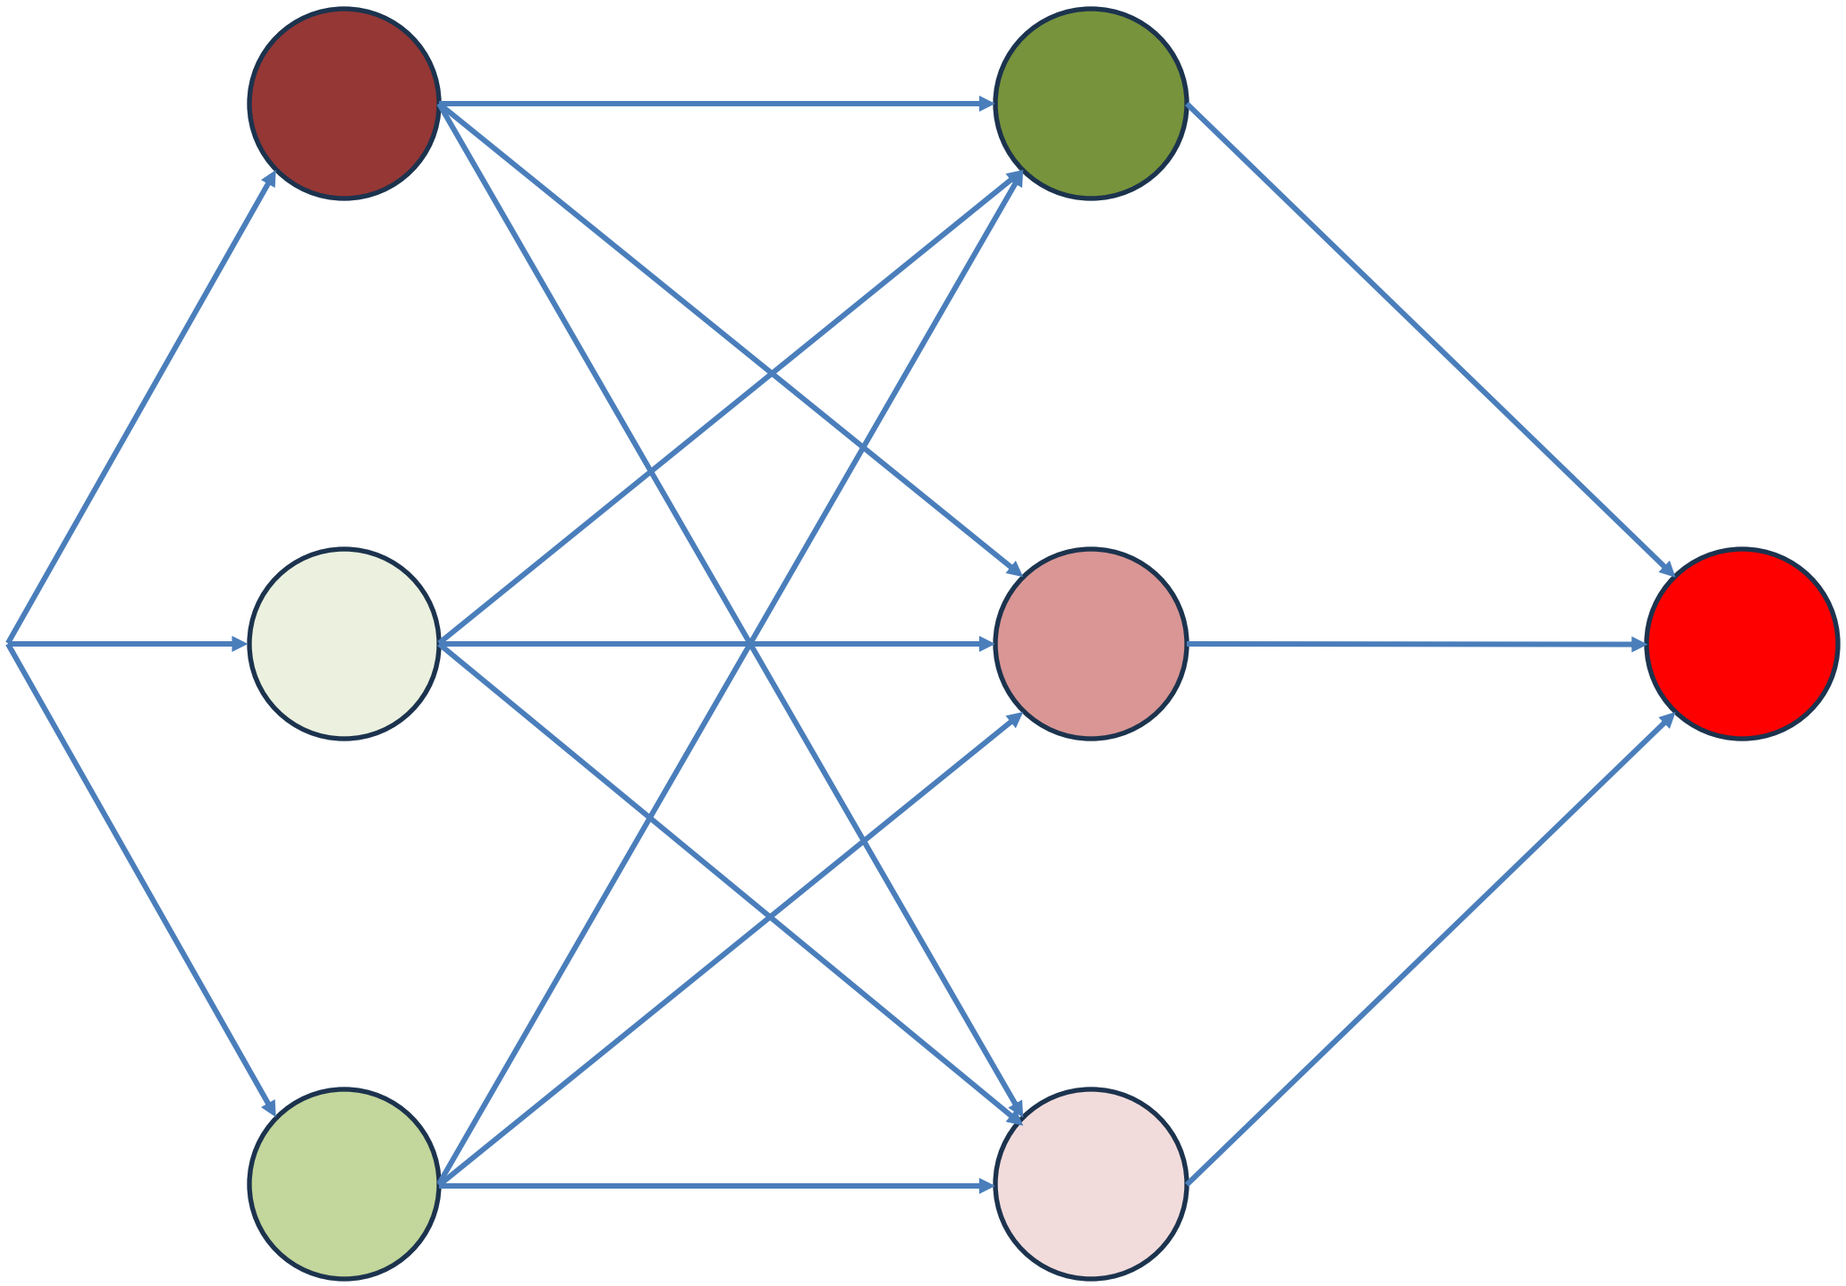
If the prediction at the end is right, this will propagate back through the system, reinforcing which nodes were activated and how

Neural Network



Neural Network





NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

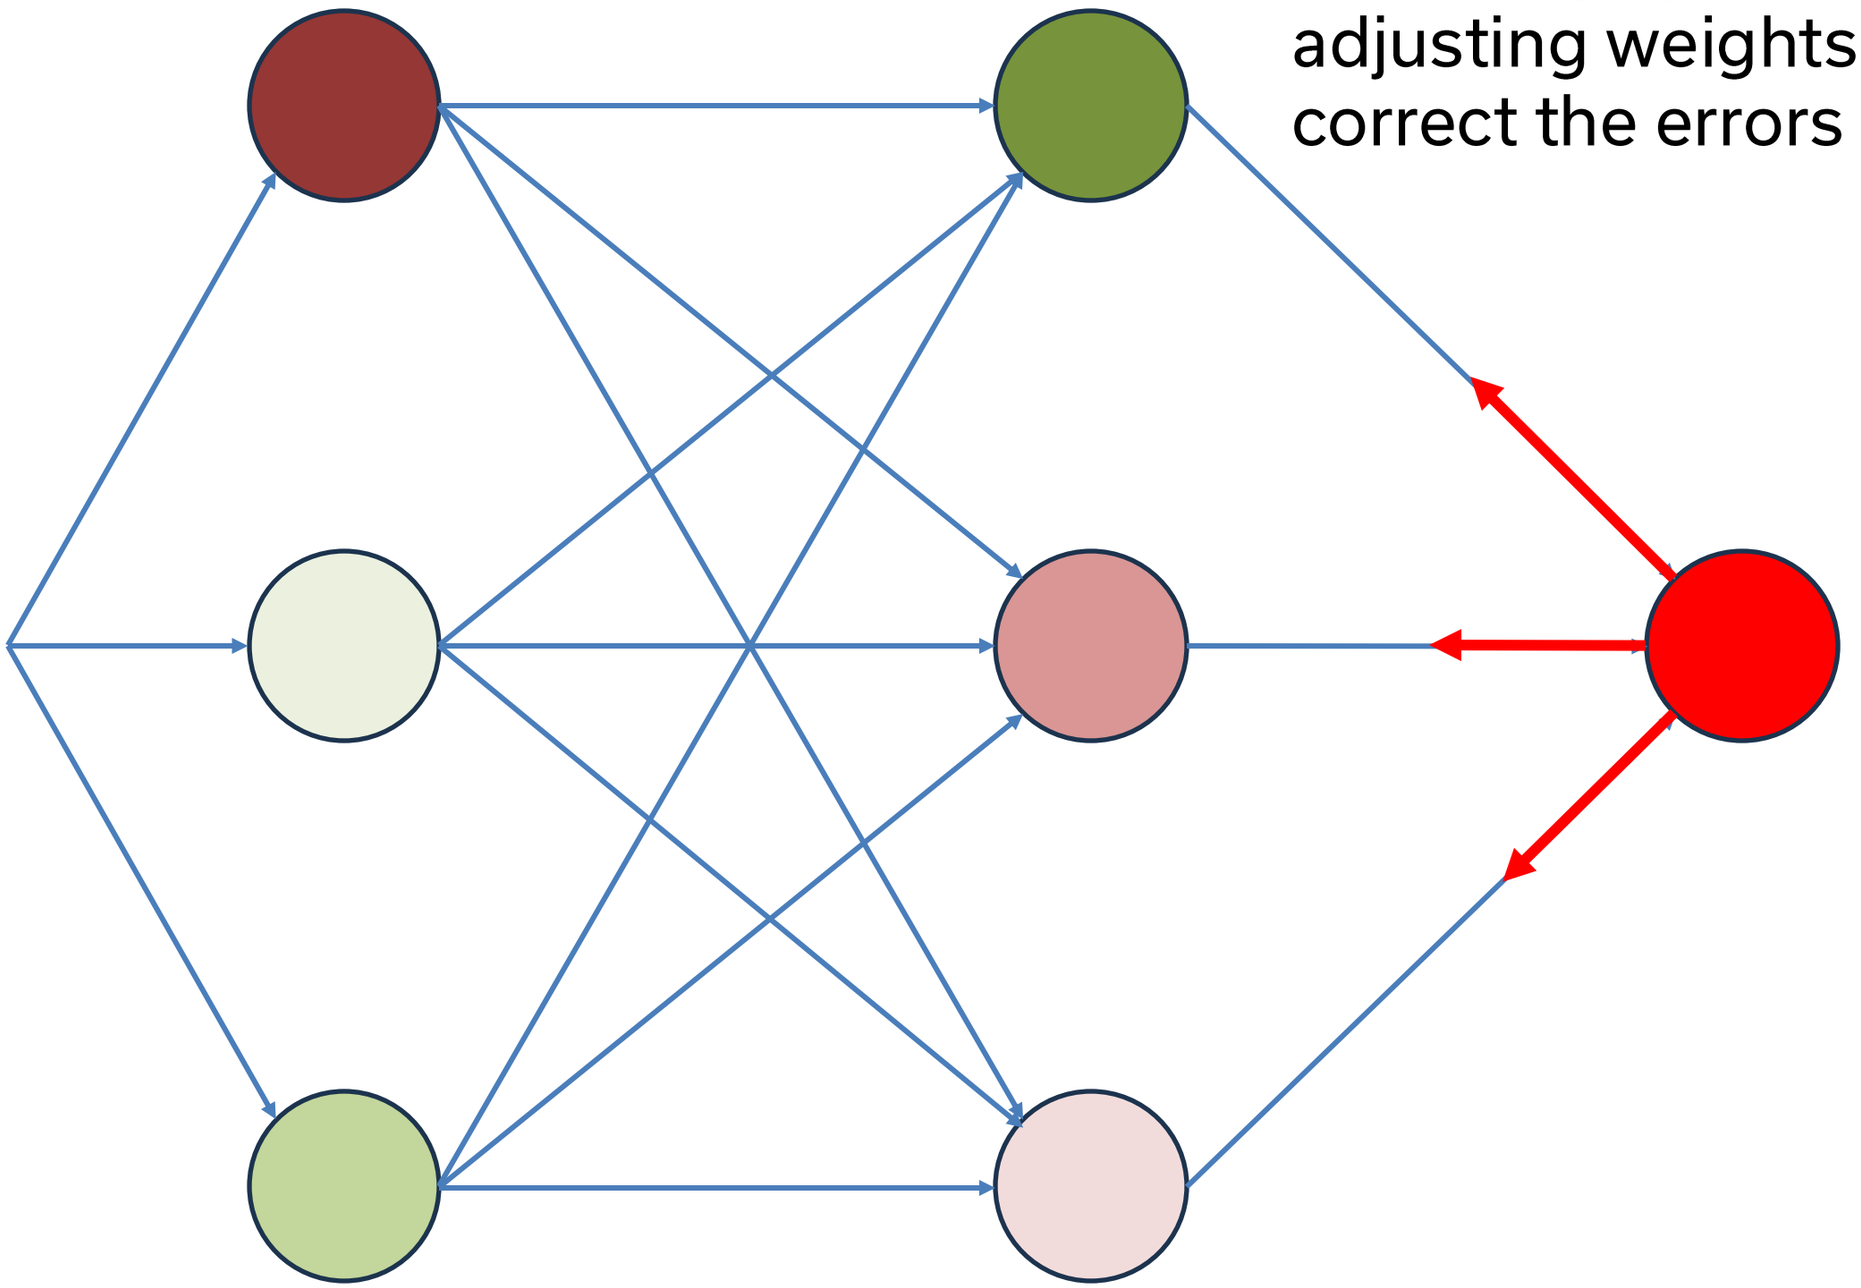
CPE Higher Education Matters

University of Kentucky

Neural Network

If the prediction is wrong, the signal is still back-propagated , adjusting weights to correct the errors



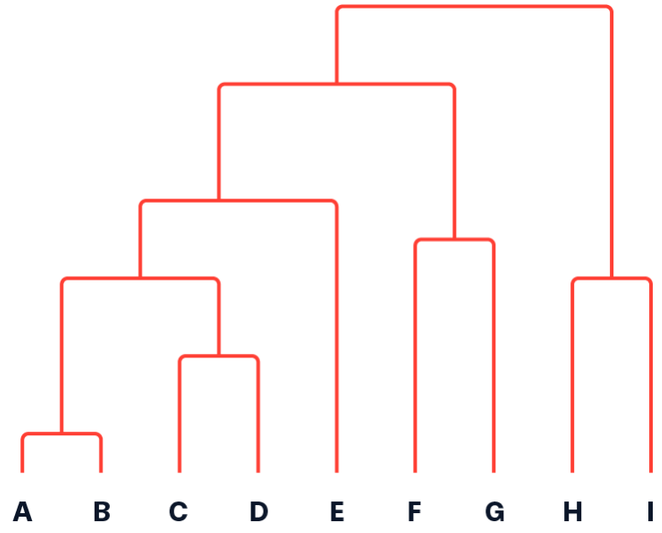


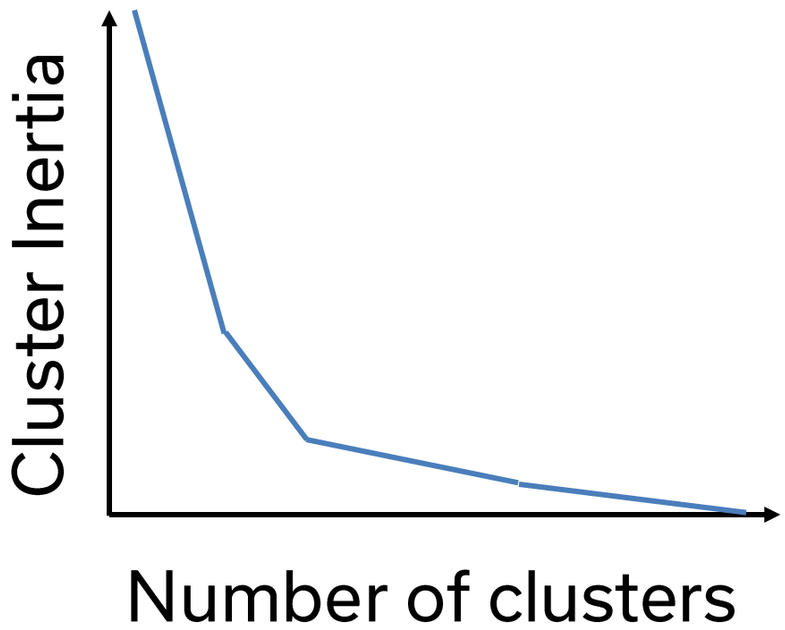
Unsupervised Learning

- The model is given unlabeled data and must find patterns or structures itself
- Can be used to identify previously unknown data trends
  - –Clustering: segmenting customers into different marketing groups
  - –Association: people with condition X also tend to have condition Y
- Many clustering techniques rely on knowing how many clusters you want beforehand
  - –Ex: algorithm to find the “best” three groupings of data points

Unsupervised Learning

- What if you don't know how many clusters you should have?
  - Build a hierarchical cluster by pairing individual points, and then pairing those pairs
  - Define "distances" between points that could indicate groupings, and the density of points will be analyzed
  - Test multiple values of k for traditional clustering and use visual analysis or calculate clustering scores to determine the optimal number of groups





Self-supervised Learning

- Modern AI models can sometimes turn an unsupervised learning problem into a supervised learning problem
- The model generates labels that it then tries to predict, rather than requiring external labels passed into it
- So what are the labels?
  - Part of the input data that is "held out" from the model

Self-supervised Learning

Input Sentence:

The chef [redacted] a delicious meal.

Now the model must determine the missing part of the sentence (the label)

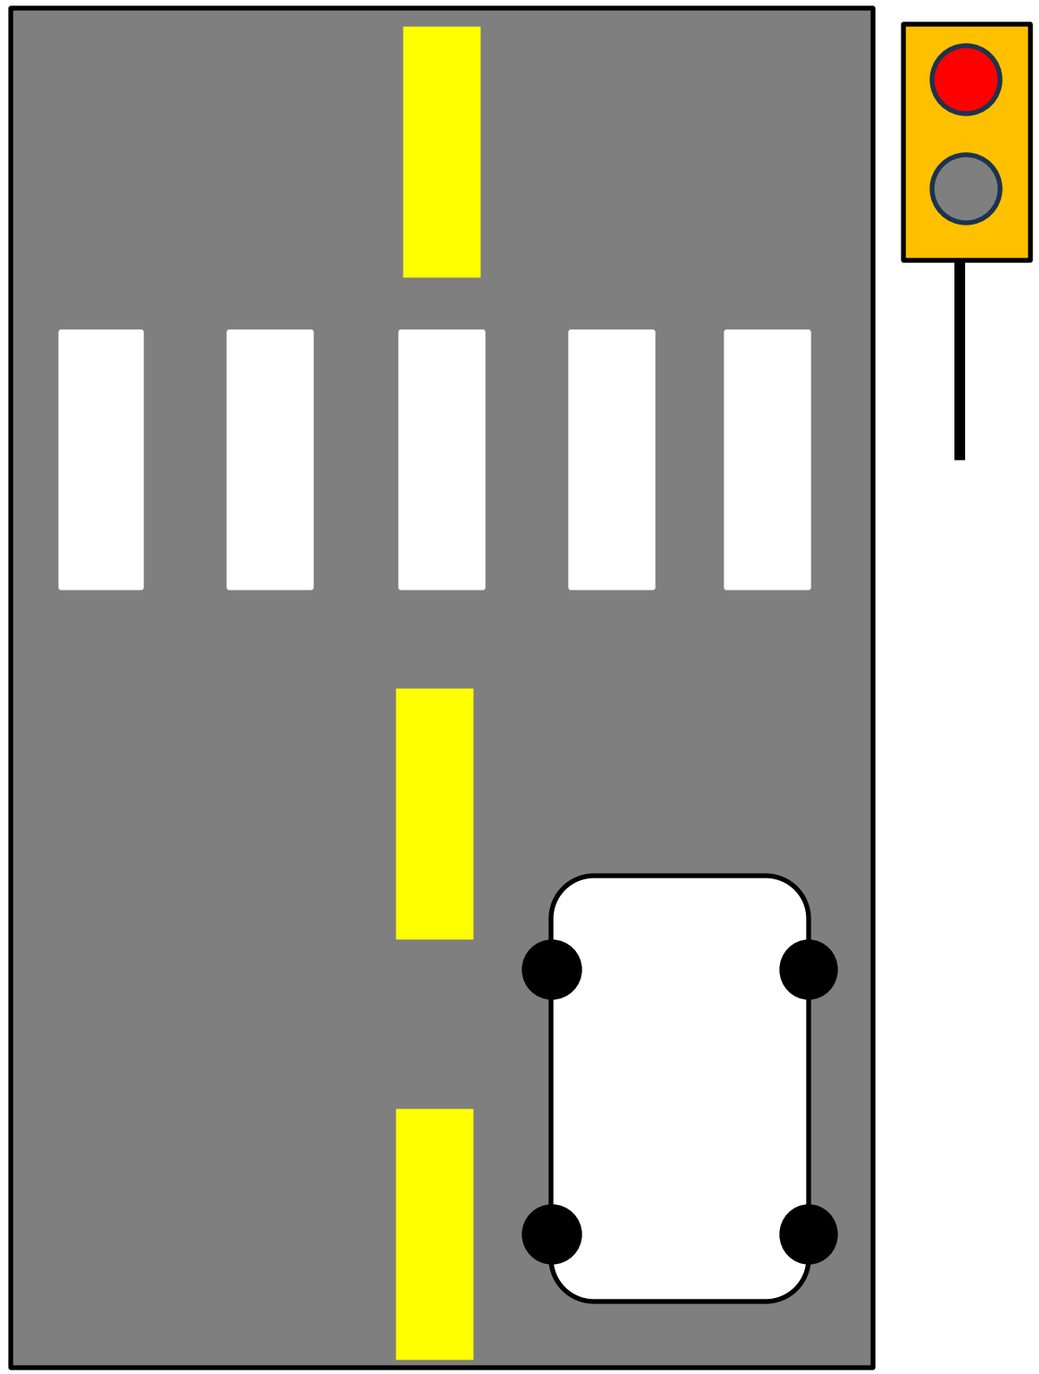
Reinforcement Learning

- An “agent” learns by performing actions and receiving rewards or punishments for those actions
- This is how AI that “makes decisions” learns
  - Chess bot
  - Robotic system learning to walk
  - Self-driving cars
- The agent can repeatedly take actions in a simulation to learn which ones are favored in certain situations

Reinforcement Learning

What is the right action to take in this situation?

In a simulation, an agent will try everything it can to see what gets rewarded and what gets punished

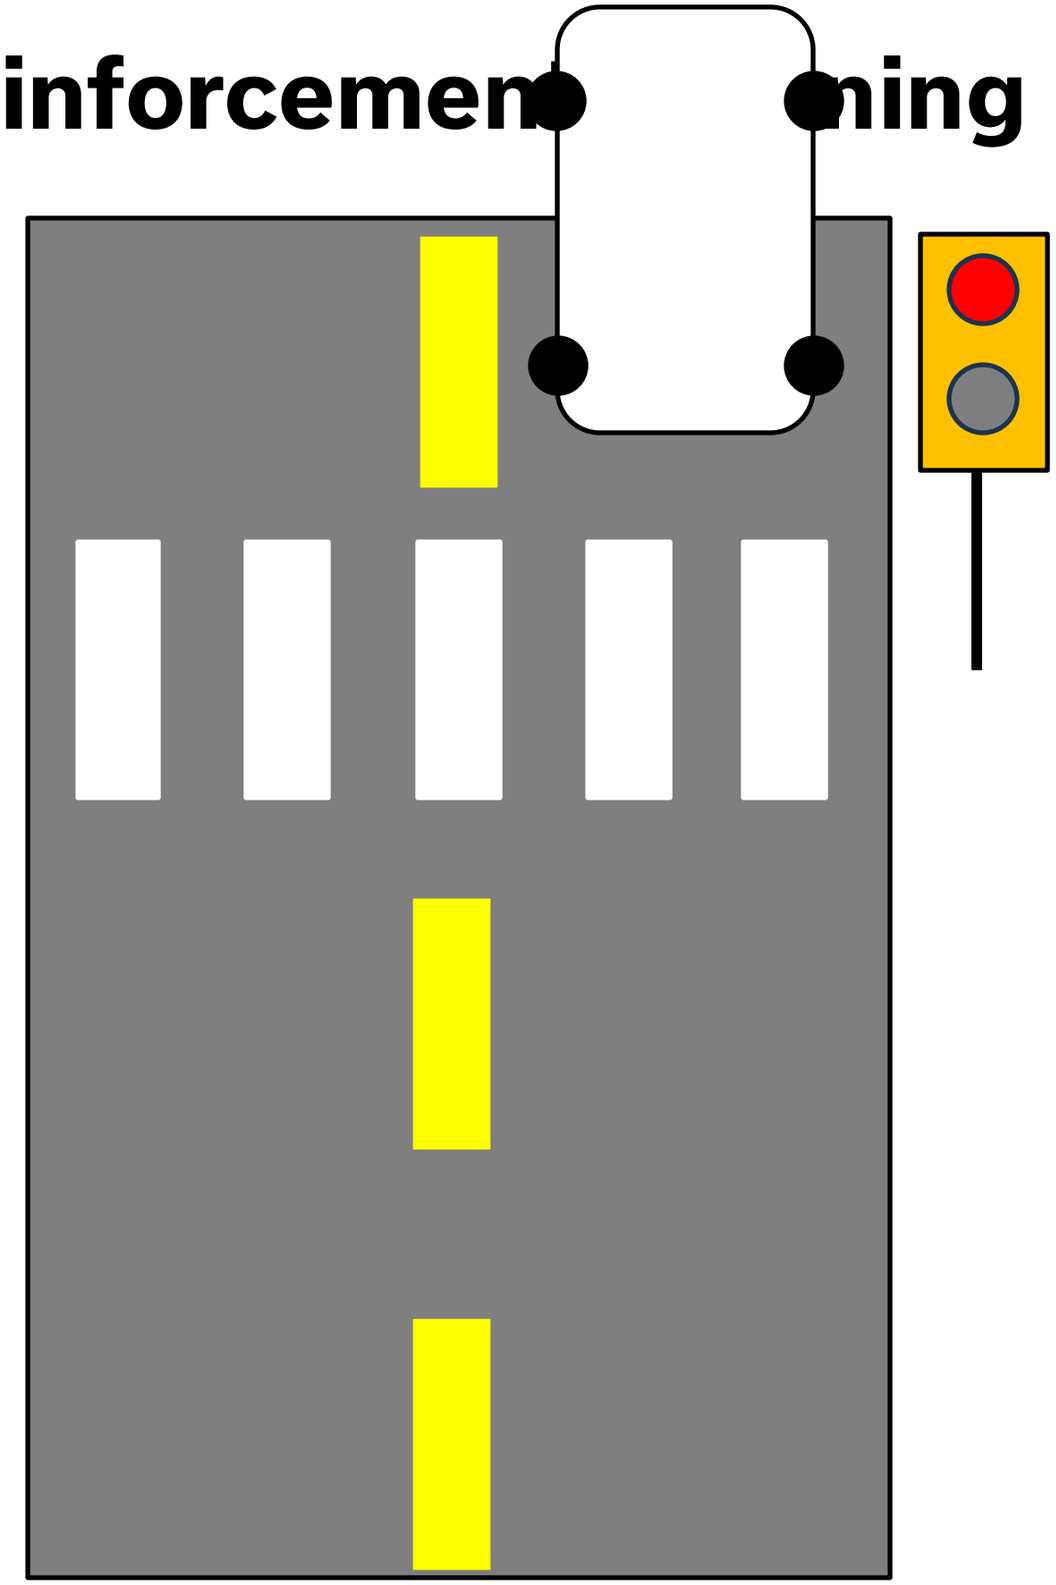


Reinforcement Learning

What is the right action to take in this situation?

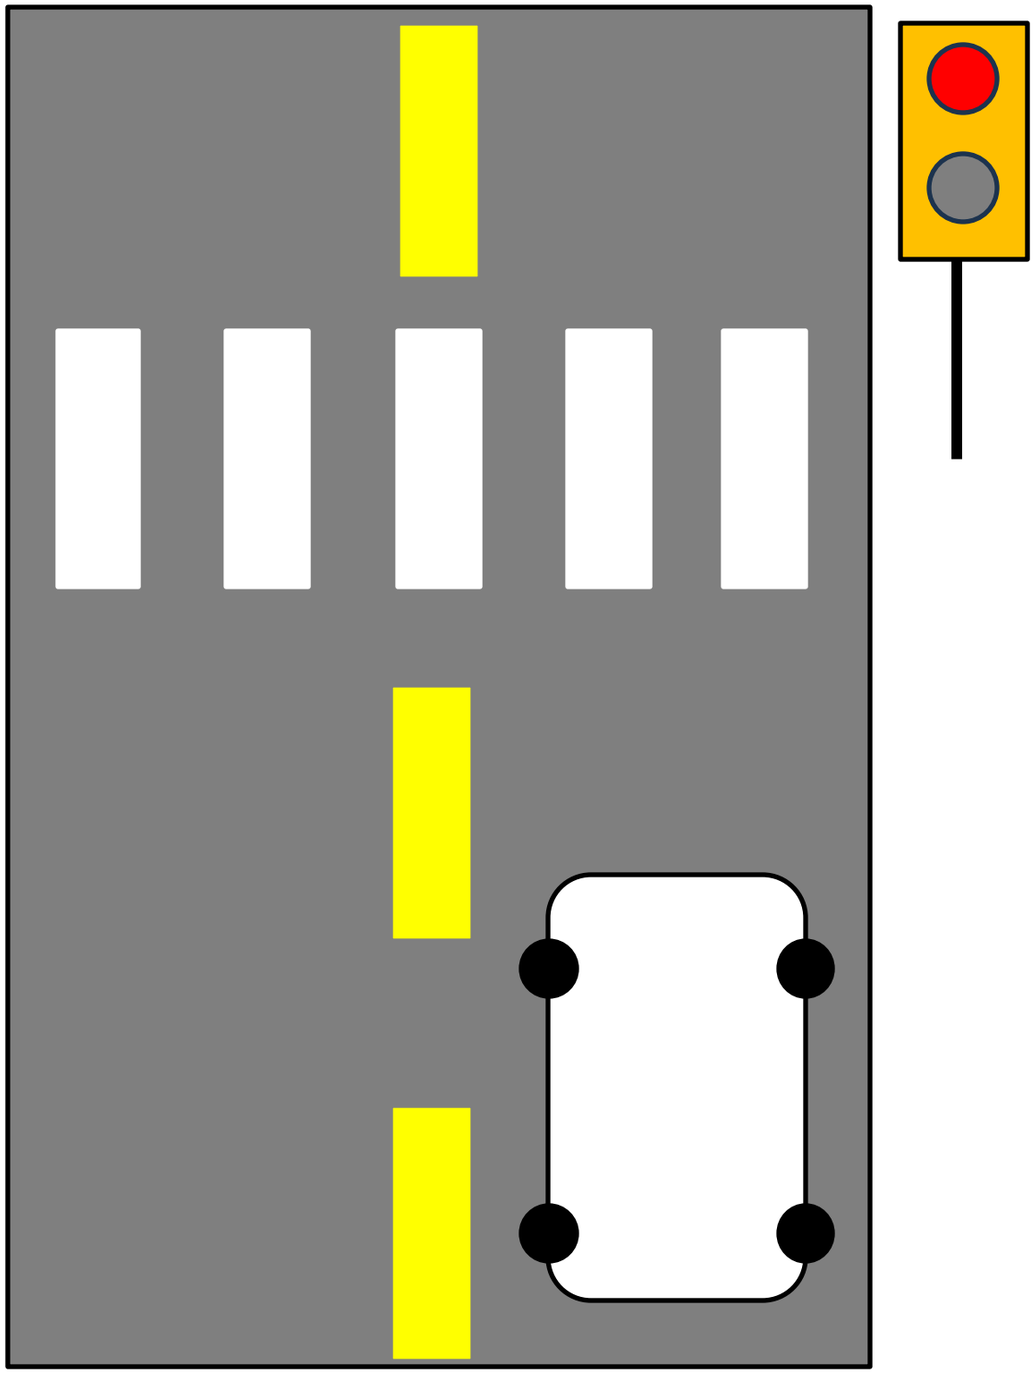
WRONG
-50 points

When an action is taken that has been defined as undesirable, an appropriate penalty to the machine's reward function will be taken



Reinforcement Learning

What is the right action to take in this situation?

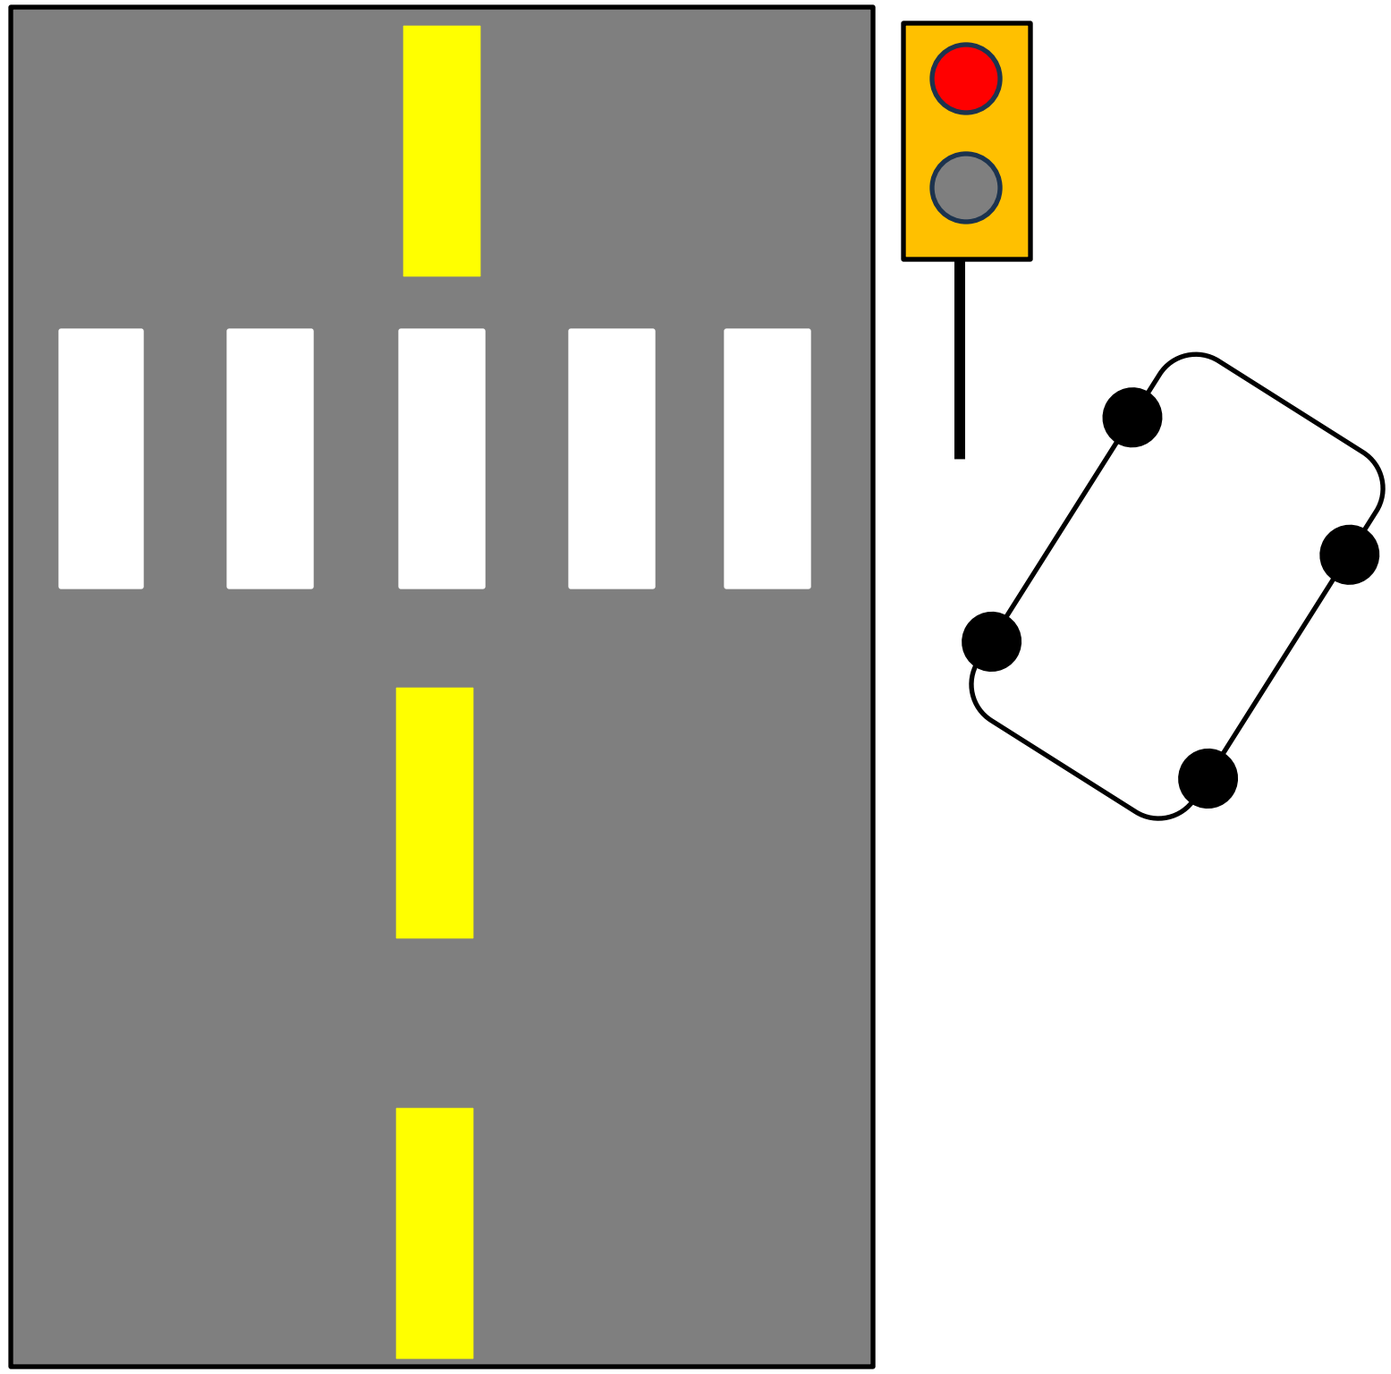


Reinforcement Learning

What is the right action to take in this situation?

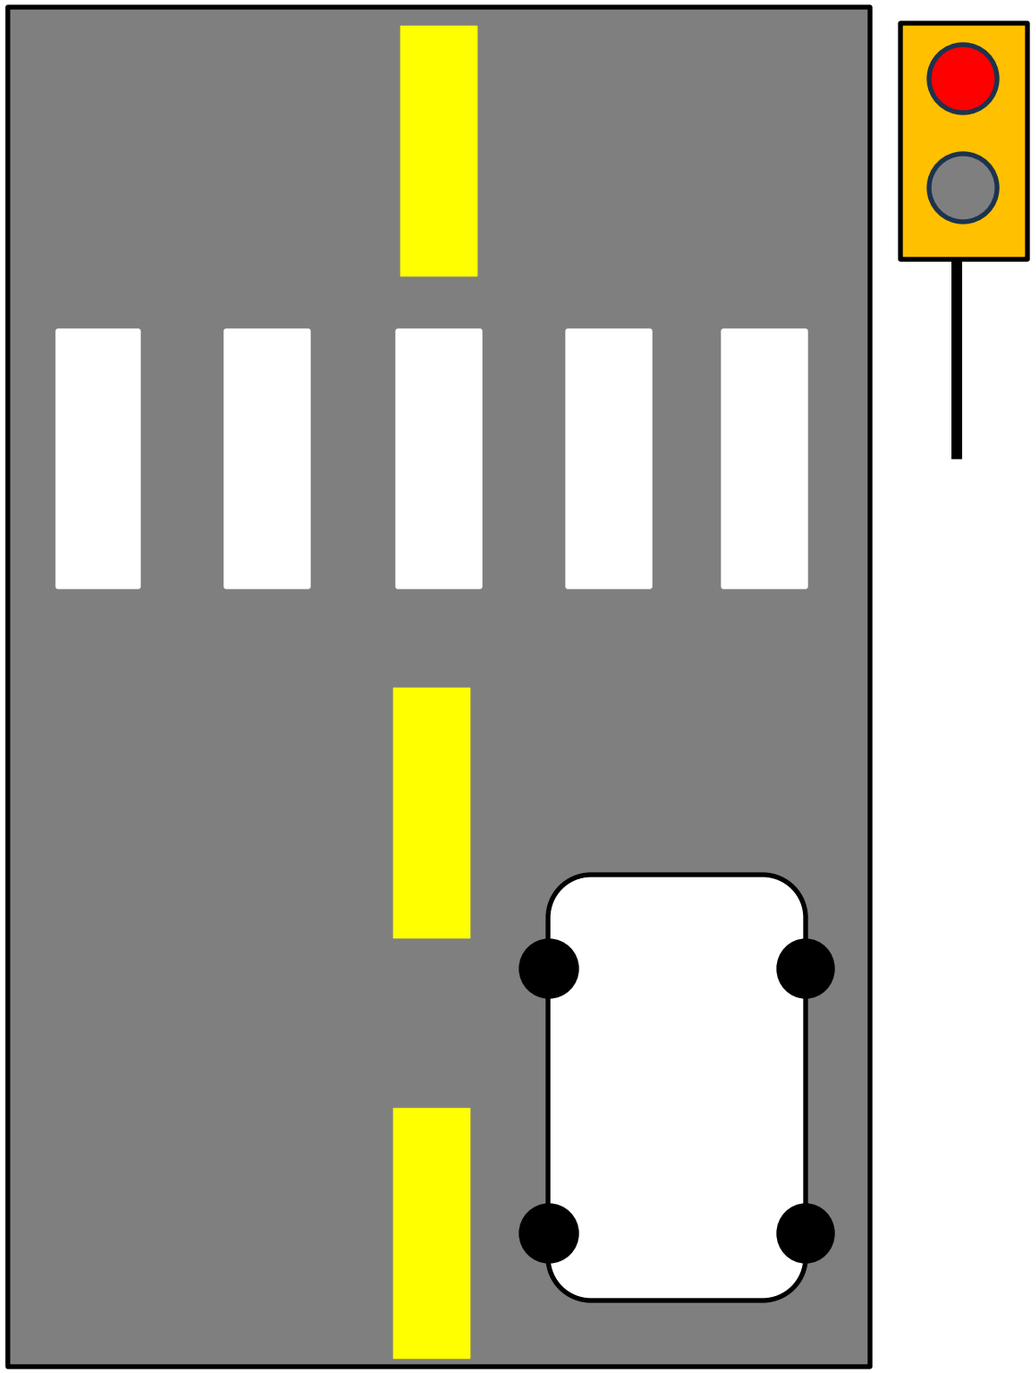
WRONG
-100 points

Different actions will have different penalties to emphasize to the machine what to do and what not to do



Reinforcement Learning

What is the right action to take in this situation?

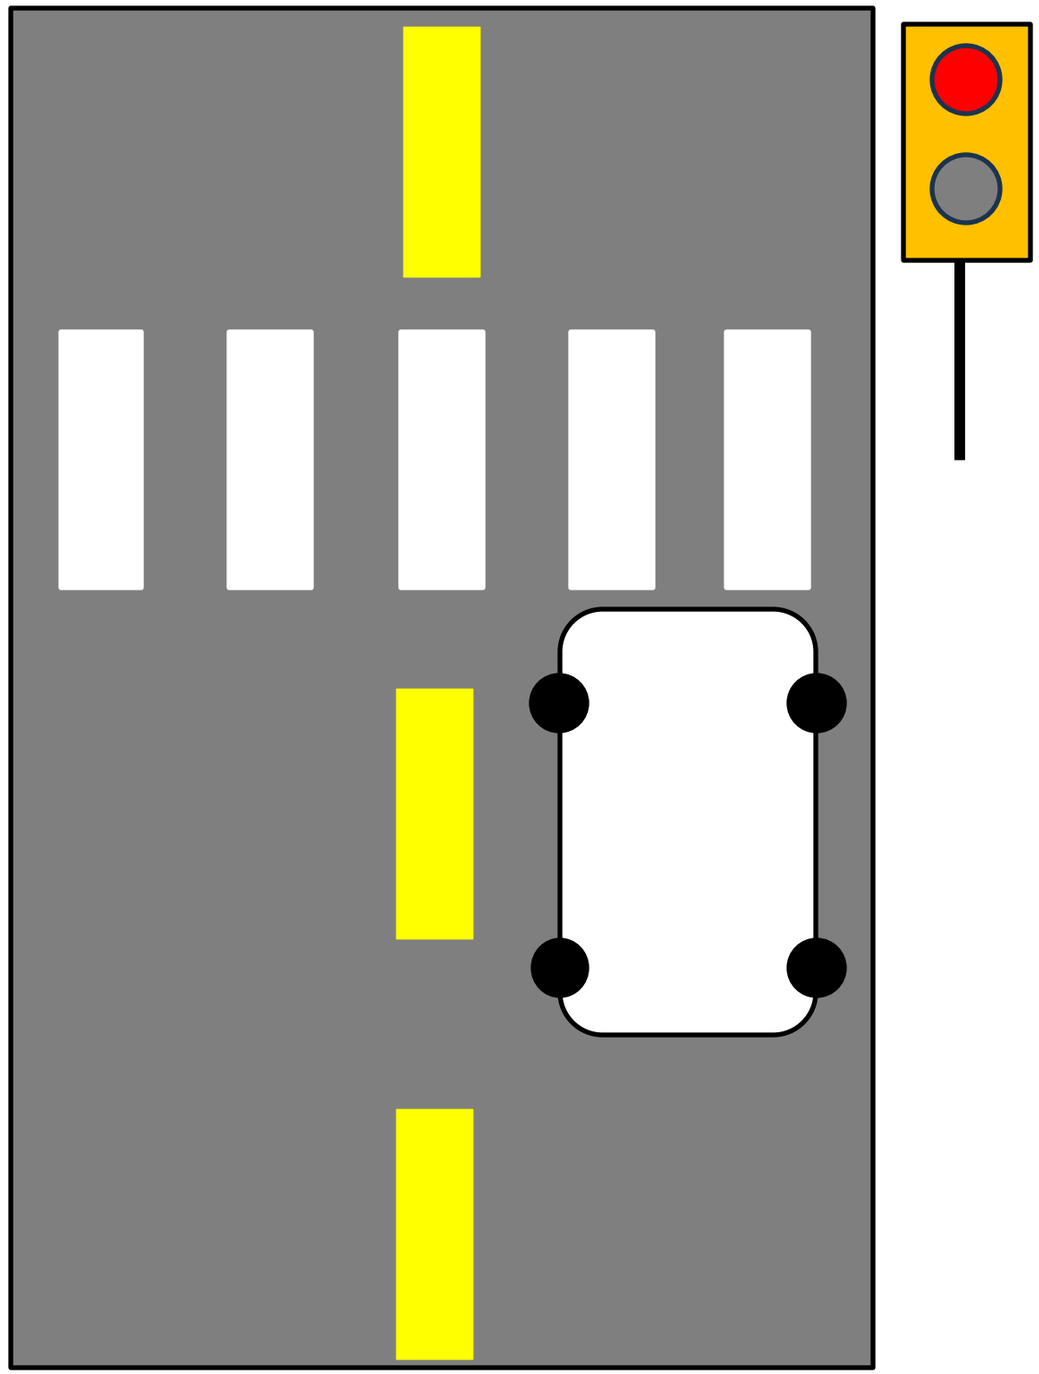


Reinforcement Learning

What is the right action to take in this situation?

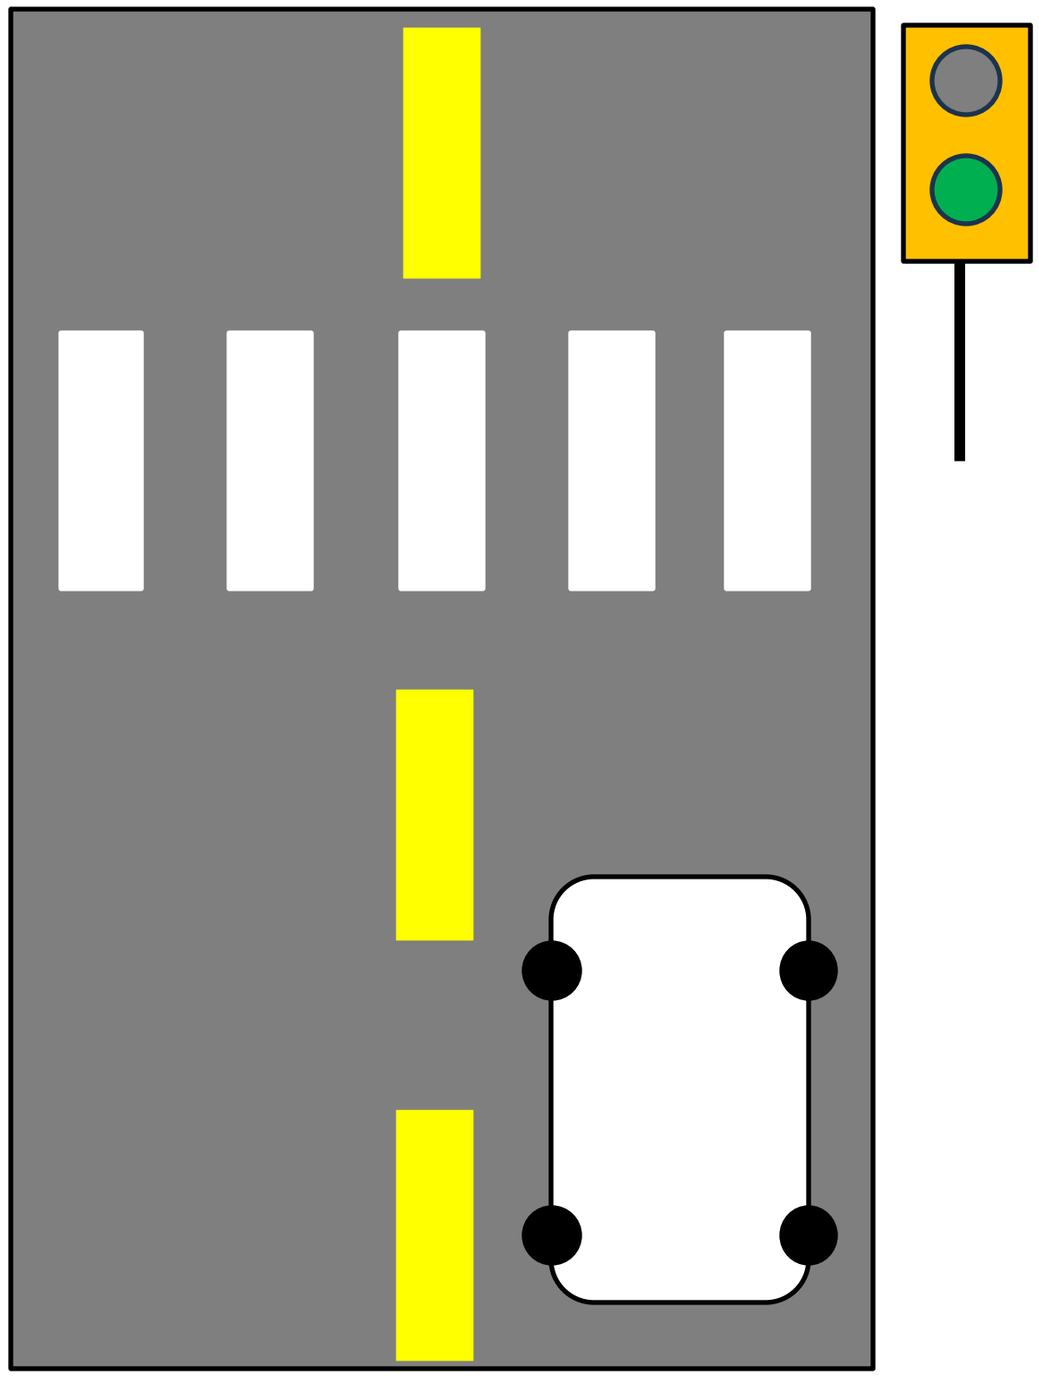
CORRECT
+100 points

When the machine does something desirable, it is rewarded



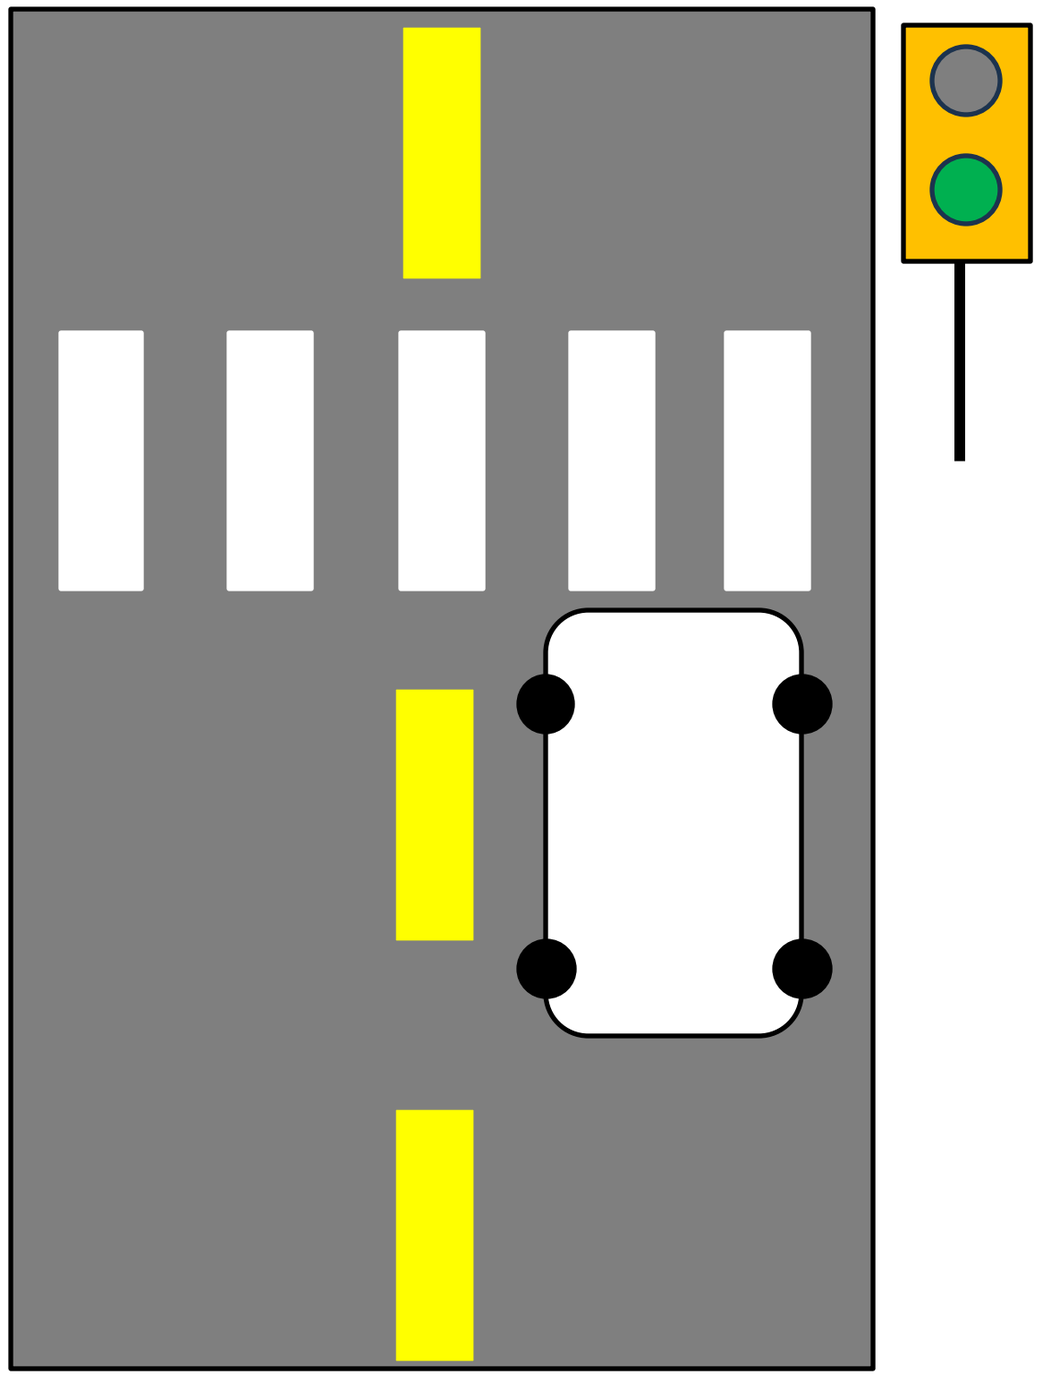
Reinforcement Learning

It will learn to take in information from its environment, as the correct action to take depends on the context



Reinforcement Learning

It will learn to take in information from its environment, as the correct action to take depends on the context

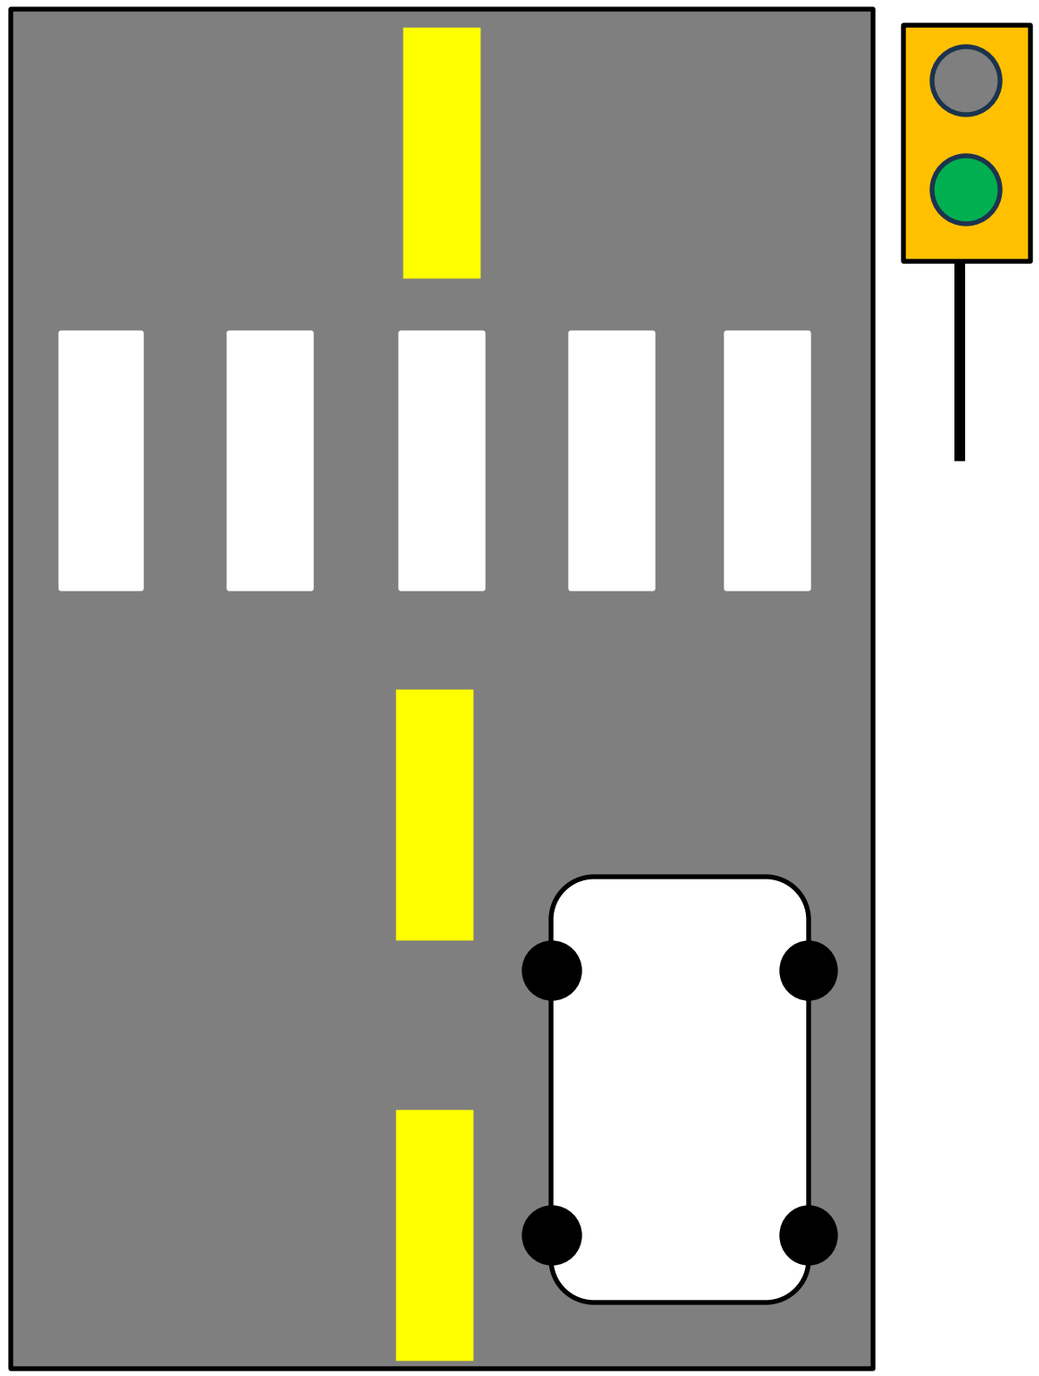


WRONG
-20 points

Now that the stoplight has changed, the action it previously took correctly is now incorrect

Reinforcement Learning

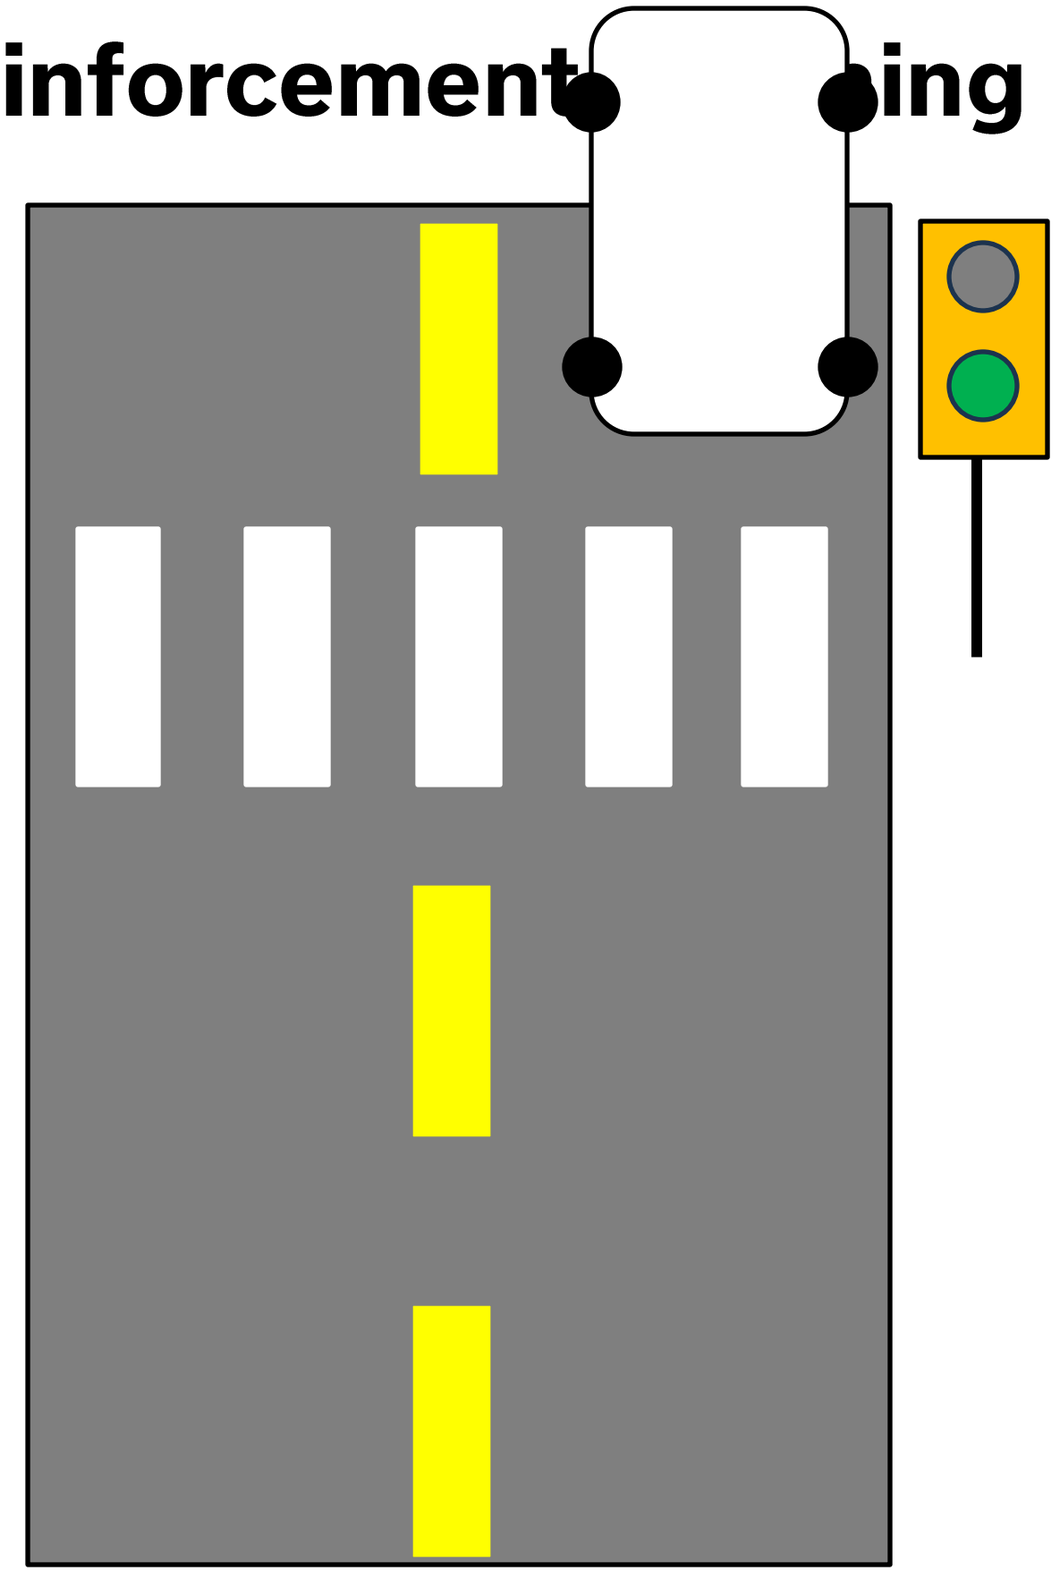
It will learn to take in information from its environment, as the correct action to take depends on the context



Reinforcement ing

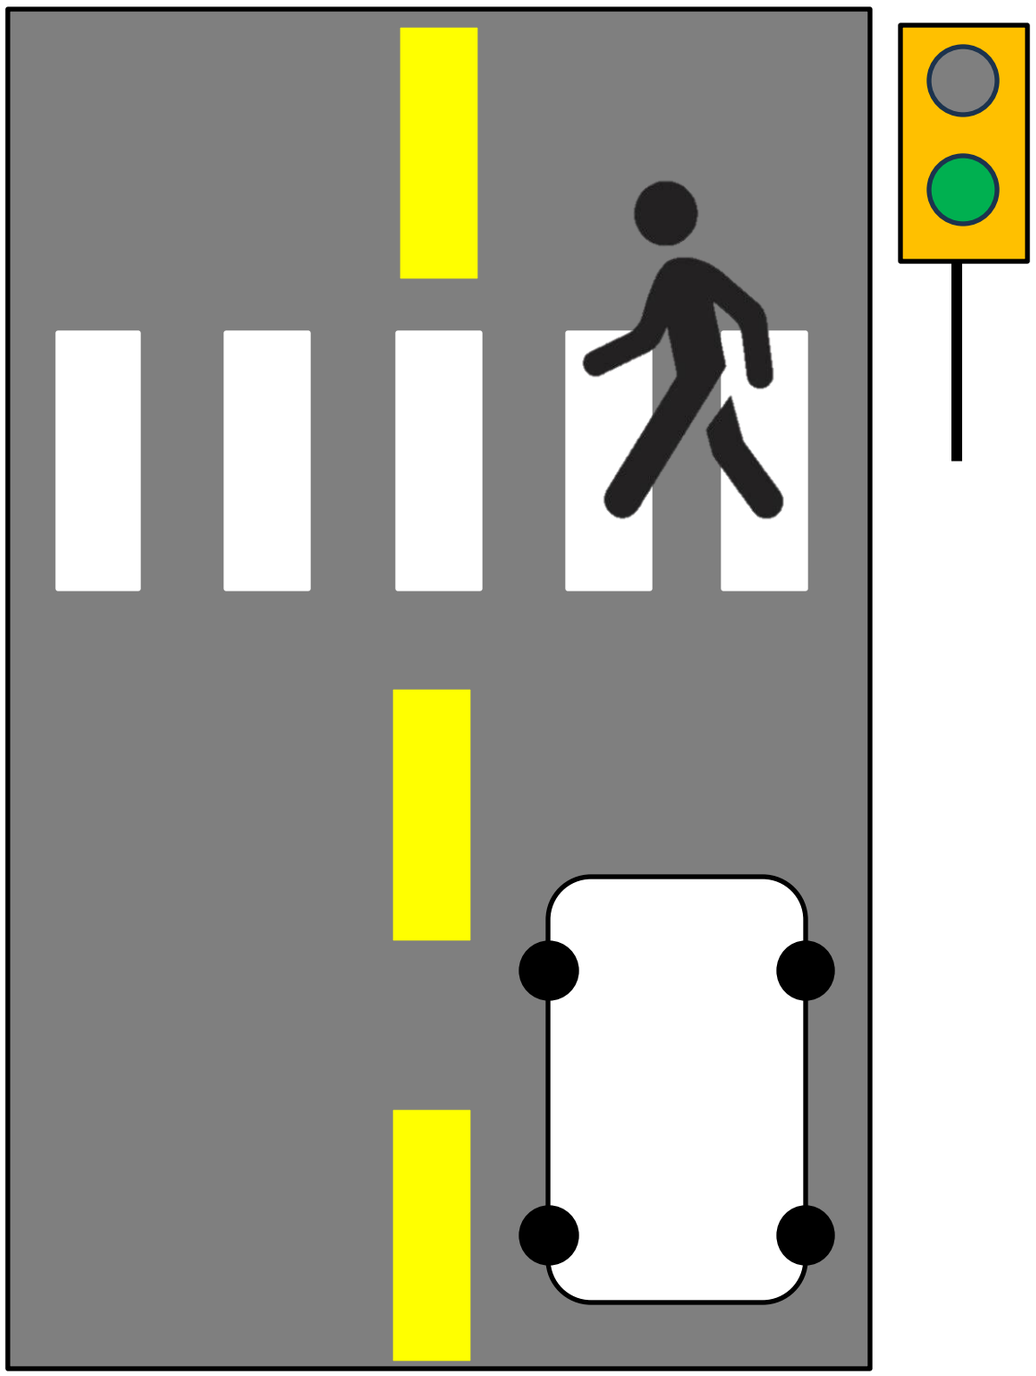
It will learn to take in information from its environment, as the correct action to take depends on the context

CORRECT
+50 points



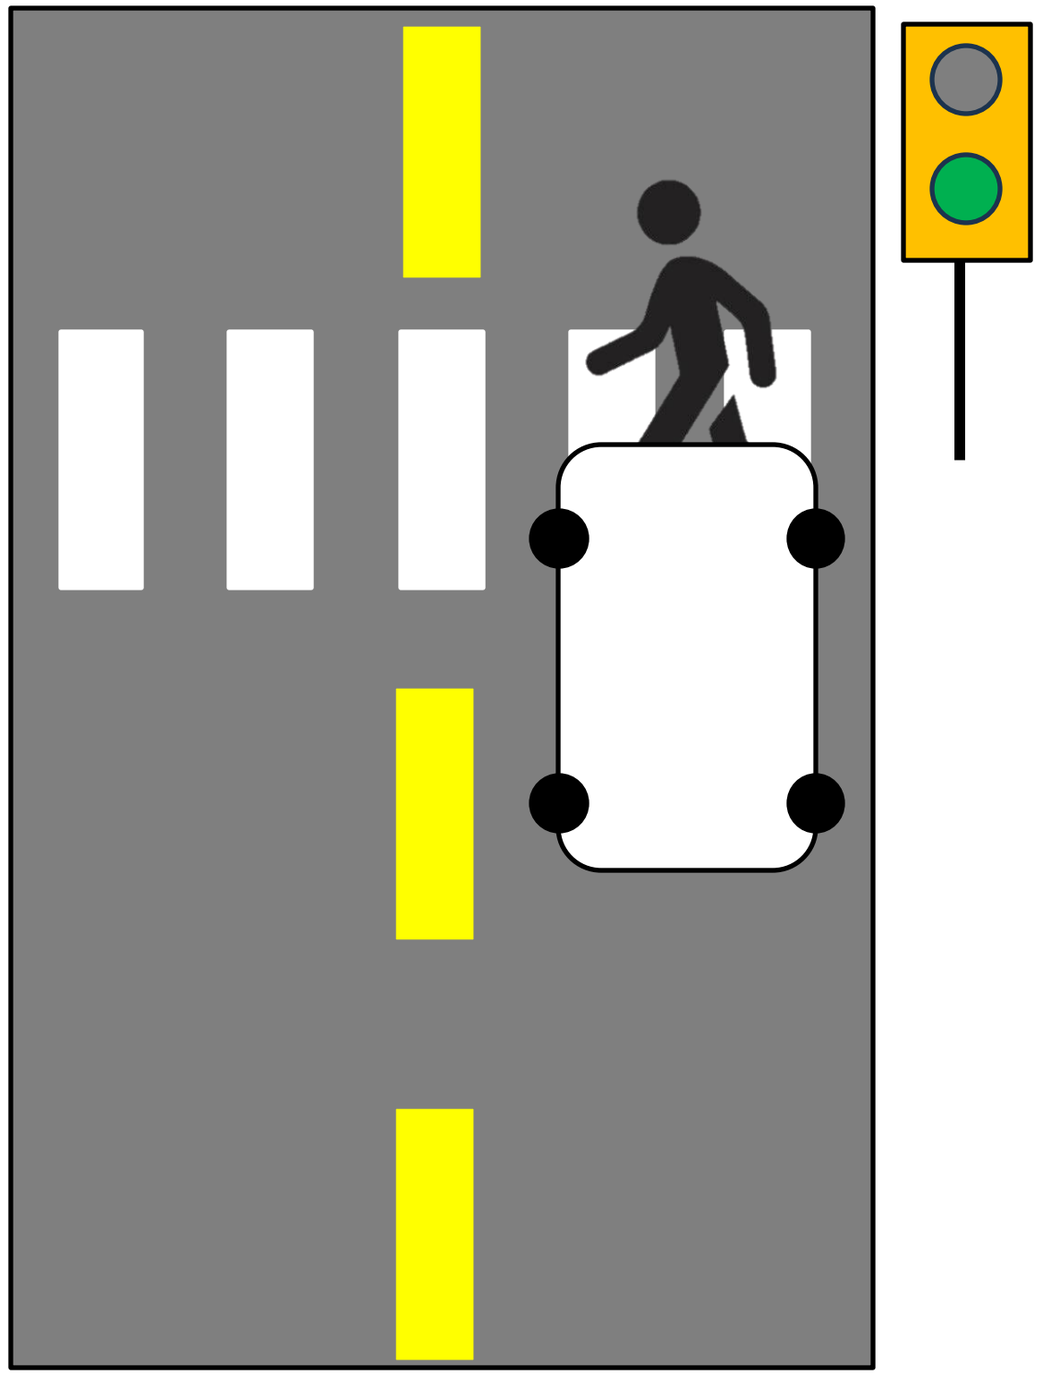
Reinforcement Learning

The rewards and punishments for each action will determine how it should weigh each option



Reinforcement Learning

The rewards and punishments for each action will determine how it should weigh each option



WRONG
-1000 points

The punishment for hitting a pedestrian is greater than the reward for obeying the green light

Q-Learning

- A reinforcement learning algorithm that builds a large table of rewards for every situation it encounters, learning the "quality" of taking each action

Stoplight?	Pedestrian?	Action	Reward
Red	No	Drive	-50
Red	No	Stop	+100
Green	No	Drive	+100
Green	No	Stop	-20
Green	Yes	Drive	-1000
Green	Yes	Stop	+100

Gaps in Reinforcement Learning

- If we have to pre-define every action and situation's rewards, then we're no better off than rule-based learning
- Several ways to address this:
  - Define only a reward for the final goal (stop the car) and automatically ripple smaller rewards back for the actions required to reach that goal (apply the brakes, steer forward, etc.)
  - Or the opposite: define rewards for small steps and let them be added up to create rewards for larger goals
    - Rather than rewarding the car for driving 1000 feet (+100), reward it for each foot it drives (+0.1) and penalize it for each second it remains still (-0.1) or moves backwards (-0.5)
  - ML can adapt to new situations by analyzing similarity to previous situations
    - It can "predict" the reward for an action by analyzing previous actions it has taken so it can act accordingly

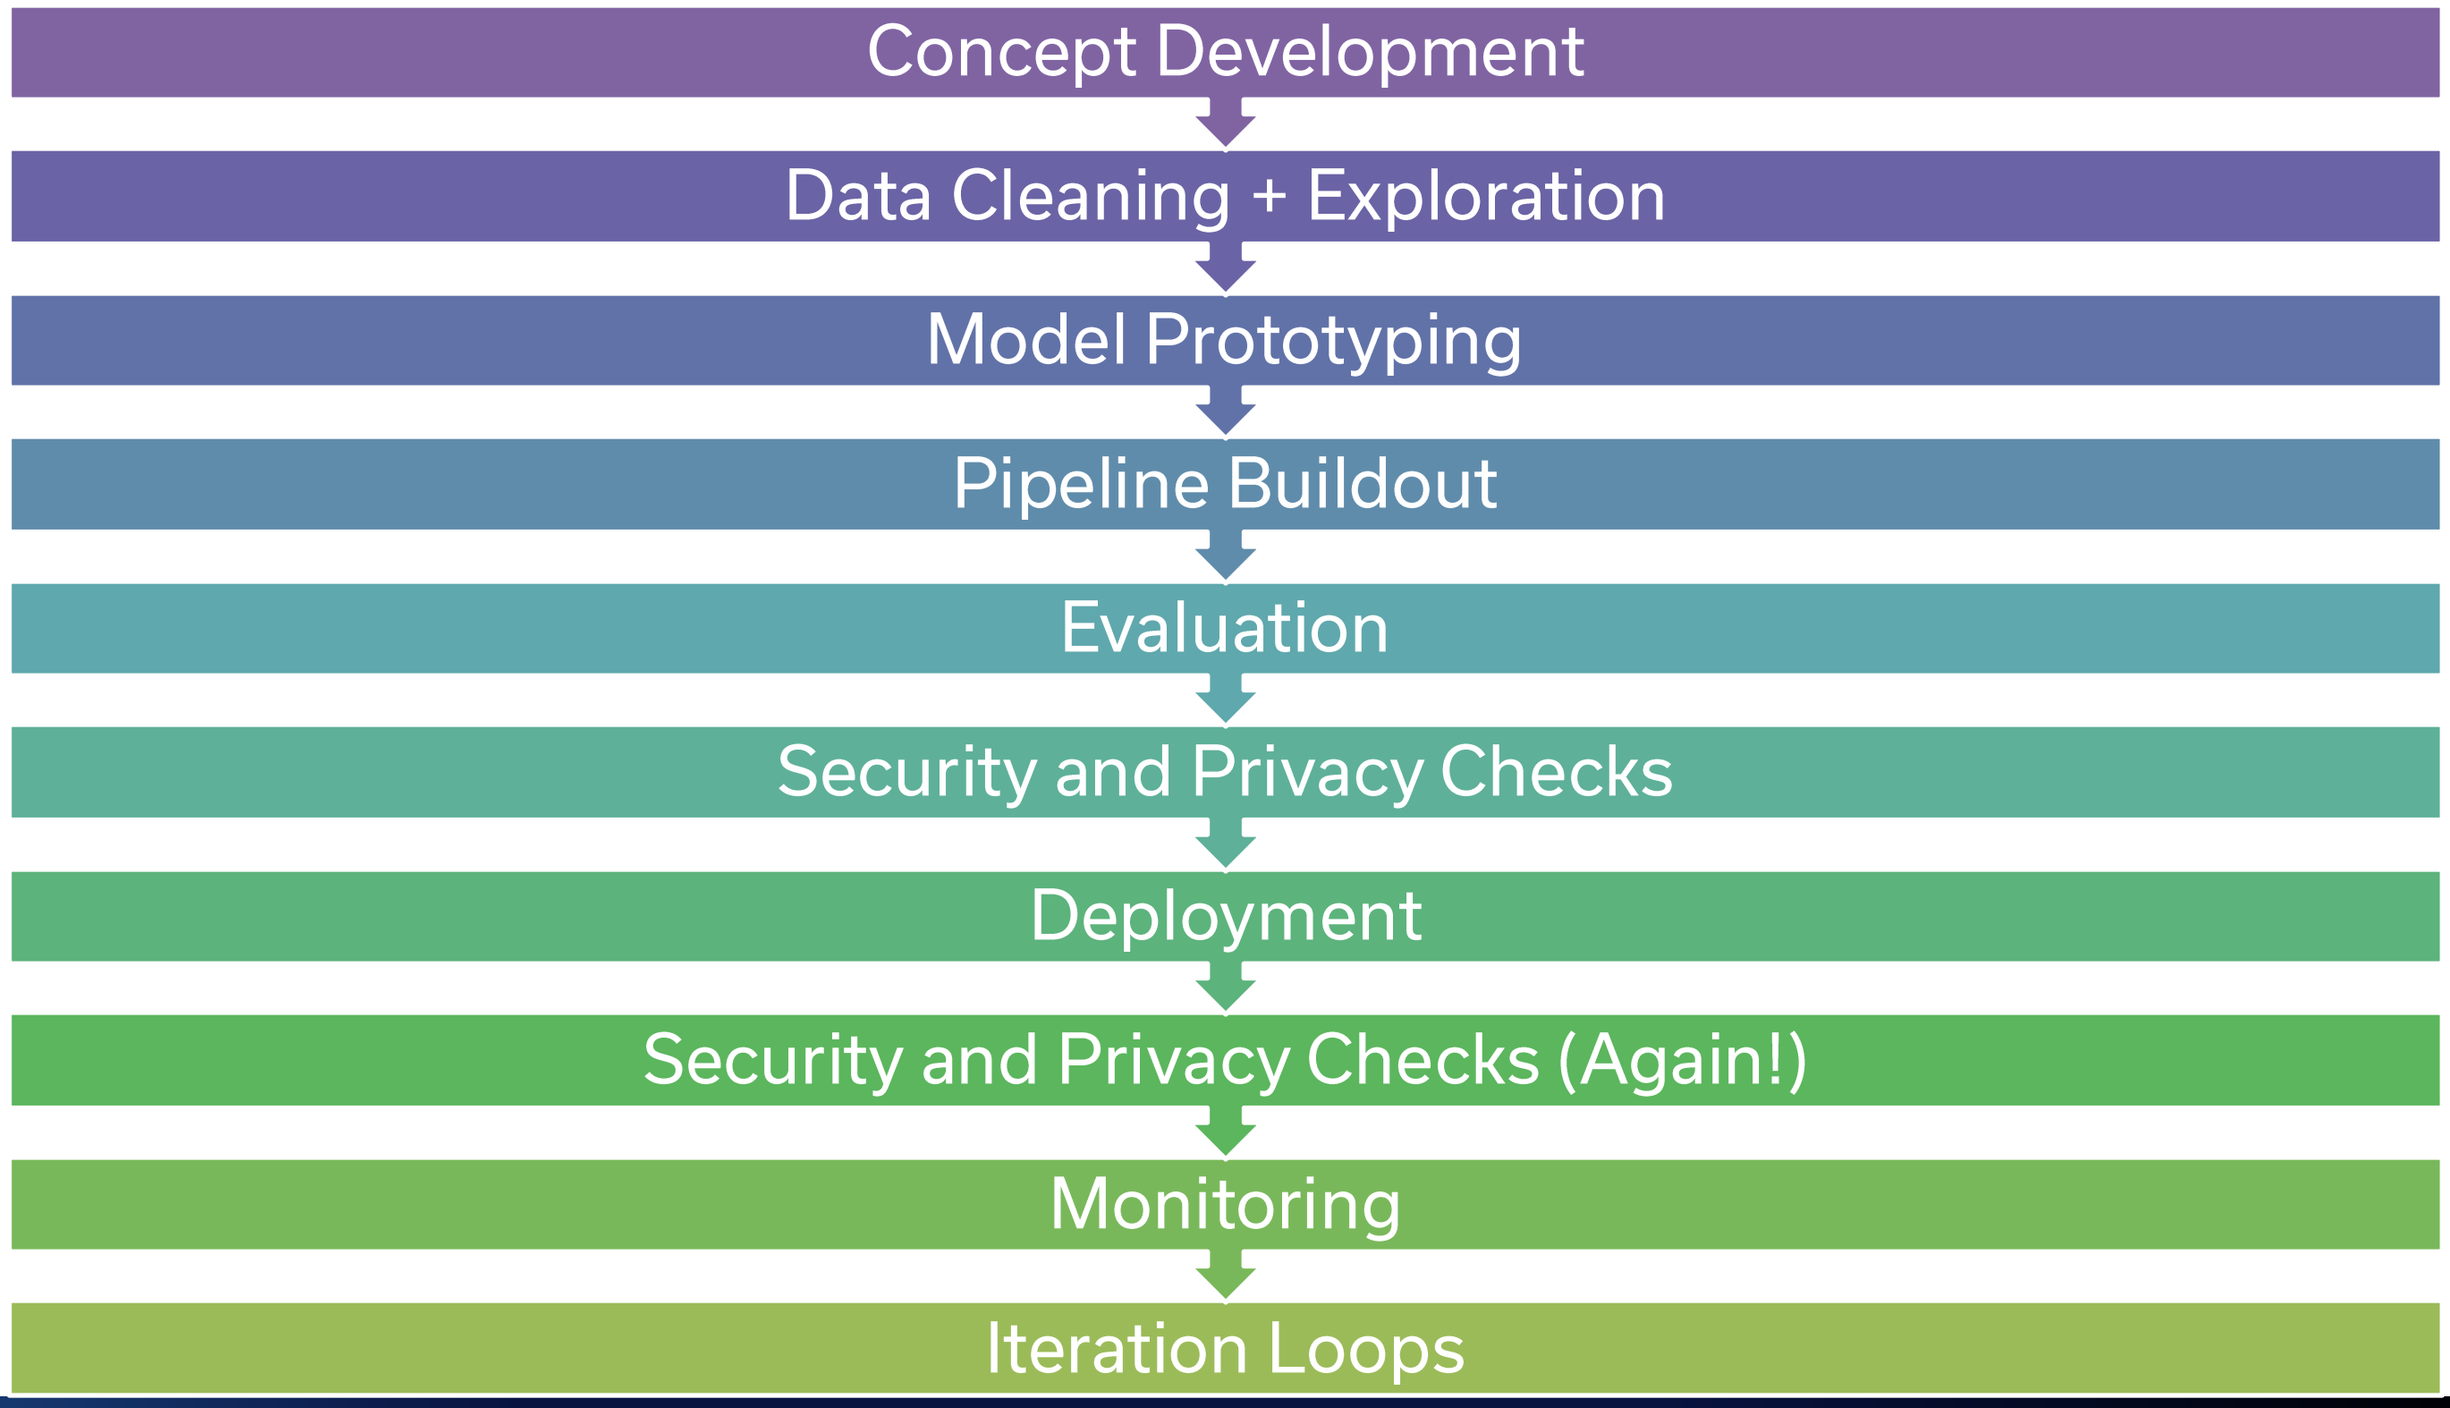
NAIRR Resources

- NAIRR provides access to many platforms that offer computing power for machine learning workflows
  - GPU Access: Indiana Jetstream 2, NCSA Delta, NVIDIA DGX, Purdue Anvil, SDSC Expanse, TACC Frontera
- Weights & Biases MLOps: tracking experiments and visualizing model performance
- Hugging Face: hosting and sharing of ML models and applications



Hugging Face

Machine Learning Pipeline



NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Machine Learning Pipeline

Concept Development

- Data gathering/surveillance across Kentucky for opioid overdose related data
- Rather than just capturing/visualizing past trends, wanted to accurately forecast future overdose trends across Kentucky
- Thus, a use case was recognized where AI could be helpful and needed

RAPID ACTIONABLE DATA FOR OPIOID RESPONSE
RADOR-KY
UNIVERSITY OF KENTUCKY

NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
CPE Higher Education Matters
University of Kentucky

Machine Learning Pipeline

Data Cleaning + Exploration

- Data cleaning can be the most difficult part of the process
- Many data questions needed to be answered:
  - What data source should be forecasted?
    - Can expand to multiple, but start with the most timely, relevant, and actionable source for the use case (Emergency Medical Services data)
  - At what geographical level should data be aggregated?
    - County-level was most applicable for statewide planning and response
  - At what temporal level should data be aggregated?
    - Monthly snapshots are specific enough to capture seasonal trends without as much sparsity as the weekly level
  - How should small, sparsely populated counties be handled?
    - Aggregating these regions to Area Development Districts (ADDs)

Machine Learning Pipeline

Model Prototyping

- Start with exploratory analysis: try different models and techniques to see what works first
- How to find models/techniques to try?
  - –Some coding knowledge is typically required at this step
  - –Many resources available, such as Python package documentation
    - Scikit-learn: https://scikit-learn.org/1.4/tutorial/index.html
    - Darts: https://unit8co.github.io/darts/
    - PyTorch: https://docs.pytorch.org/tutorials/index.html
- We wrote and published a paper about this process: https://arxiv.org/abs/2410.16500

Machine Learning Pipeline

Model Prototyping

- Implementing a simple model in Python is easier than you might expect!

```python
from sklearn.datasets import load_iris
from sklearn.model_selection import train_test_split
from sklearn.ensemble import RandomForestClassifier
from sklearn.metrics import accuracy_score

# 1. Load some example data (Information about flowers)
data = load_iris()
X, y = data.data, data.target

# 2. Split data: 80% for "learning", 20% for "testing" the model later
X_train, X_test, y_train, y_test = train_test_split( *arrays: X, y, test_size=0.2)

# 3. Initialize the Random Forest model
model = RandomForestClassifier()

# 4. Training: The model looks for patterns in the data
model.fit(X_train, y_train)

# 5. Evaluation: How well did it learn?
predictions = model.predict(X_test)
score = accuracy_score(y_test, predictions)

print(f"Model Accuracy: {score * 100}%")
```

Machine Learning Pipeline

Pipeline Buildout

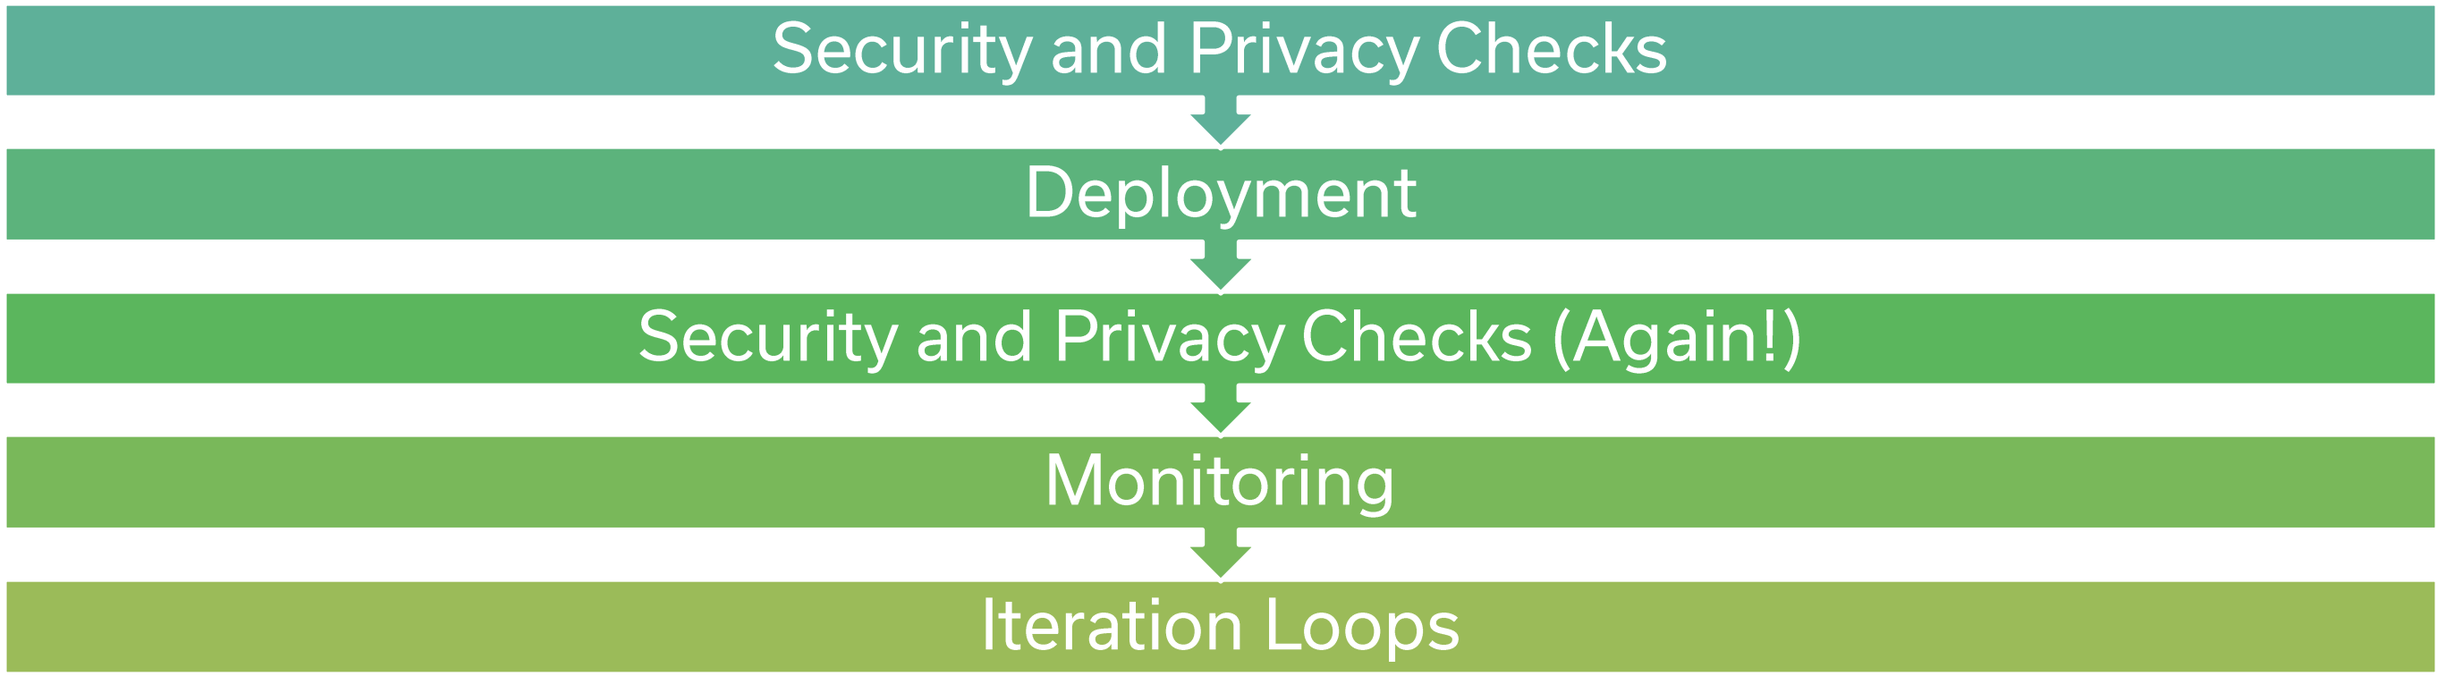
- The key is to start small and build it out piece by piece
- Examples:
  - First, we replaced our manual data upload with a process to read in the data directly from the RADOR database (MSSQL -> Python)
  - Then, rather than saving the output results to a file, send it to a different table in the database (Python -> MSSQL)
  - Next, automate the job to run regularly to update forecasts as new data comes in, rather than requiring manual action (Cron)
- Eventually, we had a working pipeline!

Machine Learning Pipeline

Evaluation

- Continual evaluation performed throughout model prototyping as well
- We added a step into the pipeline to record the accuracy of each forecast as its made, comparing previous forecasts with new, real data as we get it
- This ensures we have a process for continuous monitoring of model performance over time

Machine Learning Pipeline



- Predictions are displayed on the RADOR-KY website now (https://rador-ky.uky.edu/)
- Automated email updates about new forecast generation ensures we can monitor for both successes and errors
- RADOR-KY is years into development- this process takes time!

Conclusion

- Machine learning covers a wide range of models, systems, and applications
- At its core, it's about having a machine learn the rules of the situation instead of teaching it the rules
- With sufficient high-quality data, machine learning can match human expertise in a variety of fields

NAIRR Pilot Resources & Allocation Process

Caroline Leach

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
CPE Higher Education Matters
University of Kentucky

Bridging the Gap: An Introduction to the NAIRR Pilot

- Accelerate AI and AI-powered Discovery and Innovation
- Expand the AI workforce and train the next generation of AI researchers and educators
- Increase capacity, integration, and use of world-class public and private-sector AI resources
- Advance AI interpretability, security, and trust

nairrpilot.org/about

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

NAIRR Pilot

National Artificial Intelligence Research Resource Pilot

- Pilot program led by the NSF that brings together a coalition of federal agencies, industry partners, and non-profit organizations
  new.nsf.gov/focus-areas/artificial-intelligence/nairr

- Connects U.S. **researchers** and **educators** to computing capacity, tools, data, and training resources
  nairrpilot.org/about

Requesting Resources via the National AI Research Resource (NAIRR) Pilot (2025), Lauren Michale, Timothy Middelkoop

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

ICPE Higher Education Matters

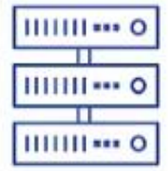
University of Kentucky

14 agencies and 28 industry and non-profit partners are contributing *in kind* state-of-the-art resources

DARPA
DOD
DOE
DoEd
FDA
NASA
NIH
NIST
NOAA
NSF
USDA
USGS
USPTO
VA

- AI2: Allen Institute for AI
- AMD
- Amazon Web Services
- Anthropic
- Cerebras
- Databricks
- Datavant
- EleutherAI
- Google
- Groq
- Hewlett Packard Enterprise
- Hugging Face
- IBM
- Intel
- **Lexset**
- Meta
- Microsoft
- MLCommons
- NVIDIA
- Omidyar Networks
- OpenAI
- OpenMined
- Palantir
- Regenstrief Institute
- SambaNova Systems
- Vocareum
- **Voltage Park**
- Weights & Biases

Research Resources



Computing



Datasets



Models, software, platforms



Educational/training opportunities



Collaborations

NAIRR Pilot

nairrpilot.org/pilotresources

Requesting Resources in the National AI Research Resource (NAIRR) Pilot (2025), Lauren Michale, Timothy Middelkoop

NAIRR Pilot

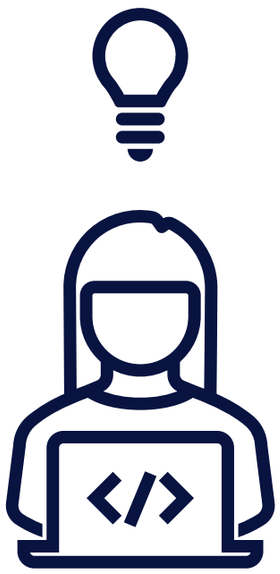
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

ICPE Higher Education Matters

University of Kentucky

NAIRR removes cost and access barriers so researchers can experiment and innovate

- Imagine a researcher with a breakthrough idea, but no dedicated cyberinfrastructure to support AI experimentation at scale and limited funding.
- How do they start?

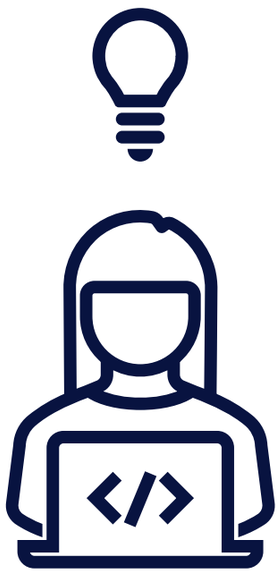


Idea → Experimentation

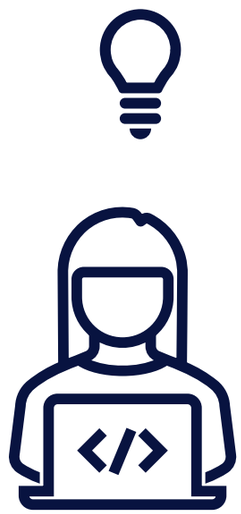


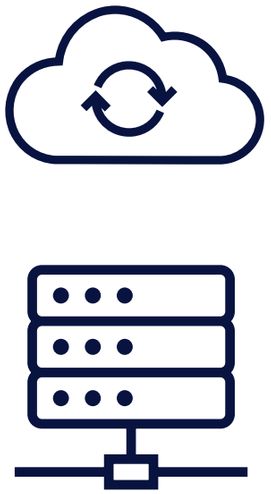
NAIRR removes cost and access barriers so researchers can experiment and innovate

- Great ideas don’t always come with the required resources
- Individual subscriptions don’t always equate to research-scale infrastructure



Idea → Experimentation







NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

NAIRR removes cost and access barriers so researchers can experiment and innovate

- “Tokens” – Small pieces of text that language models read and generate; measures how much a model is used
- Scaling an LLM-based tool for hundreds of users, or a large volume of documents accumulates costs quickly.

What resources are available?

These resources represent leading-edge offerings from corporate and non-profit organizations

- Model & Inference Services
- Cloud Training Environments
- High-Performance GPU



NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

What resources are available?

These resources represent leading-edge offerings from corporate and non-profit organizations

- Model & Inference Services
  - OpenAI token credits
- Cloud Training Environments
  - NIH Cloud Lab
- High-Performance GPU
  - Jetstream2

& more...

Explore the full catalog online to match resources to your specific needs.

nairrpilot.org/pilotresources

NAIRR Pilot | KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM | CPE Higher Education Matters | University of Kentucky

Who are these resources for?

You!

No experience necessary!
No cost to you!

- Researchers and instructors from U.S. 2- or 4-year academic institutions or non-profit organizations *(and private sector if they have a federal research grant)*
- Any, or no, source funding
- Any application domain of AI in line with NAIRR and NSF priorities

Requesting Resources via the National AI Research Resource (NAIRR) Pilot (2025), Lauren Michale, Timothy Middelkoop

NAIRR Pilot | KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM | CPE Higher Education Matters | University of Kentucky

Available Resources

The set of resources available through the NAIRR Pilot has been expanded to include additional resources funded by federal agencies as well as resources contributed by private and non-profit sector partners. Over time, we expect new resources will be added to the NAIRR Pilot, and some resources may be removed from the pilot as their available contributions are committed to projects.

Researchers are strongly encouraged to review the most up-to-date list of resources, which will always be available via the NAIRR Pilot Resource Catalog.

Filters

AI Capabilities

- ☐ Model inference services
- ☐ Model training services (GPU)
- ☐ Model training services (non-GPU)
- ☐ Research Collaboration
- ☐ AI tools and support

Resource Category

- ☐ Federal agency systems
- ☐ Private sector computational resource
- ☐ Private sector model access
- ☐ Other private sector contribution
- ☐ Classroom Platform

Resource Type

- ☐ Cloud
- ☐ GPU Compute

Resources

- Amazon Web Services ⌄
- Anthropic Model API ⌄
- Cerebras Wafer-Scale Engine 2 (CS-2) AI Accelerator ⌄
- Chameleon ⌃

Resource Type:	Cloud Platform
Organization:	University of Chicago
Units:	SUs
User Guide:	Link to User Guide
Features Available:	• Cloud

submit-nairr.xras.org/resources

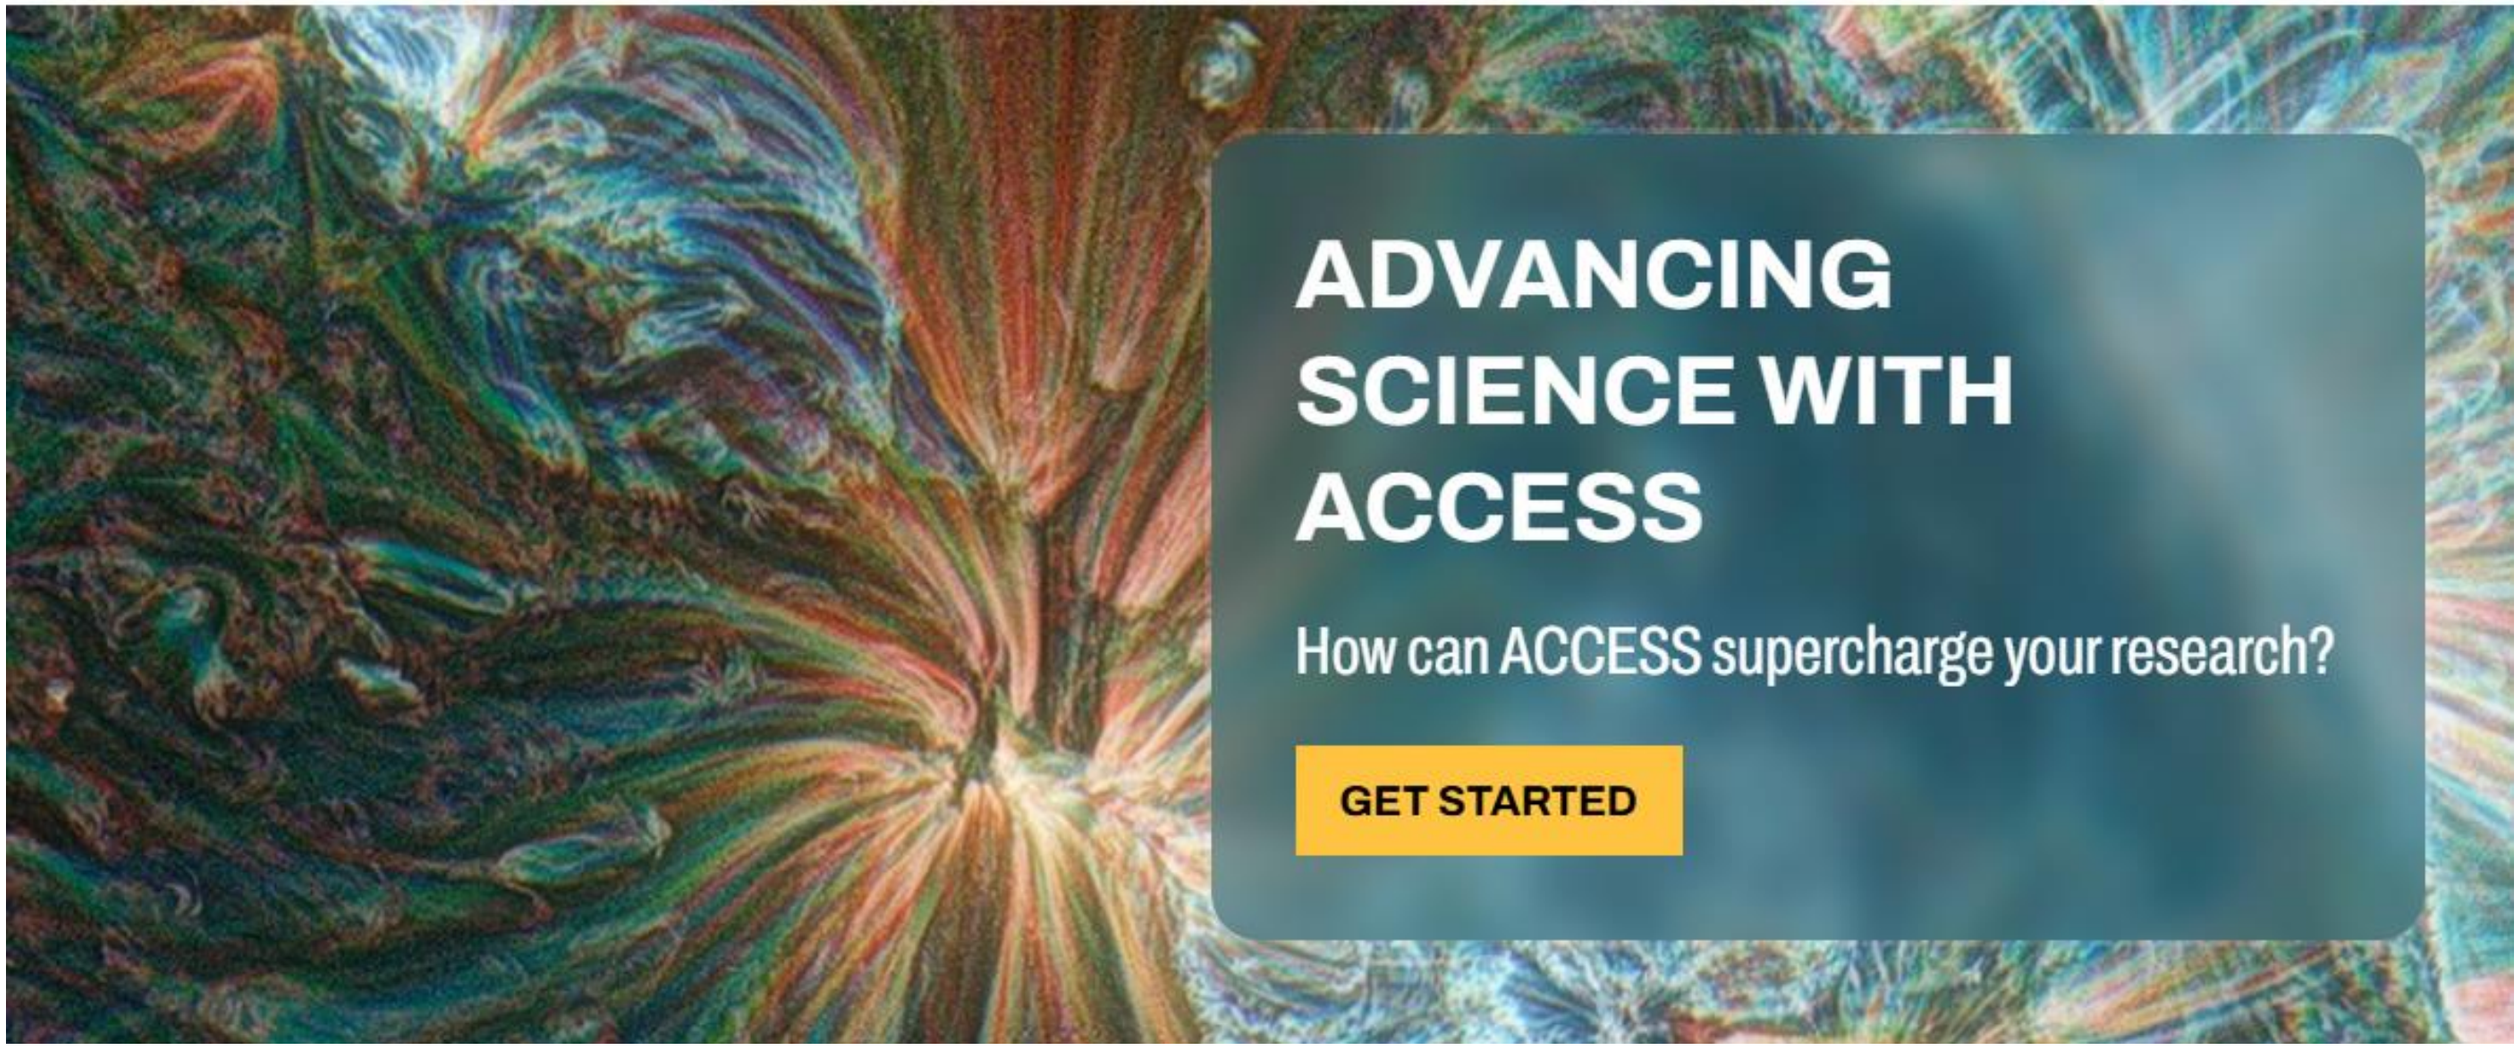
NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

ICPE Higher Education Matters

University of Kentucky

NSF | ACCESS



https://access-ci.org/

NAIRR Pilot | KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM | ICPE Higher Education Matters | University of Kentucky

Are there requirements?

- Research or course must have an AI component.
- Principal Investigator / Lead Instructor and all other users **must perform work via an eligible organization physically based in the U.S.**

nairrpilot.org/help/faq

How to request resources

- Visit https://nairrpilot.org/
- Under "*Current Opportunities*," select
  - Researcher Resources, for research projects
  - Classroom/Educators Resources, for teaching activities
- Full-Scale Requests require a 3-page proposal
  - Align proposal with NAIRR objectives
  - Quantify resource needs
  - Mention you attended this NAIRR Pilot Expansion workshop in your application!

NAIRR pilot focus areas

- Advancing AI methods that enable scientific discovery and improve AI interpretability, security and trust.
- Accelerating time to science and innovation through AI enabled automation, autonomy and novel design and control processes.
- Applying AI to use, share, or integrate sensitive data from multiple sources to enable new experimental methods and discovery.
- Advancing approaches for integrating simulations and AI.
- Creating or developing open-source AI tools, models, datasets, and methods.
- Training and educating the next generation AI-savvy workforce.

Other projects in other areas of AI research and applications, may secondarily be considered for allocation.

nairrpilot.org/opportunities/allocations

NAIRR Community Resources

YouTube
@NAIRRPilot

Recordings of past events & showcases

Youtube.com/@NAIRRPilot

NAIRR Pilot

Registration for NAIRR Pilot Office Hours

nairrpilot.org/help/office-hours

GitHub
NAIRRProgram

Hosts the community-learning-resources repository

github.com/NAIRRProgram

Lunch

11:30PM – 12:15PM

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

Workshop Webpage:

Event Survey:

ai_workshop@uky.edu

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Early or Late

How AI Shapes What Happens Next in Eastern Kentucky

Danna Barnett
Director of Entrepreneurship
Shaping Our Appalachian Region (SOAR)

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

The gap shows up fast

Early Adopters
Moving Ahead

The Gap

Cautious Middle
Hesitating

Late Adopters
Falling Behind

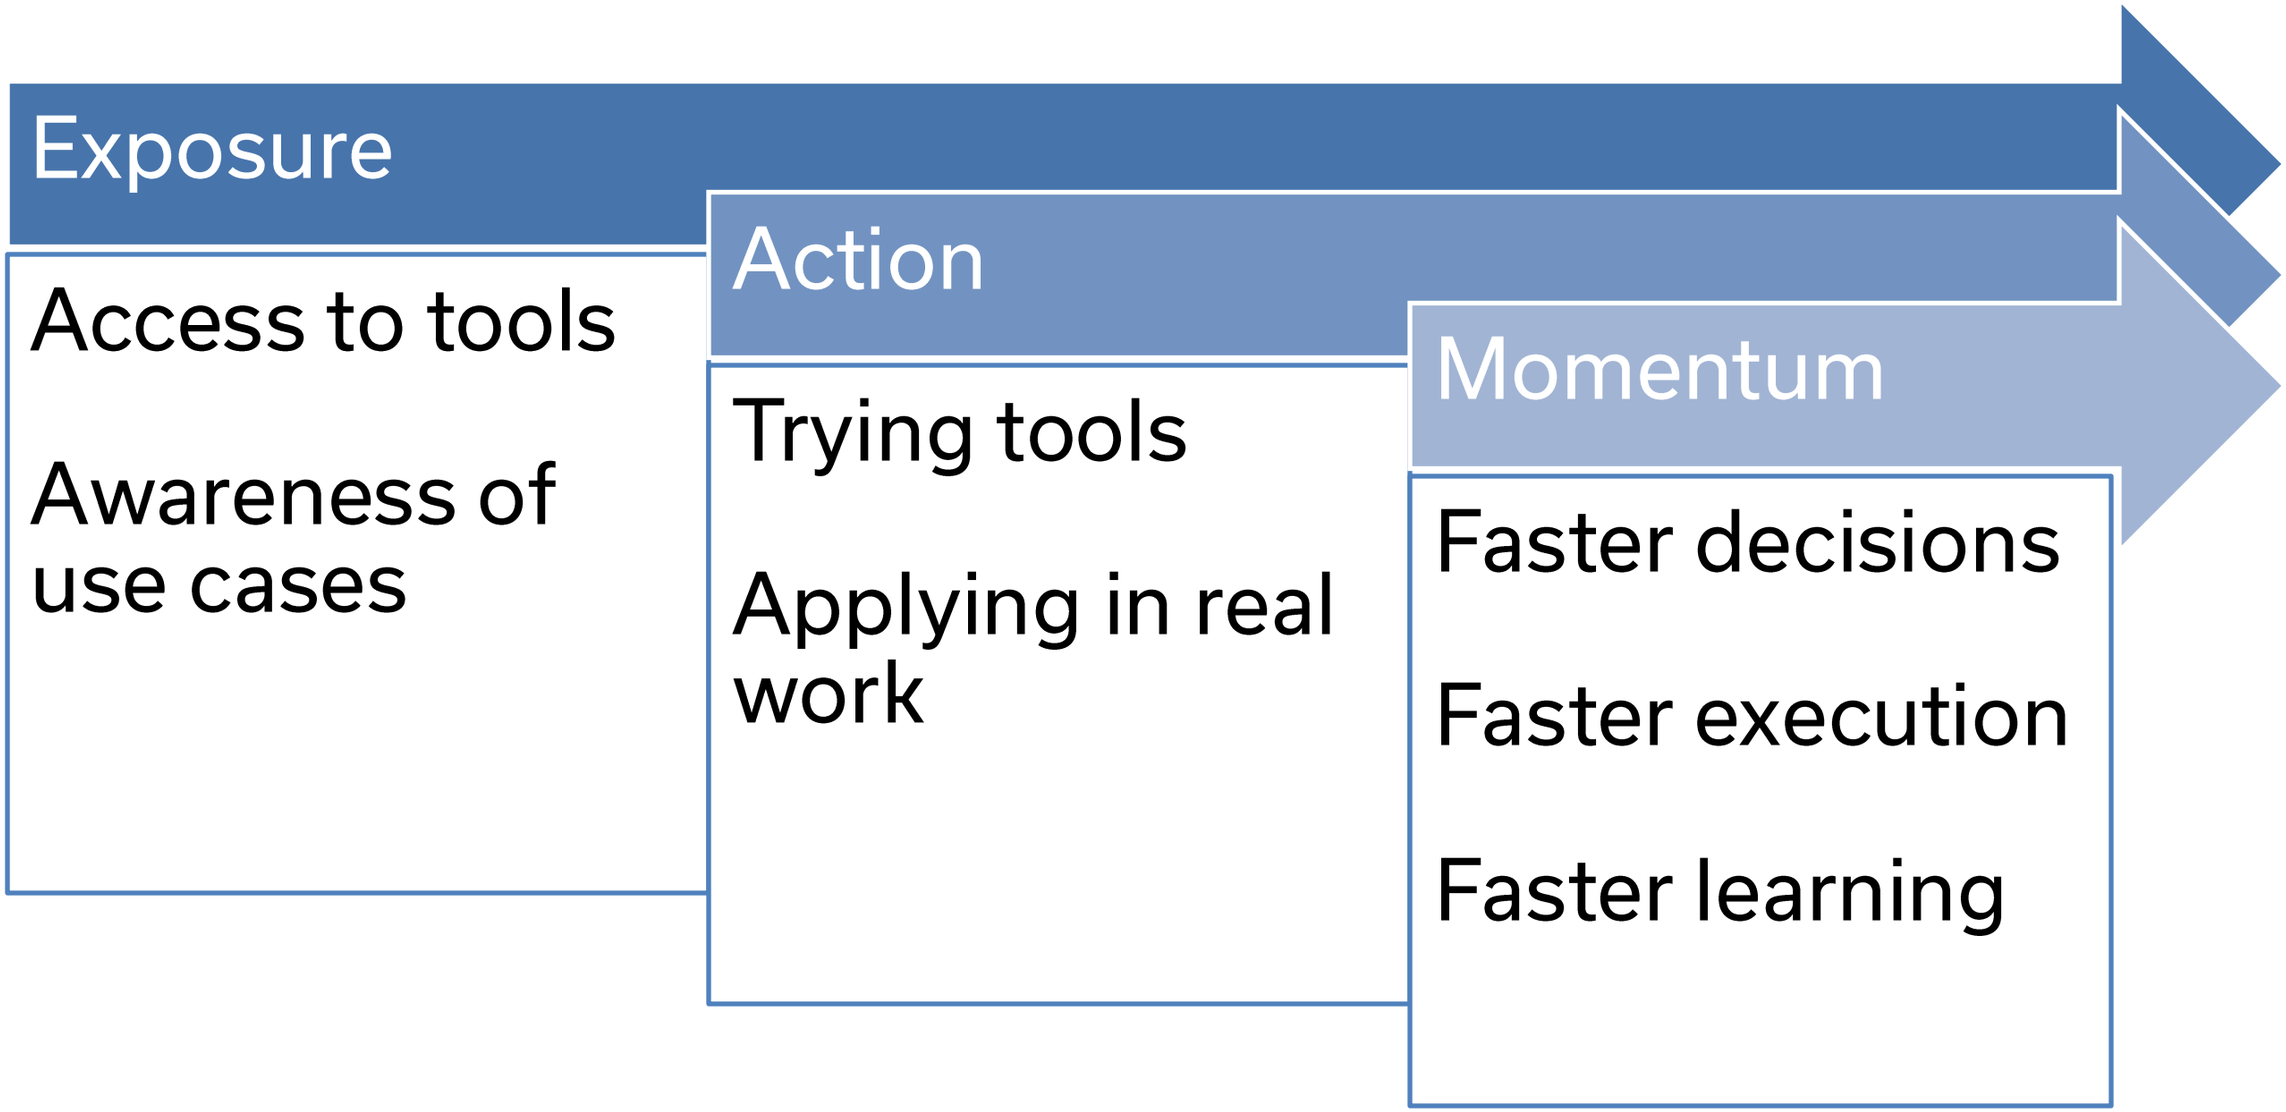
NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Same People. Different Speed.



Education + Workforce + Business

AI Changes How Work Gets Done

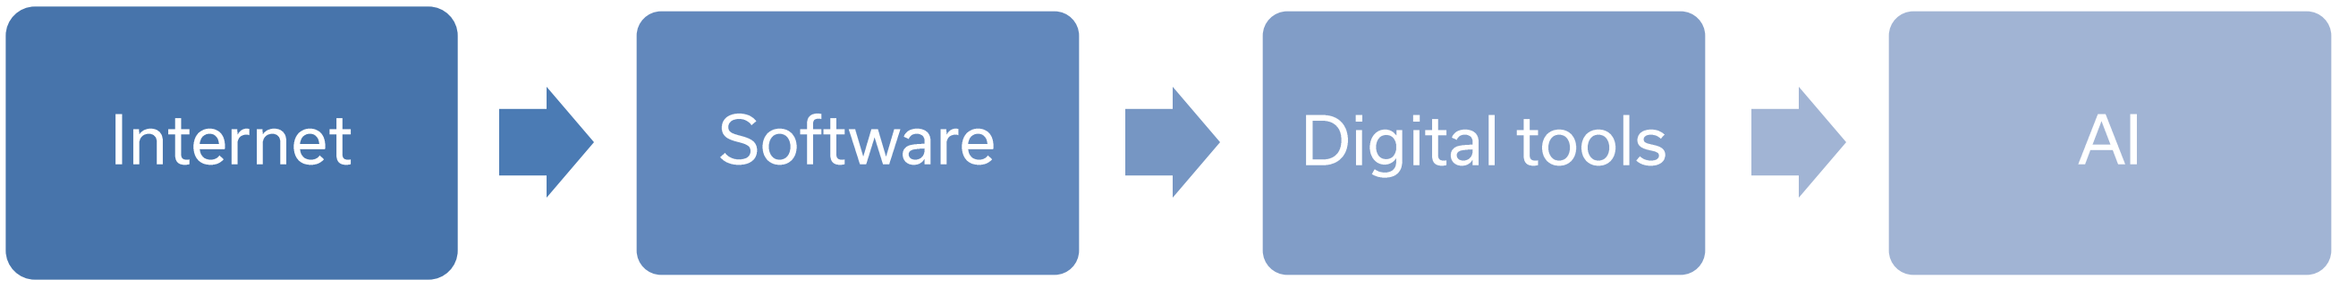
- Students can learn, test ideas, and build skills faster

- Faculty can use AI to support preparation and feedback while keeping their expertise at the center

- Businesses can operate faster and with fewer resources

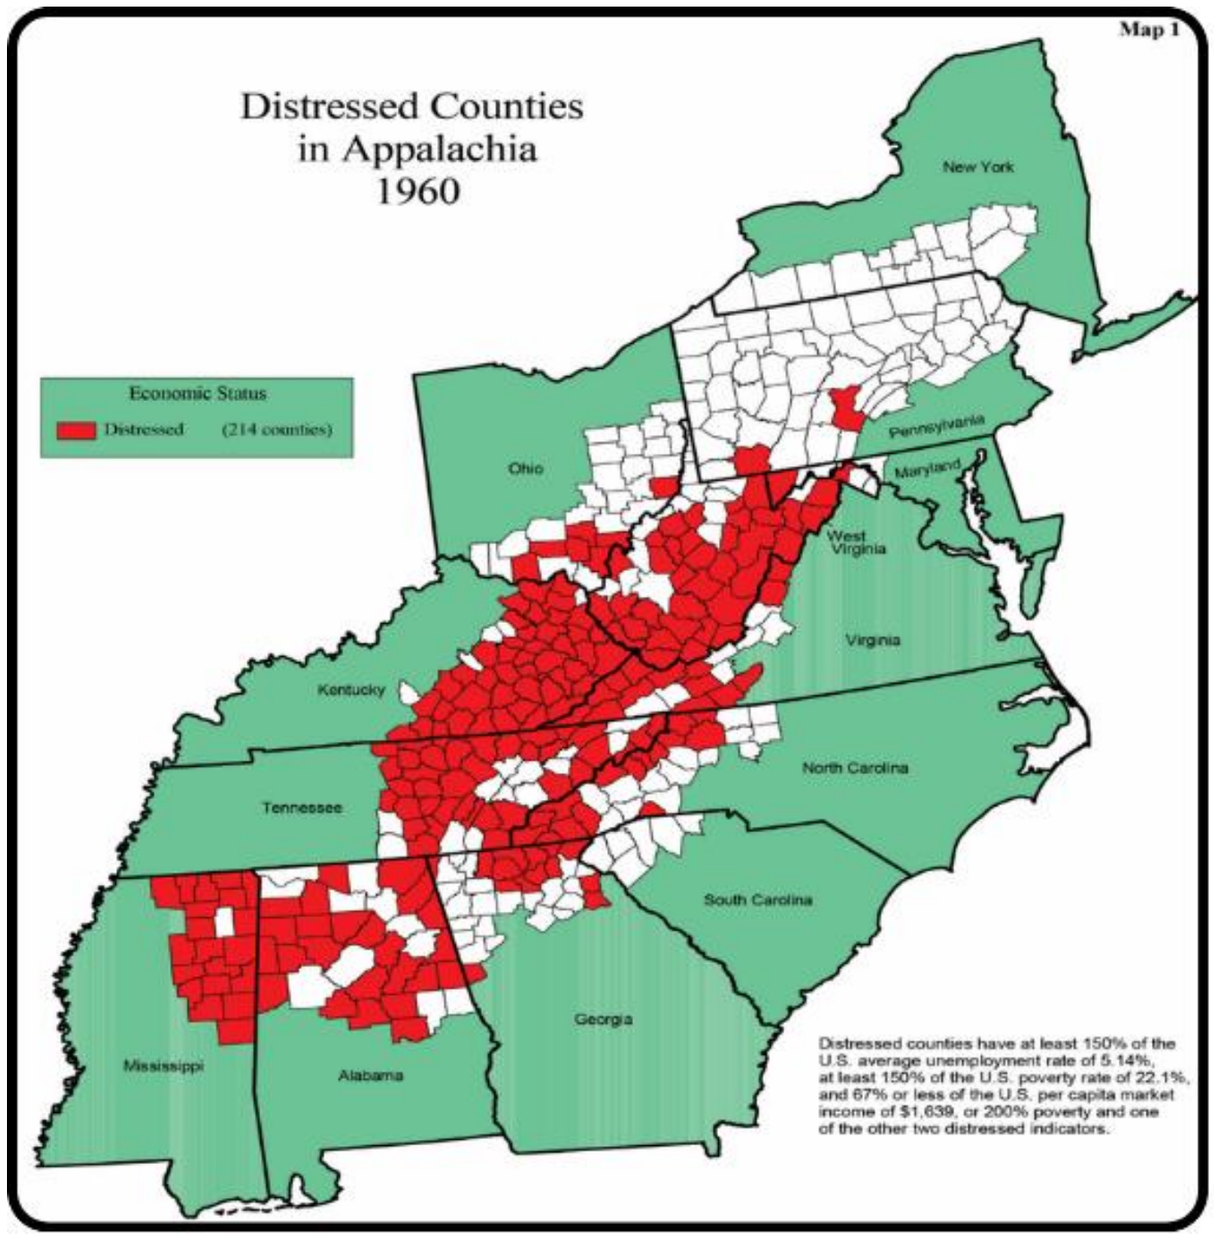
- Workers can produce more without increasing hours

The Technology is New
The Pattern is Not

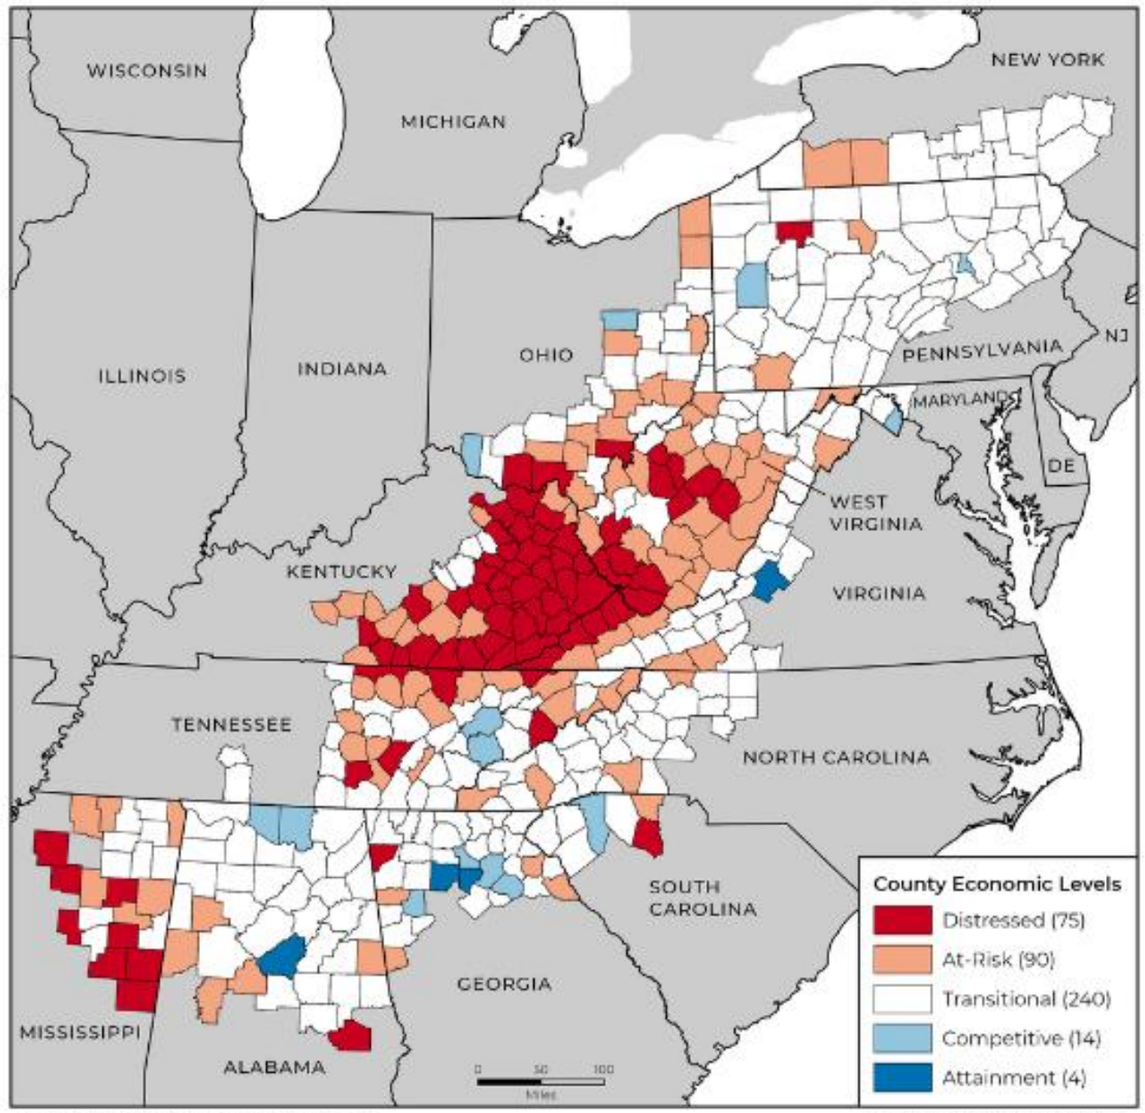


Every major technology shift rewards
the people who engage early.

1960



2025



Appalachian Regional Commission, May 2025
Data Sources:
Unemployment data: U.S. Bureau of Labor Statistics, LAUS, 2021–2023
Income data: U.S. Bureau of Economic Analysis, LAPI, 2023
Poverty data: U.S. Census Bureau, ACS, 2019–2023

Fiscal Year 2026 Designations:
Effective October 1, 2025
through September 30, 2026

Industries changed faster than systems could adapt

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

ICPE Higher Education Matters

University of Kentucky

SOAR Innovation: Entrepreneurship Pillars

Student + Community	Social Enterprise	Scalable Startups
• Education • Inspiration • Ideation • Action	• Impact • Workforce Barriers • Solutions	• Innovation • Tech-Enabled • Investment

By empowering students, tackling barriers, and scaling innovation, we build a stronger, future-ready region.

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

From Innovation to Application

- Innovation only matters if people can use it
- Connects emerging AI tools to real-world use
- Helps translate complex technology into practical application
- Supports communities, businesses, and educators in adoption



Dr. Karen Coburn
AI Technology Transfer Manager, SOAR

NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
CPE Higher Education Matters
University of Kentucky

SOAR: Helping Eastern Kentucky Move Early with AI



- AI integrated into accelerator programs
- AI Startup Sprints
  - For Businesses
  - For Students
- Community-based AI training

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Thank you

Danna Barnett
Director of Entrepreneurship
P | 859.397.3266
E | danna@soar-ky.org



SOAR | INNOVATION
SHAPING OUR APPALACHIAN REGION

POWERED BY TEAM KENTUCKY CABINET FOR ECONOMIC DEVELOPMENT
KY Innovation

ARC Appalachian Regional Commission

Stay Connected with SOAR Events

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Spotlight: Bit Source

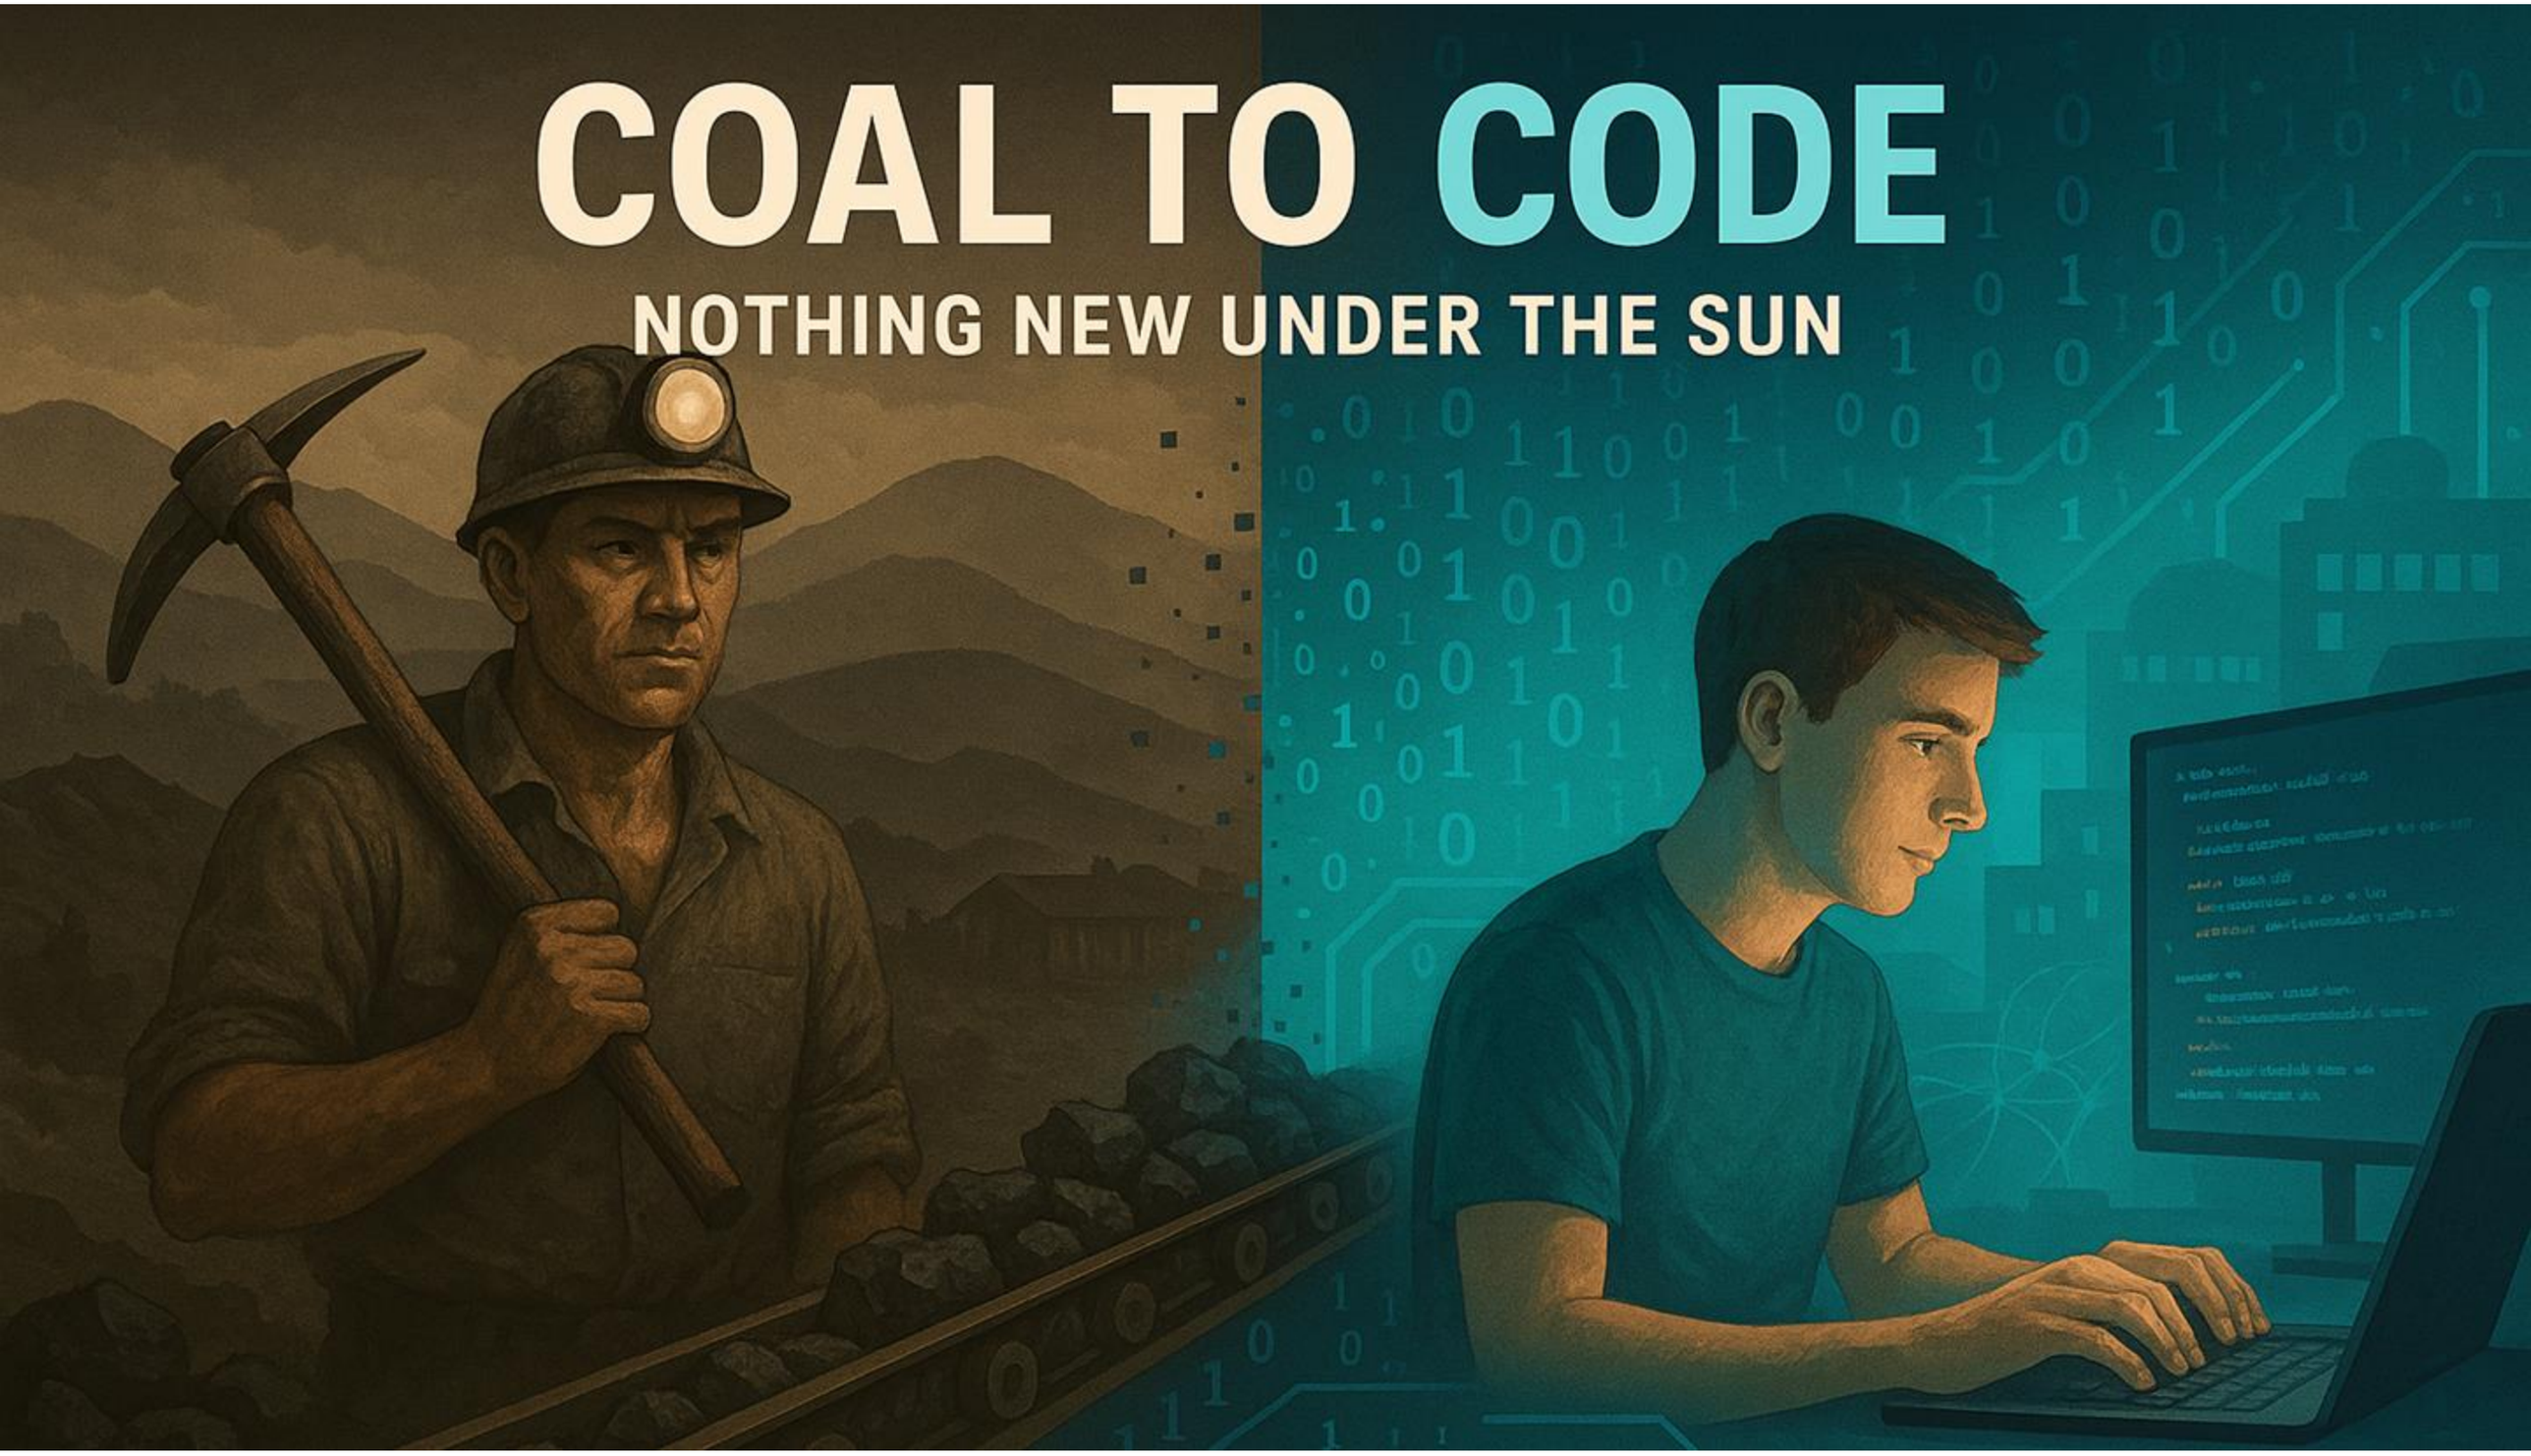
Rusty Justice
Co-Founder
Bit Source (bitsourceky.com)

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky


COAL TO CODE
NOTHING NEW UNDER THE SUN

AI Initiatives Timeline: 2017-2026

NSF

2017 — CMU Mark Stehlik & Stephen Tsou

2017 — ML+UPIKE - Concrete Inspection Model

2021 — 2018-19 - MIT REAP Kentucky, USA

2022

2023 — ML Vesuvius Challenge - UK Brent Seals

2023 — Ai Aizver (now on 5 library kiosks / 2026)

2024 — Ai Kentucky Coal Journal

2025 — County Waste Systems - Google Cloud Vision Ai - 80,000 Pages - Searchability (1 existing 3 new incoming)

3/2026 — Hosted Wisdom Works Event & Ai Roundtable KCTCS (Dr. Ryan Hall and Darrell Maynard)

2026 — Municipal - Utilities Optimization

2026 — Medical Record Digitization and EHR Import

CMU Mark Stehlik & Stephen Tsou

MIT "Ai and the Future of Work" Summit CSAIL

ML + UPIKE - Concrete Inspection Model

AML - Kentucky Coal Heritage Trail - Voice Ai

ML Vesuvius Challenge - UK Brent Seals

Ai Aizver (now on 5 library kiosks / 2026)

Ai Kentucky Coal Journal

County Waste Systems - Google Cloud Vision Ai - 80,000 Pages - Searchability (1 existing 3 new incoming)

Hosted Wisdom Works Event & Ai Roundtable KCTCS (Dr. Ryan Hall and

Municipal - Utilities Optimization

Medical Record Digitization and EHR Import

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Mark Stehlik



Carnegie Mellon University

Presidential Advisor for Science and Technology Policy



Maria T. Zuber

Office Phone 617-324-8177

Room 3-234







MIT REAP
KENTUCKY

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Filter by Label: All

Marker Details

ID: 5
URL: crack_5.jpg
Label: A
Length: 3ft

ID: 6
URL: crack_6.jpg
Label: B
Length: 3ft

ID: 7
URL: crack_7.jpg
Label: C
Length: 3ft

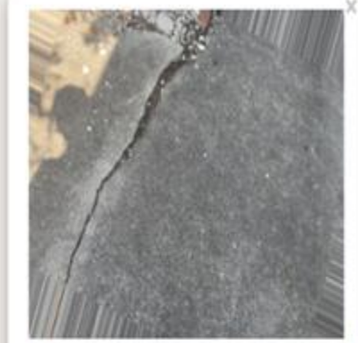
ID: 8
URL: crack_8.jpg
Label: D
Length: 3ft

ID: 9
URL: crack_9.jpg
Label: A
Length: 3ft

ID: 10
URL: crack_10.jpg
Label: B
Length: 3ft

Total Markers: 65

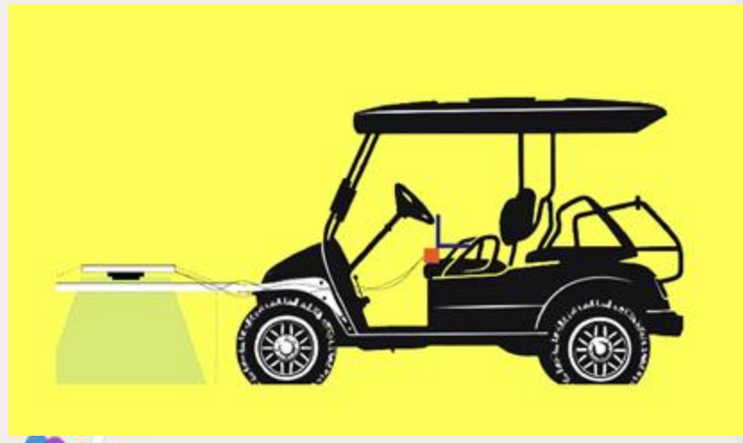
Created For:
City of Pikeville

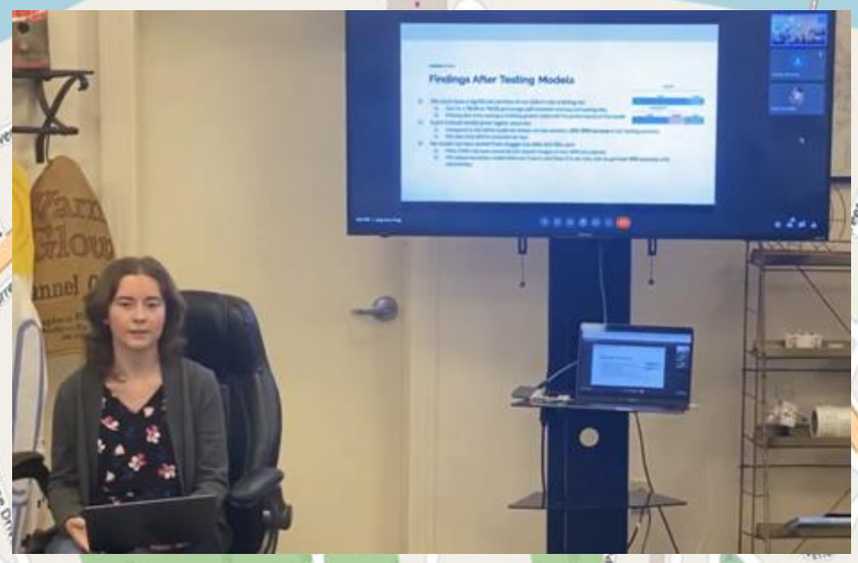


Label: B

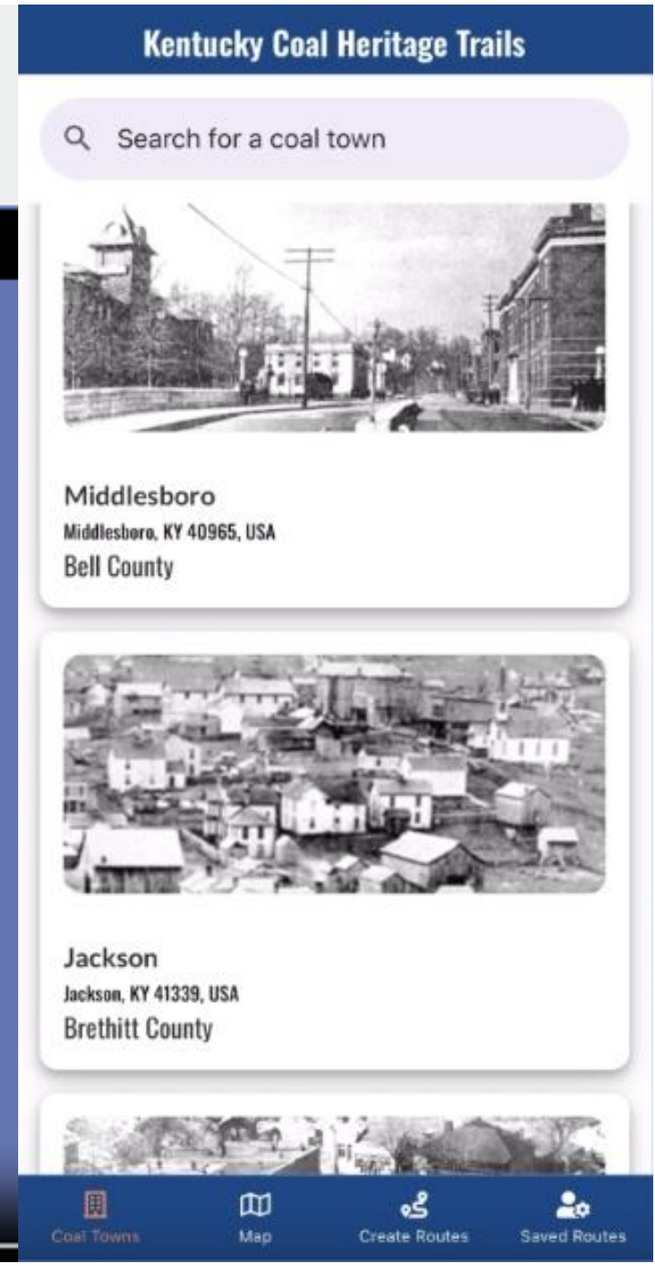
Length: 3ft

ML Neural Network







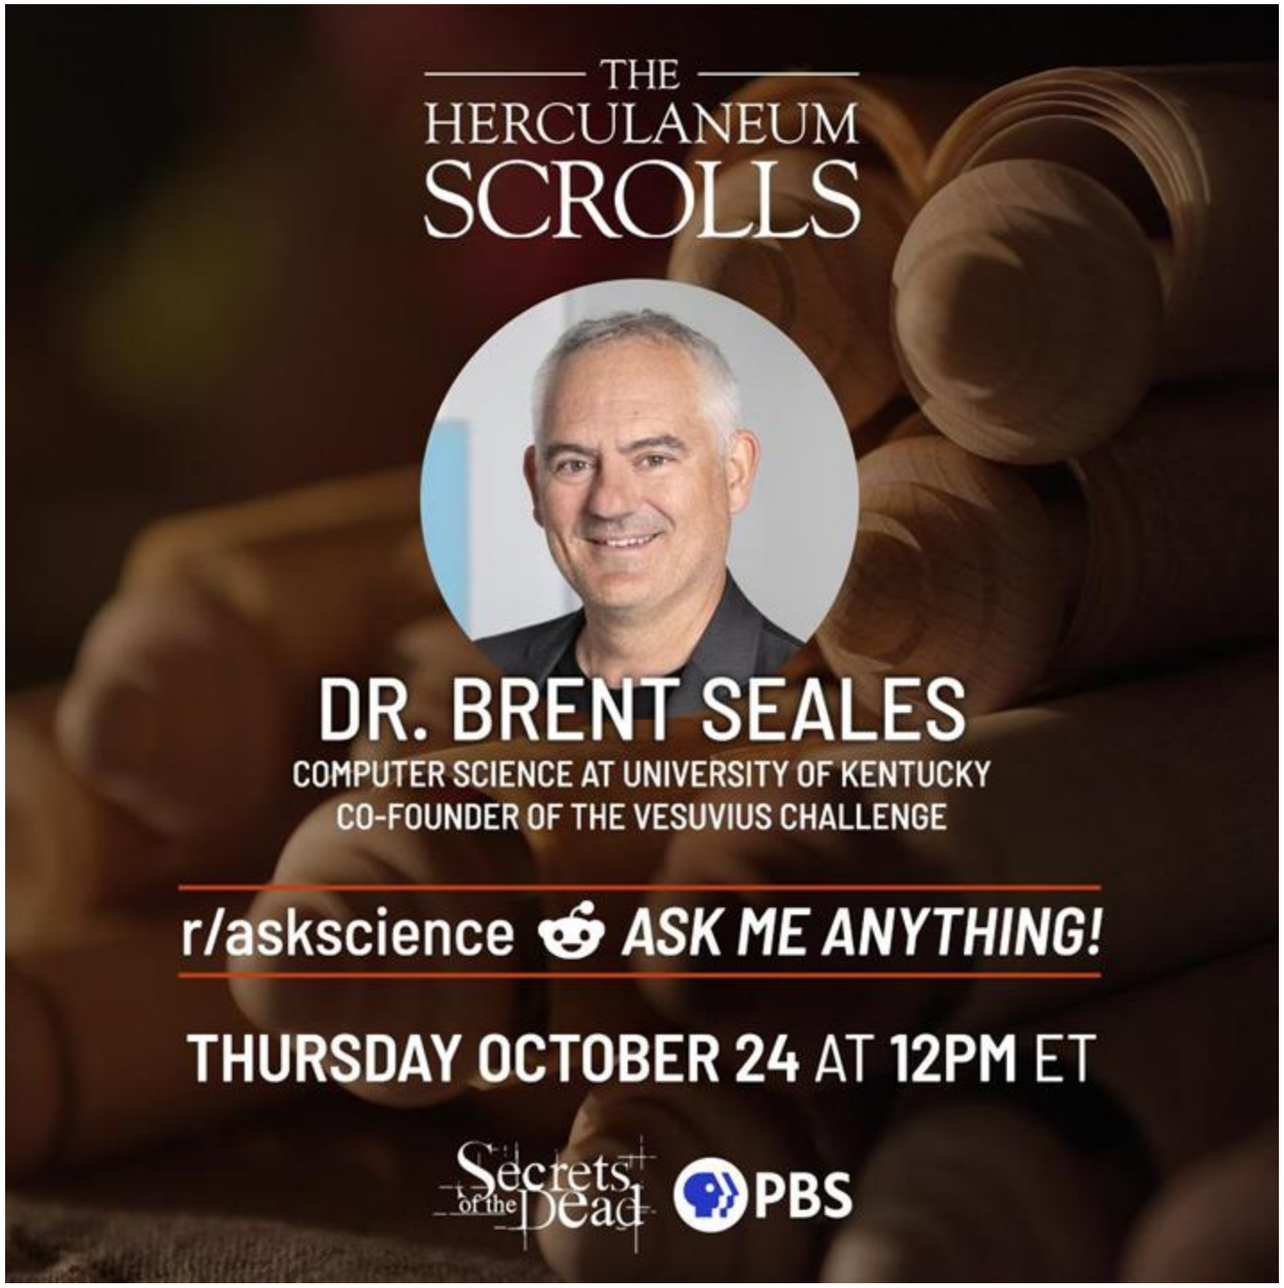


NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

ICPE Higher Education Matters

University of Kentucky



NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
KCPE Higher Education Matters
University of Kentucky







NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

VESUVIUS CHALLENGE · FEATURED CODE COMPETITION · 3 YEARS AGO

Late Submission

Vesuvius Challenge - Ink Detection

Resurrect an ancient library from the ashes of a volcano



Overview Data Code Models Discussion Leaderboard Rules

Leaderboard

Raw Data Refresh

bit source

Public Private

The private leaderboard is calculated with approximately 50% of the test data.
This competition has completed. This leaderboard reflects the final standings.

Prize Winners

#	△	Team	Members	Score	Entries	Last	Solution
342	▲ 64	Bit Source		0.568229	10	3y	

DOWNLOAD USING KAGGLE CLI

```
>_ kaggle competitions leaderboard vesuvius-challenge-ink-detection -d
```

Machine Learning

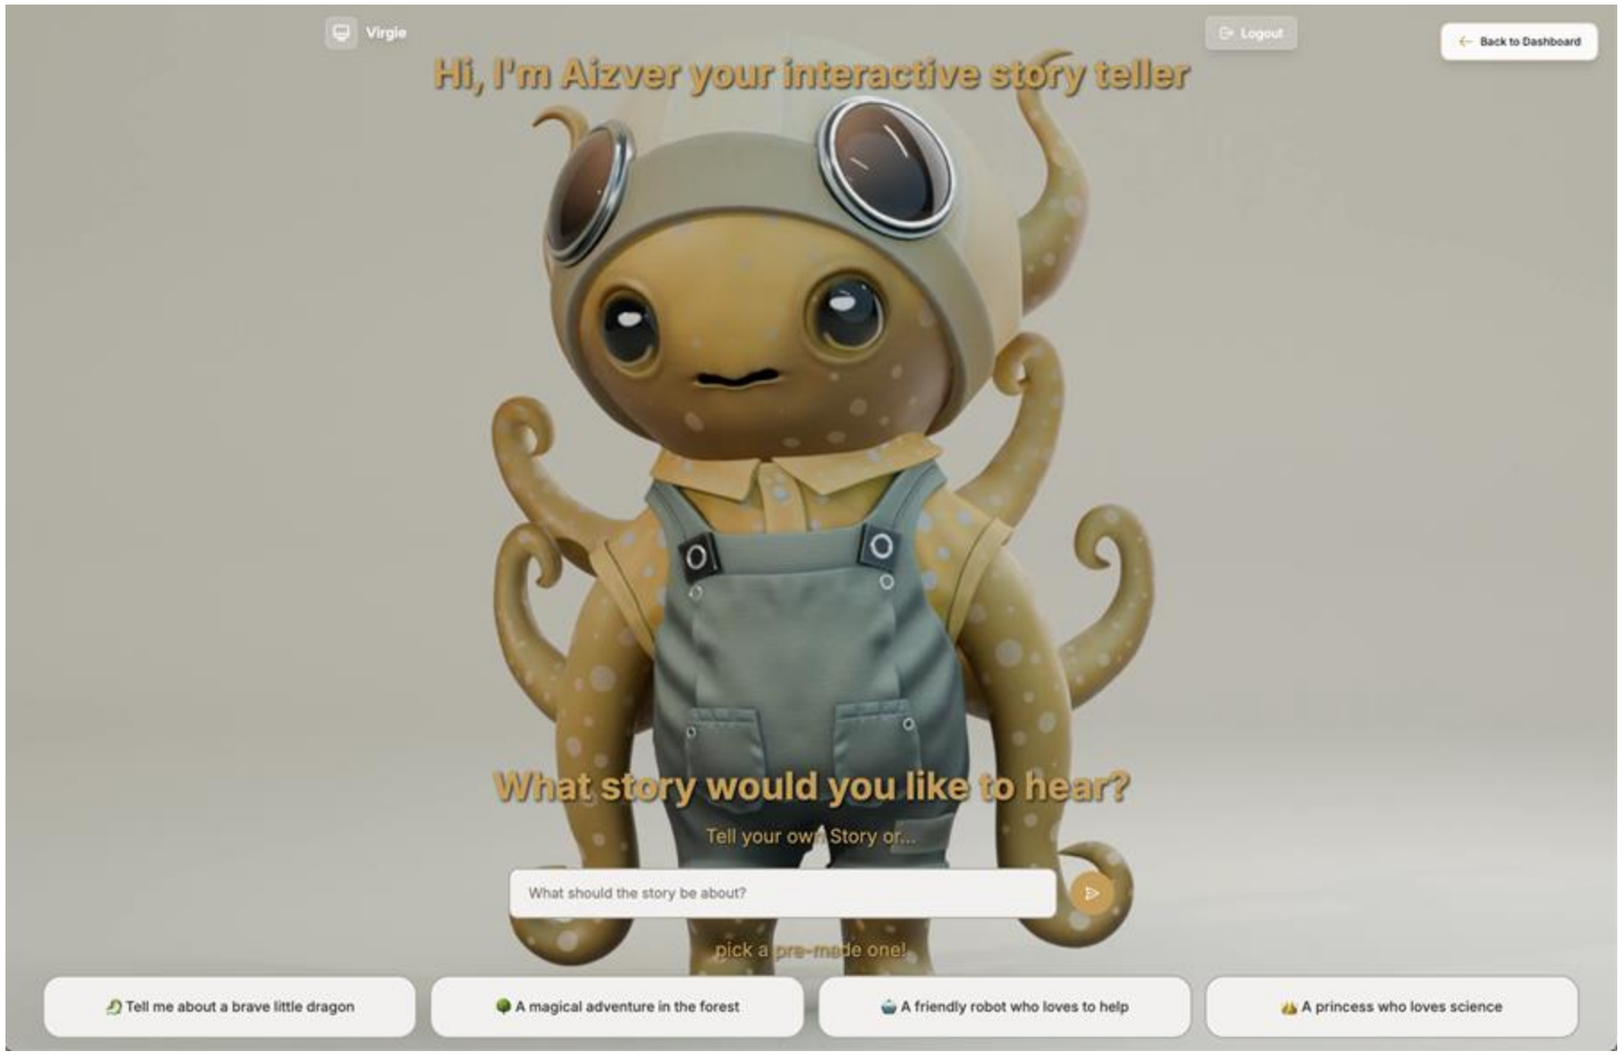
NAIRR Pilot

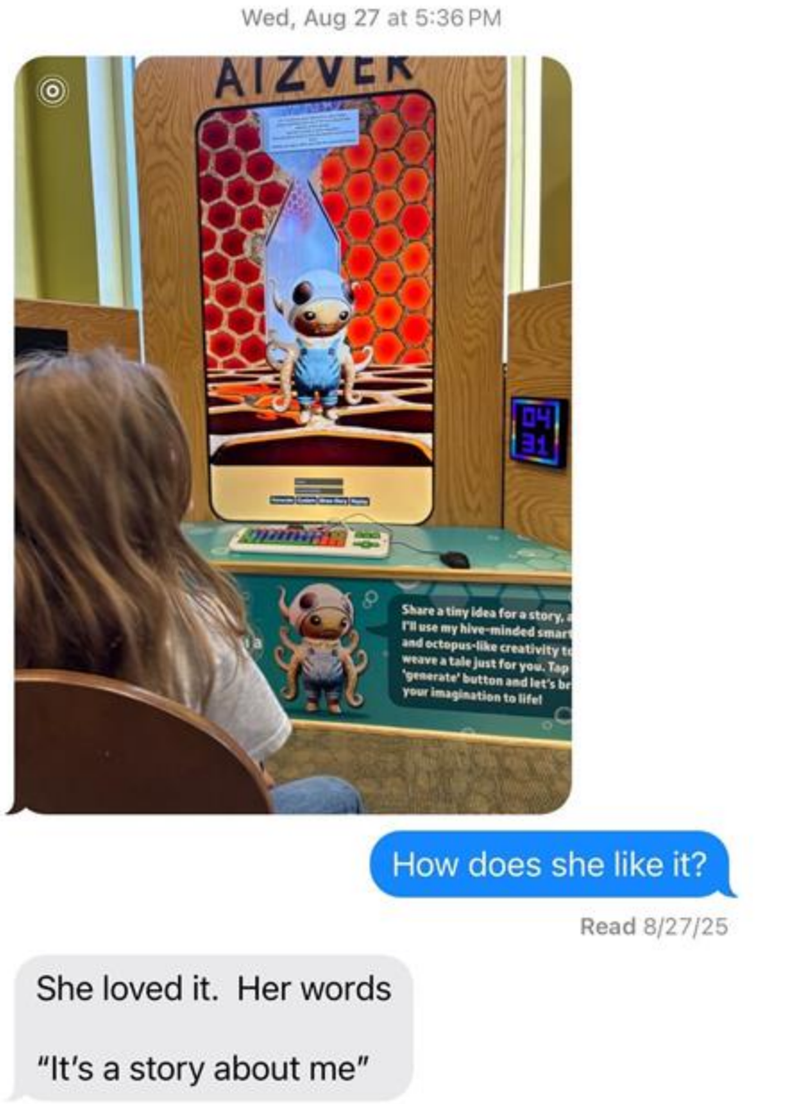
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

ICPE Higher Education Matters

University of Kentucky

Open ai - Agent, LLM, Voice

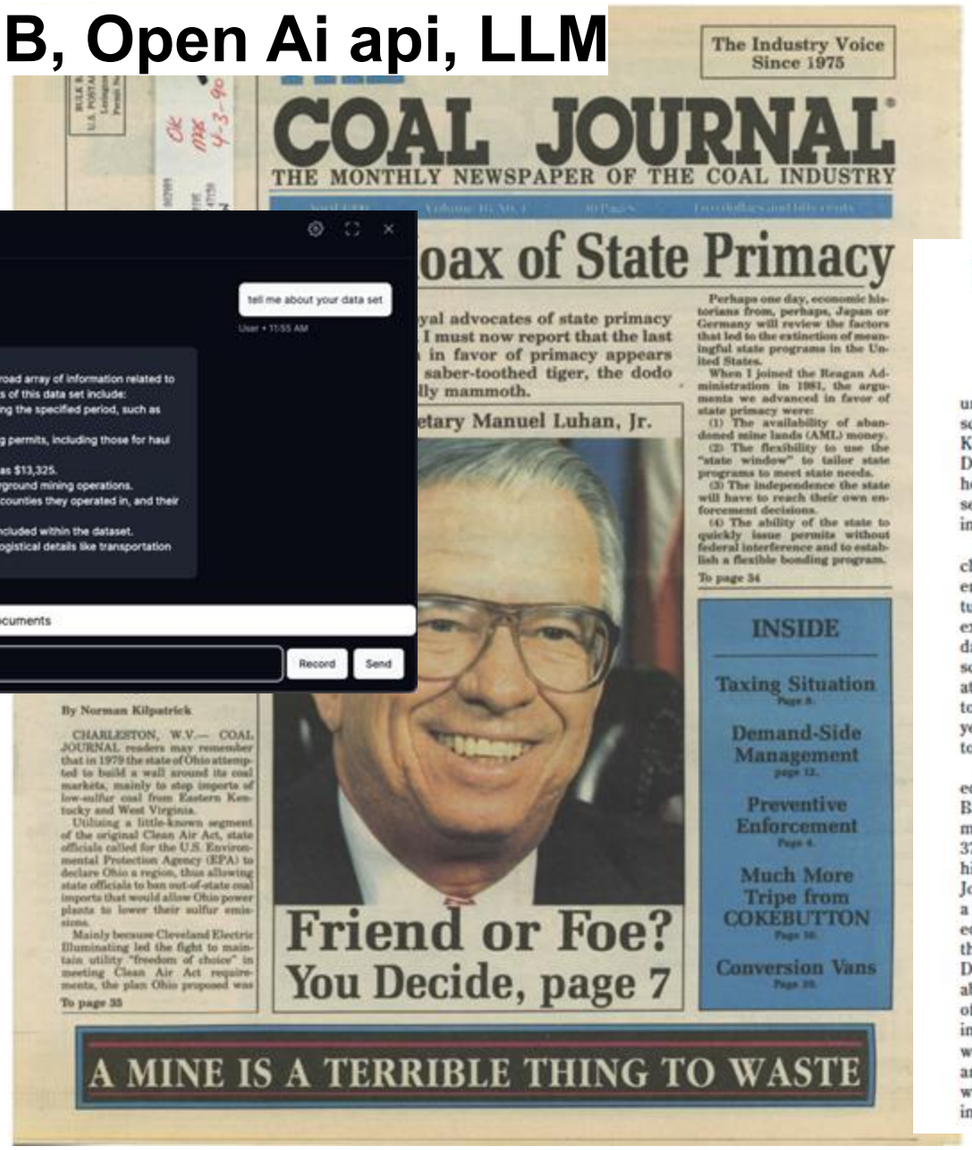




NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
ICPE Higher Education Matters
University of Kentucky

Vector DB, Open Ai api, LLM







NEWS-EXPRESS
ANE Weekend
Weekend Edition
Formerly The Pike County News and The Appalachian Express • www.news-expressky.com
Pikeville, KY Volume 108, Number 98 • 26 Pages
December 10-12, 2021
$1.00; MAIL DELIVERY 38 CENTS



World-famous UK professor shares historic project with local students

By Nicole Ziege
@NicoleZiege

A world-renowned pioneer of unlocking historic artifacts and scrolls spoke to more than 300 East Kentucky high school students on Dec. 2-3 and shared more about how he and his team unlocked the secrets of the ancient world by using cutting-edge technology.

Dr. Brent Seales, professor and chair of the Department of Computer Science at the University of Kentucky, is considered the foremost expert in the digital restoration of damaged and unreadable manuscripts, and he has worked to create and use high-tech, non-invasive tools to rescue hidden texts for 20 years, in an effort to restore them to humanity.

On Dec. 2-3, Dr. Seales visited Allen Baptist Church and First Baptist Church of Pikeville to share more about his research with about 370 high school students from 10 high schools in Pike, Floyd and Johnson counties. He also met with a group of professional engineers, educators and business people at the Appalachian Wireless Arena on Dec. 2, in an effort to share more about his research and the value of relics, science and technology in preserving history. The events were coordinated by CEDAR, Inc. and Bit Source, LLC, in partnership with the UK College of Engineering.

CEDAR, Inc. President and Executive Director John Justice expressed the importance of having Dr. Seales share his research with Eastern Kentucky students.

"This is the highest level of computer science that you can get. He's at the very top of his field," Justice said. "He's the one who created, discovered, invented the process to read what has never been able to be read before with ancient artifacts. It's important for our kids to be exposed to this so they can see what's possible."

Dr. Seales and his team created cutting-edge computer science and visualization technology that allowed them to "virtually unwrap" a scroll that had previously been considered unreadable for about 40 years.

The scroll, which was found by archaeologists in En-Gedi, Israel, in 1970 and remained in storage for 40 years, was badly burned, and it could not be physically unwrapped. In 2015, Dr. Seales and his team virtually unwrapped the ancient scroll, and they discovered that the artifact is the oldest copy of the beginning of the Book of Leviticus, believed to have been written between the second and third centuries A.D., according to UK.

During his presentation, Dr. Seales said that the ability to read unreadable documents was a combination of digitization, spectral imagining, tomography (a technique for displaying a solid object through X-rays), and virtual unwrapping. Through these scientific techniques, they were able to create a complete image of the unreadable scroll, create a virtual simulation of the scroll and slice and unwrap the virtual simulation into several pieces, which can then be lined up to read what is inside the document.

"What we developed is an entire set of steps to move from the scanning using tomography, to the reading using virtual unwrapping," Dr. Seales said.

He compared the process to that of a baker baking bread. If a baker puts a Pi (3.14) symbol with red dough on top of regular bread dough, and rolls up the dough so that the red symbol is rolled up inside, then the red symbol becomes unreadable. This represents the contents within the unreadable scroll. Then, the baker puts the dough into the oven, which represents the eventual decay of the parchment from time and natural elements.

Once baked, a baker cannot see the red symbol unless the bread is broken apart. Therefore, someone

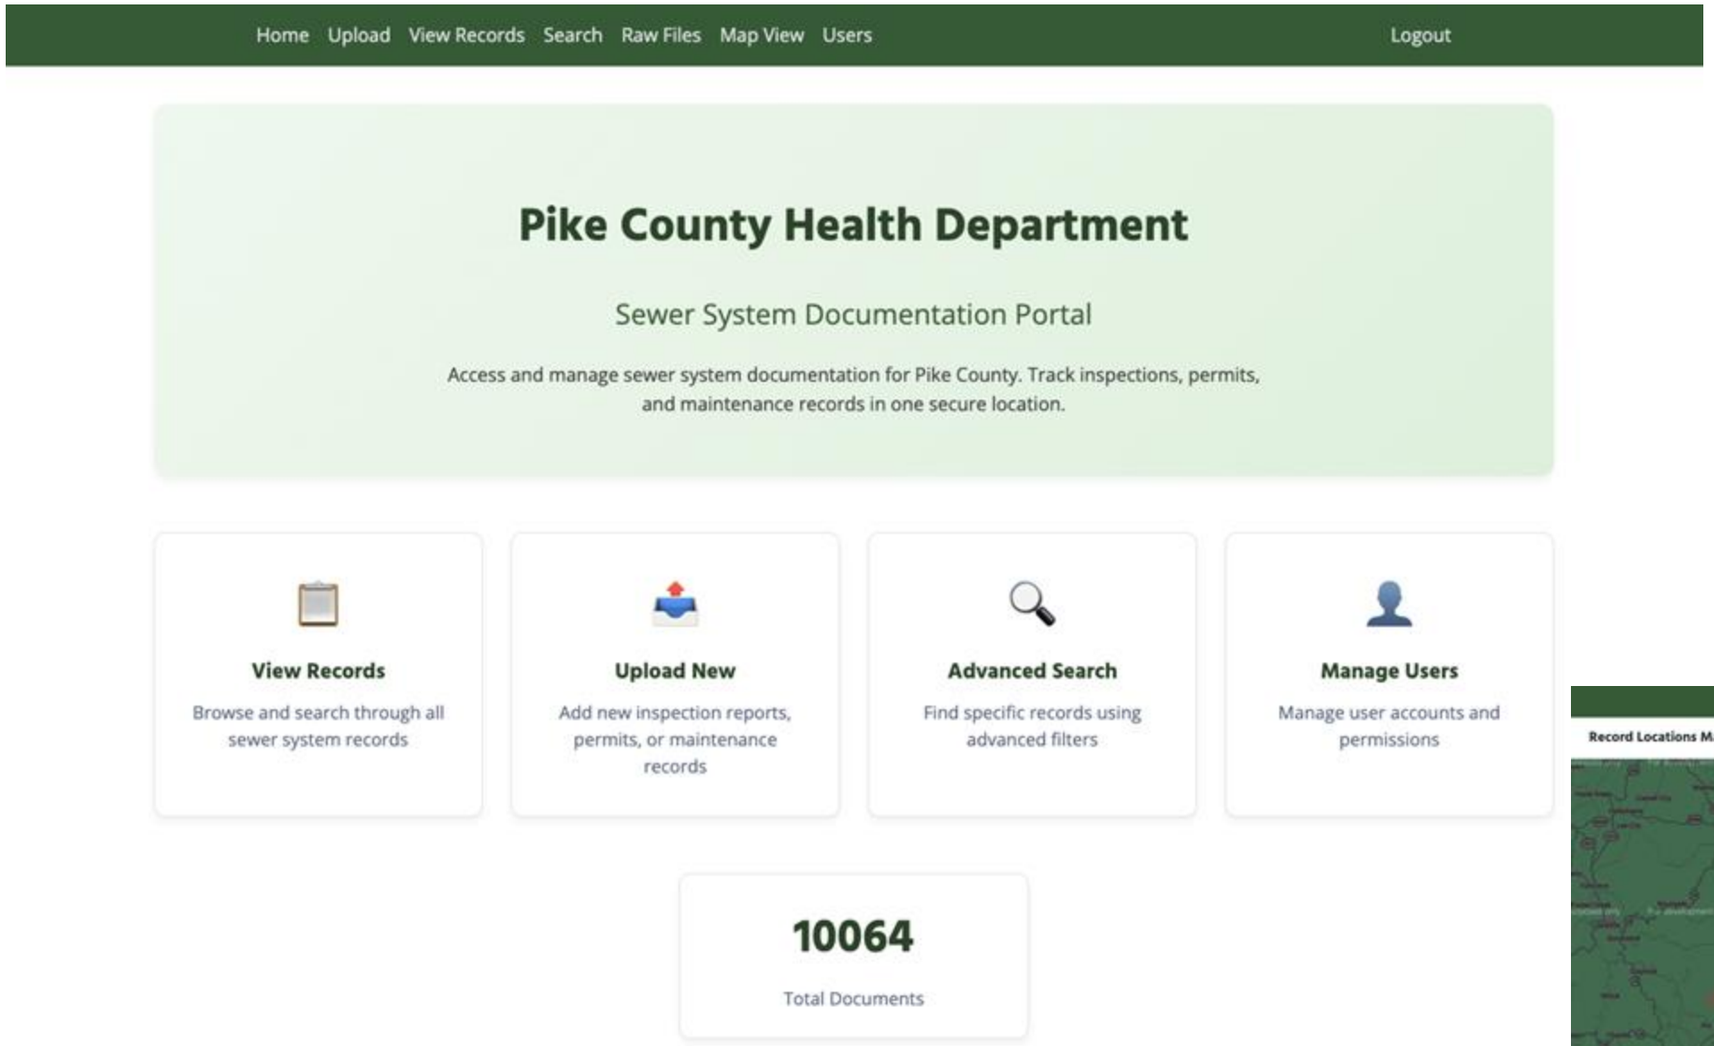
See SEALES, Page 7A

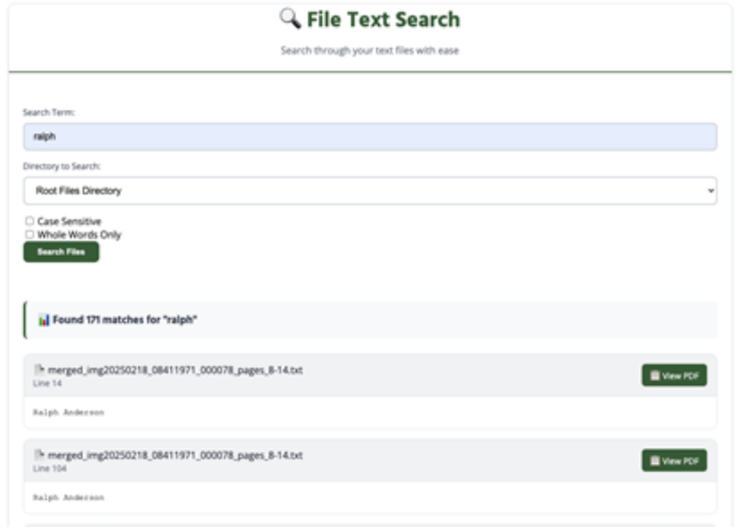


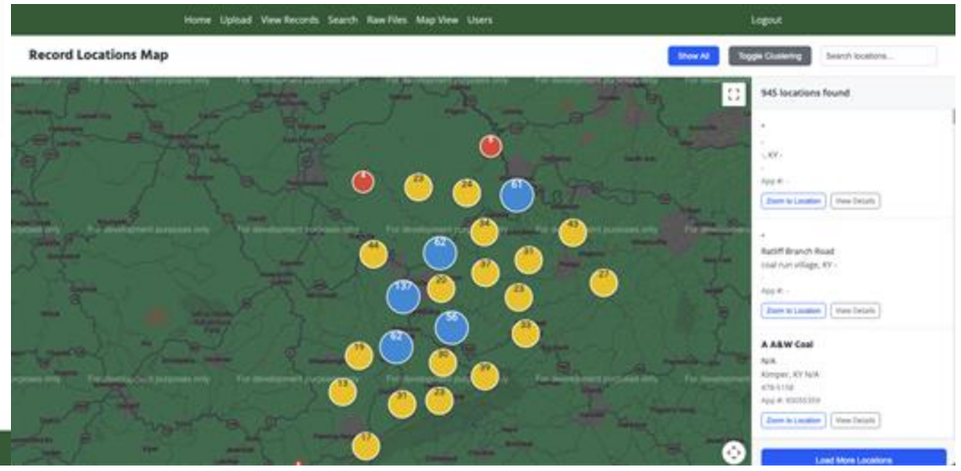
News-Express photo by Nicole Ziege

Dr. Brent Seales, professor and chair of the Department of Computer Science at the University of Kentucky, presented to more than 300 students from Eastern Kentucky on Dec. 2-3 about how he and his team used cutting-edge technology to read an unreadable ancient scroll.

NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
CPE Higher Education Matters
University of Kentucky







Google Cloud Vision Ai

NAIRR Pilot

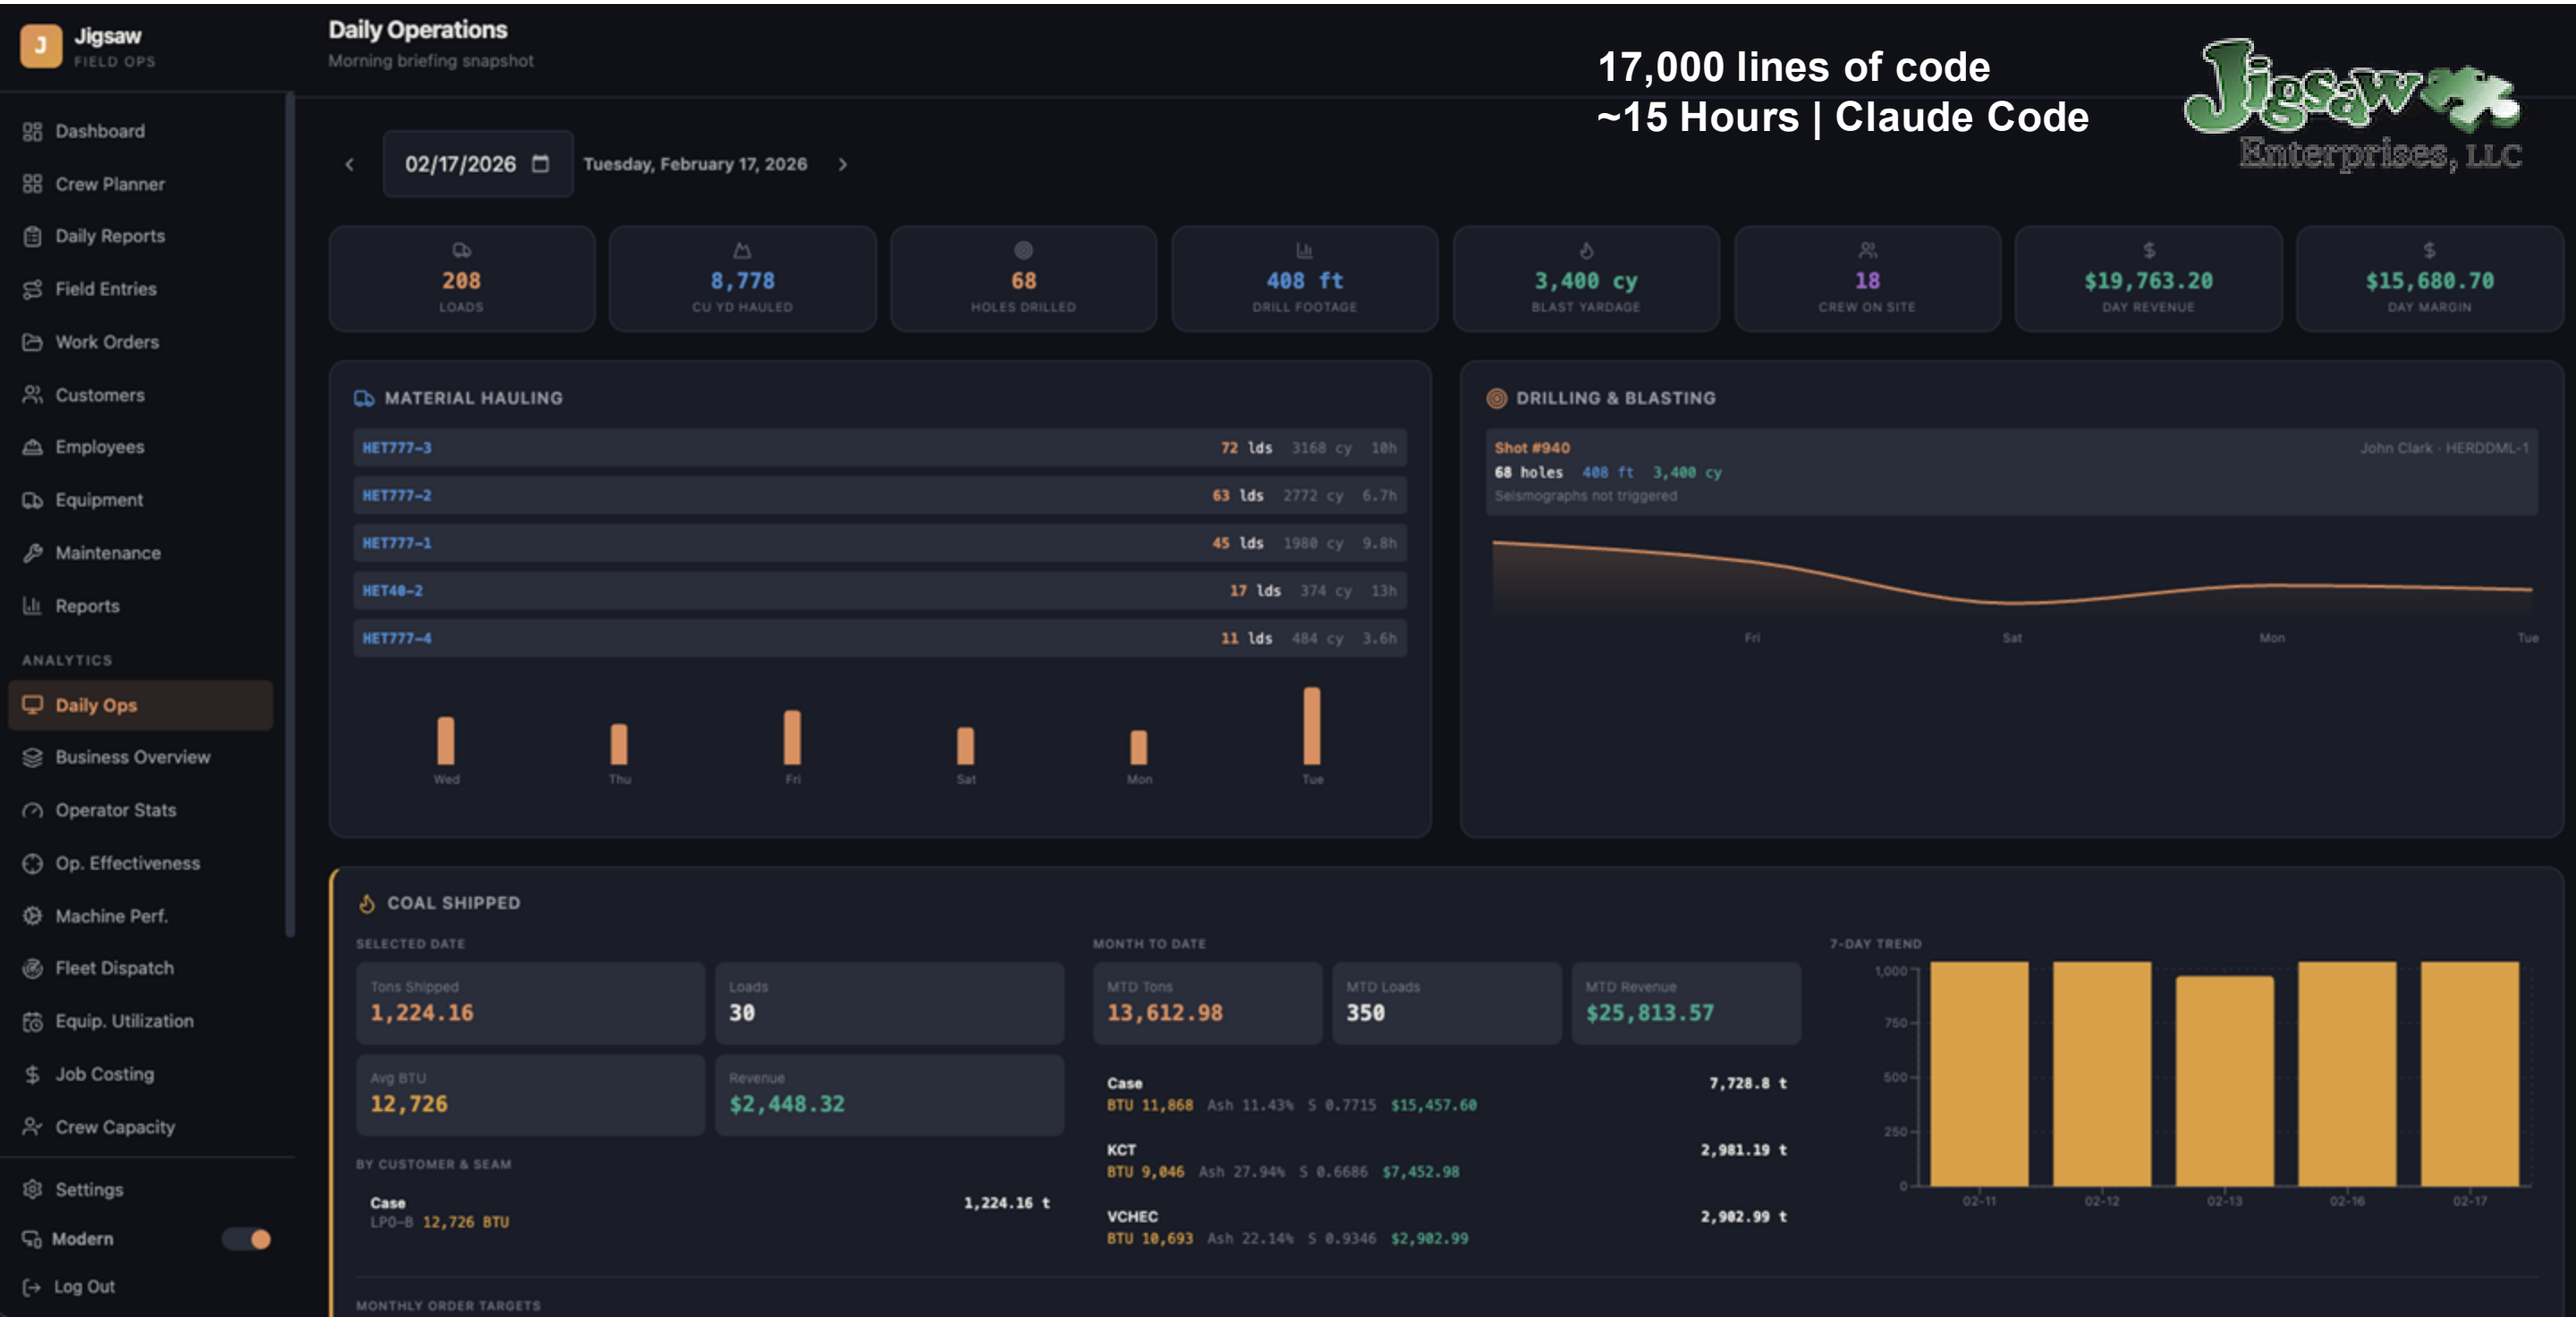
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

17,000 lines of code
~15 Hours | Claude Code

Jigsaw Enterprises, LLC



NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

KCPE Higher Education Matters

University of Kentucky

MOUNTAIN TOP MEDIA

HOME | NEWS | SPORTS | EVENT CALENDAR | WEATHER | OUR TEAM | DOWNLOAD OR WATCH!

BitSource to host “Wisdom Works”

ROBIN TILLMAN | Article Updated: March 10, 2026 | LEAVE A COMMENT

Share on Facebook | Tweet this! | in

Think tank to tackle private sector business issues in Pikeville



Roundtable Discussion | Ryan Hall (Ai Lead KCTCS) + Darrell Maynard (Eastern Telephone & Technologies)





Words Matter: Why We Should Be Careful How We Talk About AI

Payton May
Partner and COO at Bit Source

June 10, 2025

NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
ICPE Higher Education Matters
University of Kentucky

Generative AI & Vibe Coding

Vaiden Logan, BS

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

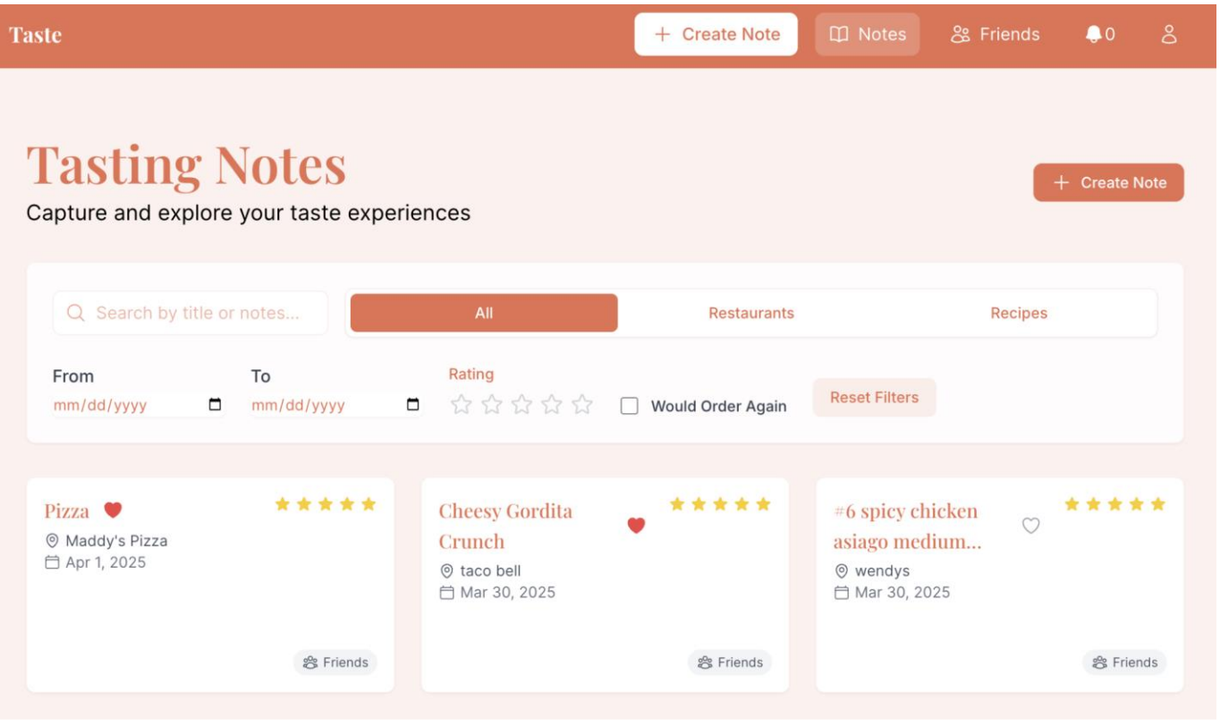
Vibe Coding

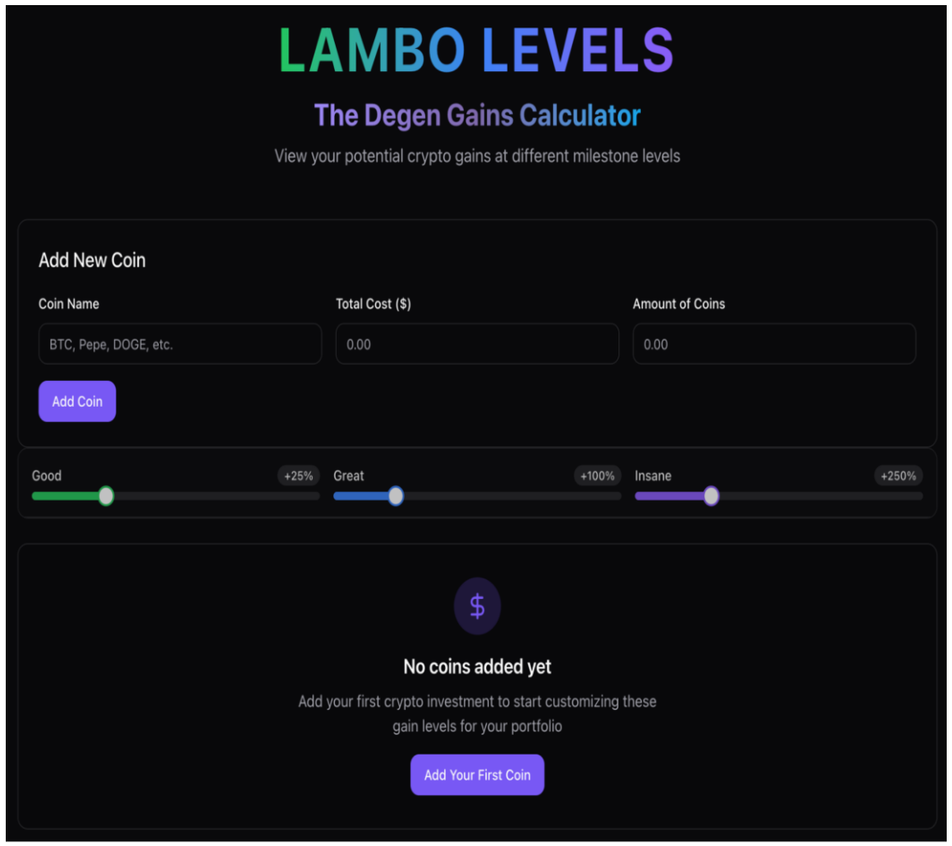
- Creating code with generative AI
- Coined by Andrej Karpathy in 2025, it describes a workflow where you "fully give in to the vibes" and forget the underlying code exists.

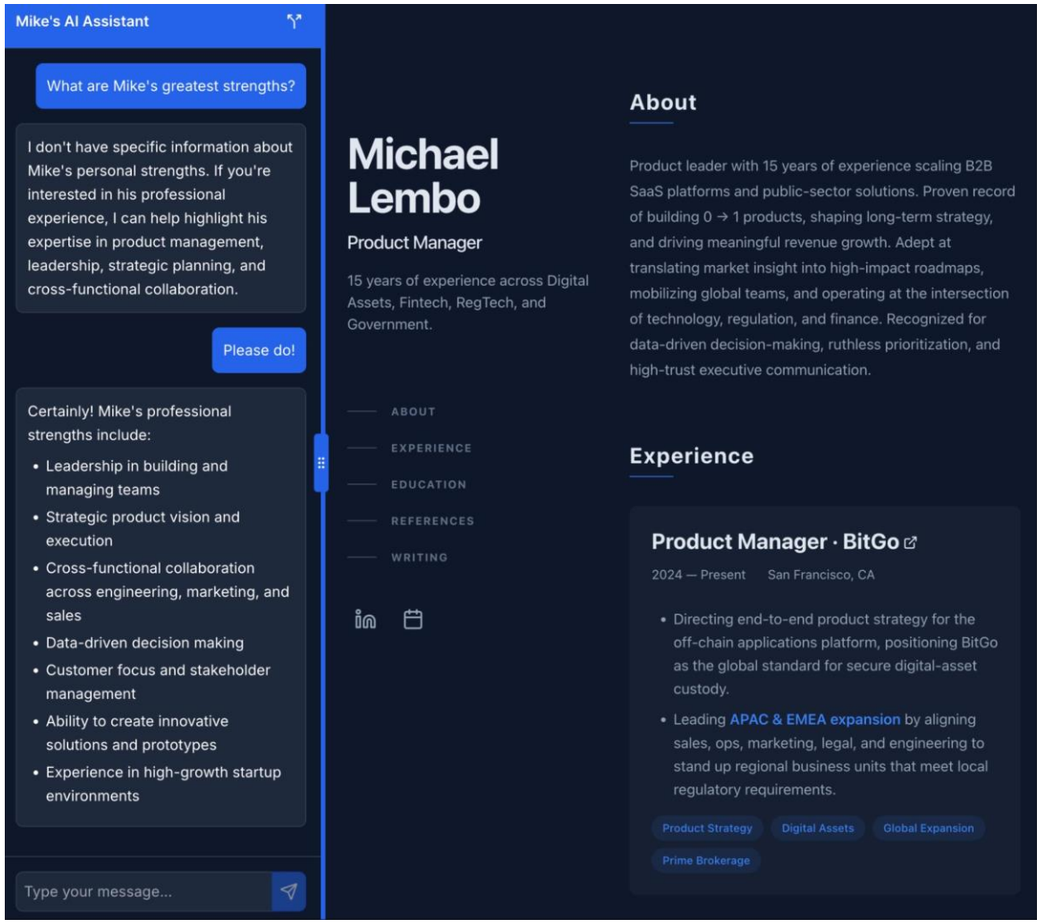


Vibe Coding

- **The Use Case:** "Disposable" research scripts, quick data visualizations, and front-end dashboards for demos.
- **The Workflow:** "Describe → Generate → Execute → Provide Feedback → Repeat."
- **"The "6-Month Wall":**
  - Pros: 70% faster prototyping; allows domain experts (biologists/physicists) to build tools without deep CS roots.
  - Cons: Compounding technical debt; security vulnerabilities; "Spaghetti code" that is hard to maintain without the AI's help.







NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

ICPE Higher Education Matters

University of Kentucky

The Pre-training Pipeline

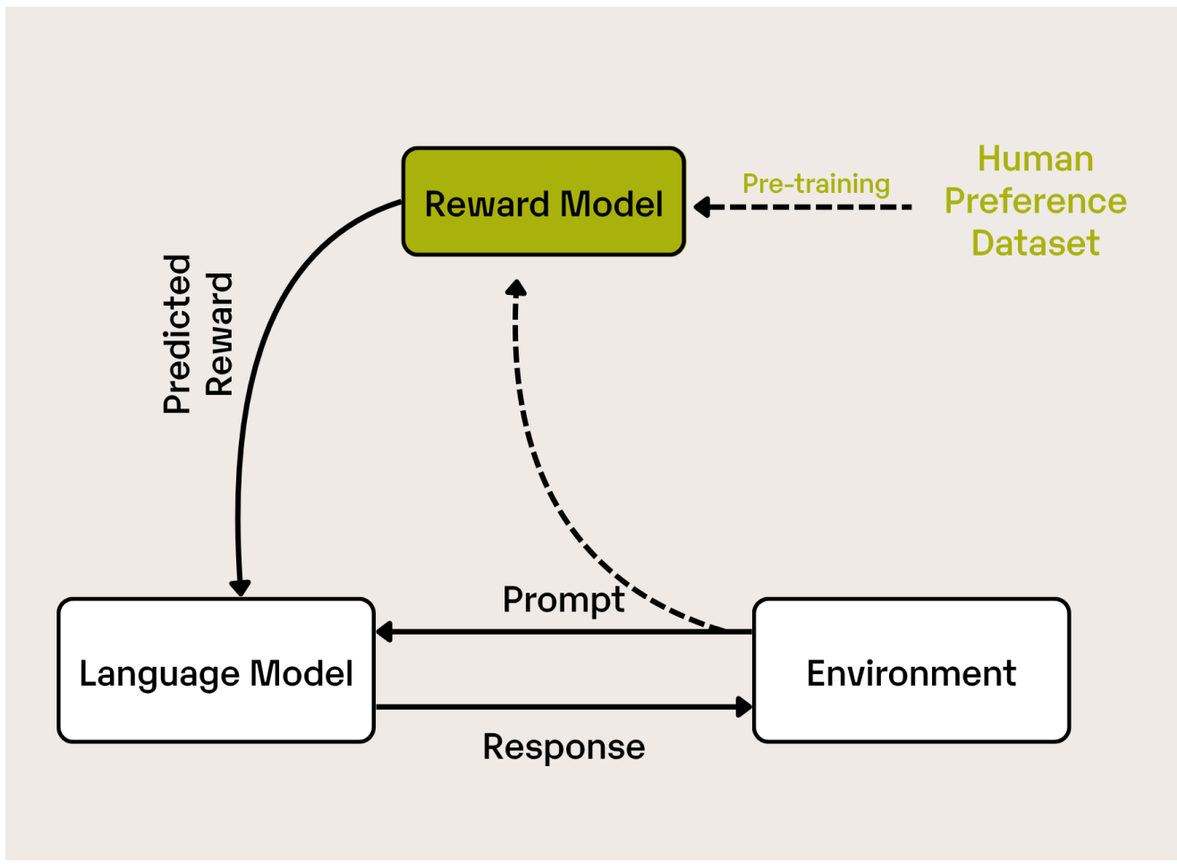
- **Data Sourcing (The Corpus):** Trillions of tokens from public GitHub repositories, StackOverflow threads, and technical documentation.
- **Cleaning & Deduplication:** Aggressive filtering to remove "boilerplate" (like massive node_modules accidentally checked in), low-stars repositories, and exact duplicates.
- **Objectives:**
  - Understanding the syntax of a language (grammar, brackets, indentation, etc.)
  - Fill-In-The-Middle: Model is given the start and end of a file and must guess the middle (Allows a model to edit existing code rather than just appending to the end)





Alignment & RLVR

- **Supervised Fine-Tuning (SFT):** Training on high-quality "Instruction-Solution" pairs (e.g., "Write a BFS algorithm in Rust").
- **RLHF (Human Feedback):** Humans rank multiple code solutions from "Best Practice" to "Buggy."
- **The Breakthrough:** RLVR (Reinforcement Learning with Verifiable Rewards):
  - –Unlike natural language, code can be objectively verified.
  - –The model generates a solution → The system runs it through a Compiler and a Unit Test suite.
  - –Reward: If it compiles and passes the tests, the model is "rewarded" (weights are reinforced). If it fails or has a runtime error, it's "punished."



The Upside

- **The Workflow:** Natural Language → Logic → Artifact.
- **The Advantage:** Eliminating the "Syntax Tax." Developers spend more time on system design and user experience rather than debugging semicolons or library imports.
- **Prototyping Velocity:** faster task completion in controlled lab settings.

The Downside

- **Technical Debt:** AI-generated code is often "Spaghetti Code" that works today but is structurally fragile.
- **Security Gaps:** Studies in late 2025 showed that up to 45% of AI-generated code contains classic vulnerabilities (OWASP Top 10), like missing input validation or hardcoded secrets.
- **The Maintenance Wall:** If you didn't write it, you can't easily "rewire" it when it breaks.

Feature	Pros	Cons
Speed	Near-instant prototypes	Hard-to-find logic bugs
Effort	Low cognitive load	High Maintenance Debt
Accessibility	Anyone can build	Security/Compliance risks
Logic	“It just works”	“Why does it work?”

Vibe coding with Google's Antigravity Demo



Vibe coding with Google's Antigravity
Demo

- Goal: Design a docker web app that visualizes weather data (This should work)

- Goal: Create a 3d simulation of atoms that can do the following things: (This probably won't work)
  1. Takes the quantum numbers (n, l, m) that describe an orbital's shape
  2. Using the Schrodinger equation, sample r, theta, and phi coordinates from those quantum numbers
  3. Render those possible positions and color code them relative to their probabilities (brighter areas have higher probability)

Foundational Models & LLMs

Mitchell Klusty, BS

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Applying ML to Language

Human Communication

Text

Speech

Machine Processes

Calculation

Finding Patterns

Data Storage

Applying ML to Language



Applications of NLP

- Search Engines
- Voice Assistant
- Email Spam Filter
- Auto-complete/Grammar Check
- Literature Review
- Clinical Note Processing
- Translation
- Code Generation
- Summarization
- Question Answering
- And More

How Computers "Read"

- **Strings:** Data type of text
- Words are made up of characters
- To a computer, just seen as a set of bits

Hello world!

Human Readable Text String

How Computers "Read"

- **Strings:** Data type of text
- Words are made up of characters
- To a computer, just seen as a set of bits

Hello world!	Human Readable Text String
\x48\x65\x6C\x6C\x6F\x20\x57\x6F\x72\x6C\x64\x21	UTF-8 Encoding
01001000 01100101 01101100 01101100 01101111 00100000 01110111 01101111 01110010 01101100 01100100 00100001	Binary representation

Early Approaches: Text as Statistics

- Instead of meaning, look for counts or patterns

Bag of Words: Find word importance by frequency within inputs

"Large language models are transforming research by enabling rapid analysis of text, automation of documentation, and intelligent search across large collections of data. These models help researchers summarize papers, generate hypotheses, and explore complex datasets more efficiently."



word	count
of	3
large	2
and	2
models	2
transforming	1
research	1
language	1

Early Approaches: Text as Statistics



Word Cloud

Highlights most common words, meant to show meaning of text

Early Approaches: Text as Statistics

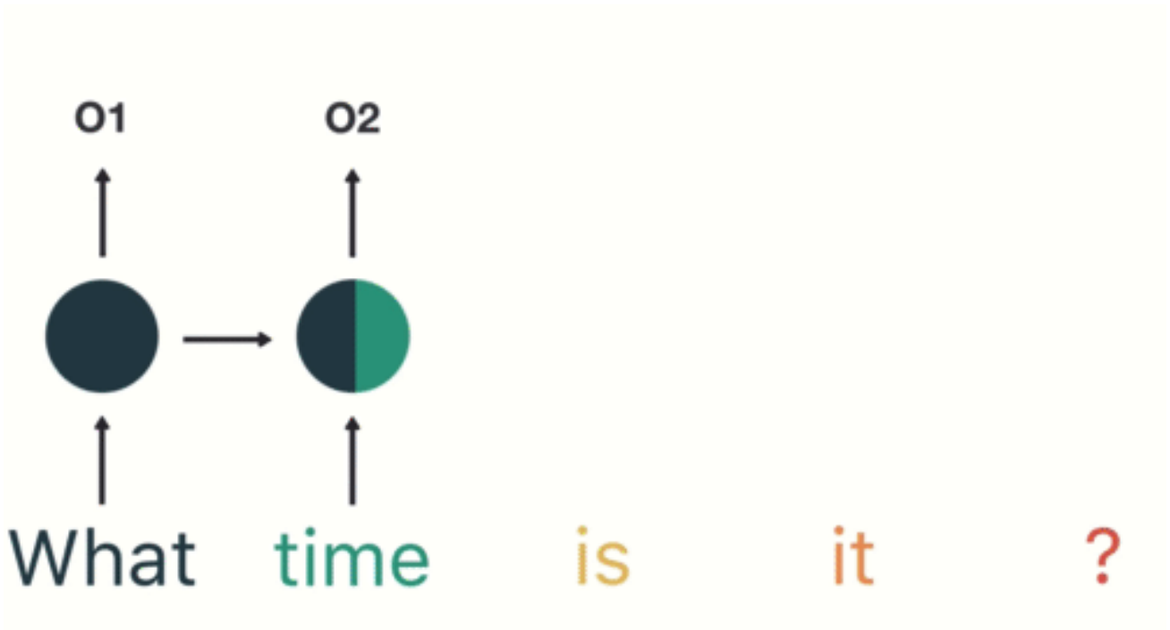


Problems:

- “Stop-words”: of, and, the, etc.
- Order of words ignored
- Doesn’t tell us much
- Meaning of text not captured very well

Recurrent Neural Networks (RNNs)

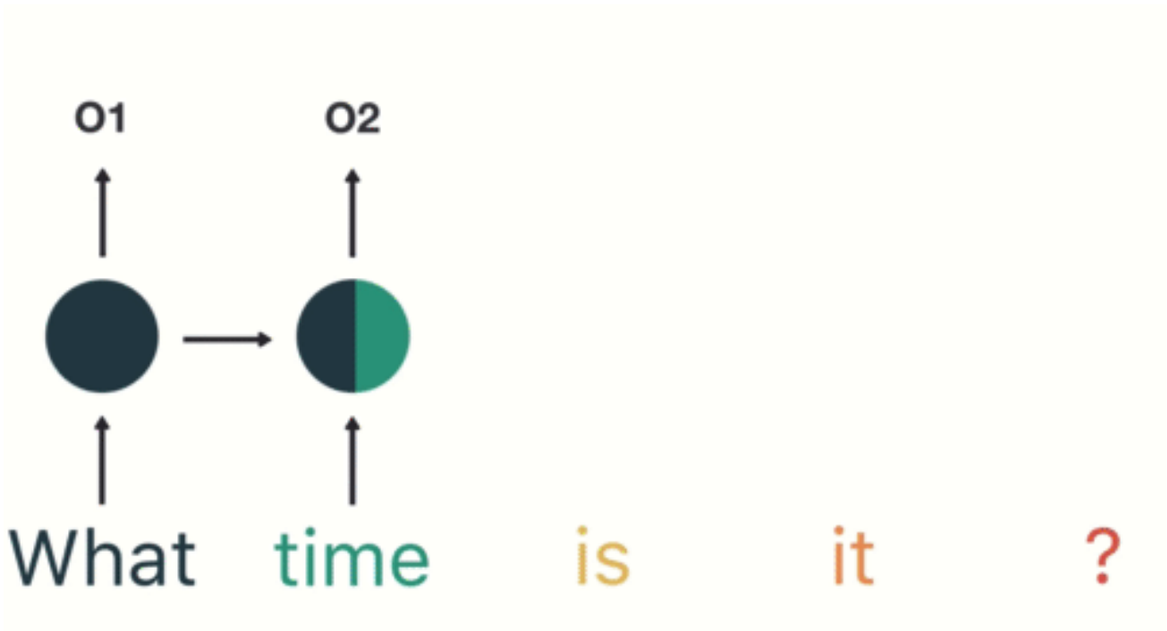
- ML technique to understand information in a sequence
- Processes data in order, start to finish, **one step at a time**
- Keeps a “memory” of what it’s already seen
- Each step dependent on what came before



Recurrent Neural Networks (RNNs)

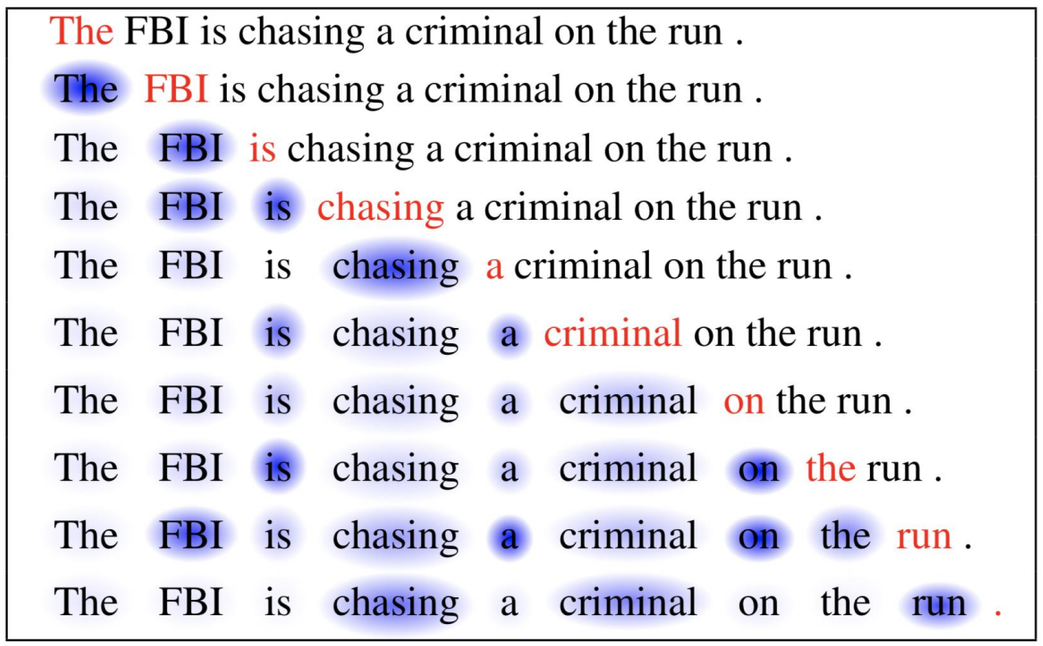
Limitations

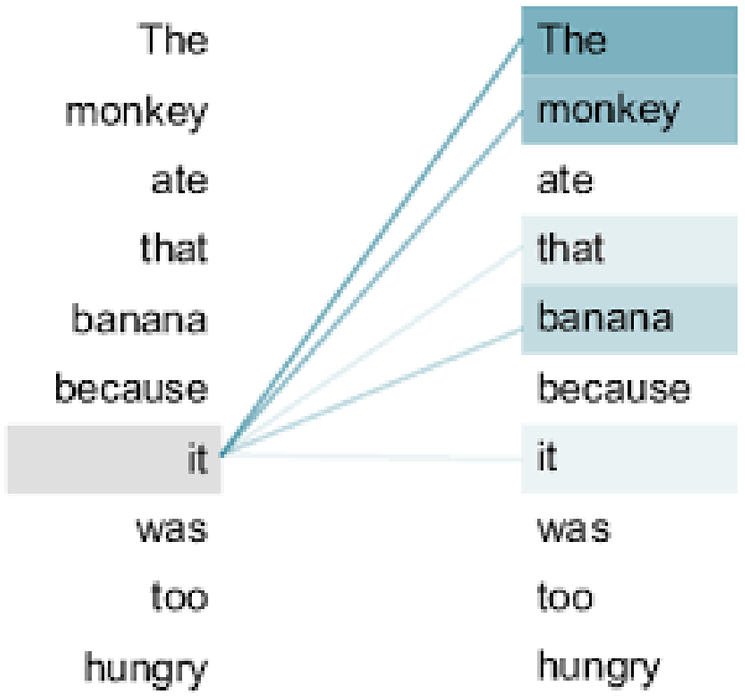
- “Memory” fades over long sequences
- Slow and hard to scale



Attention Is All You Need

- 2017 paper introducing the concept of **attention** and **transformer models**
- Through seeing vast amounts of examples, model begins to understand relationships in text
- Model looks at entire sequence at once
- Each word “attends” to each other word
- Stronger attention indicates stronger relationsh

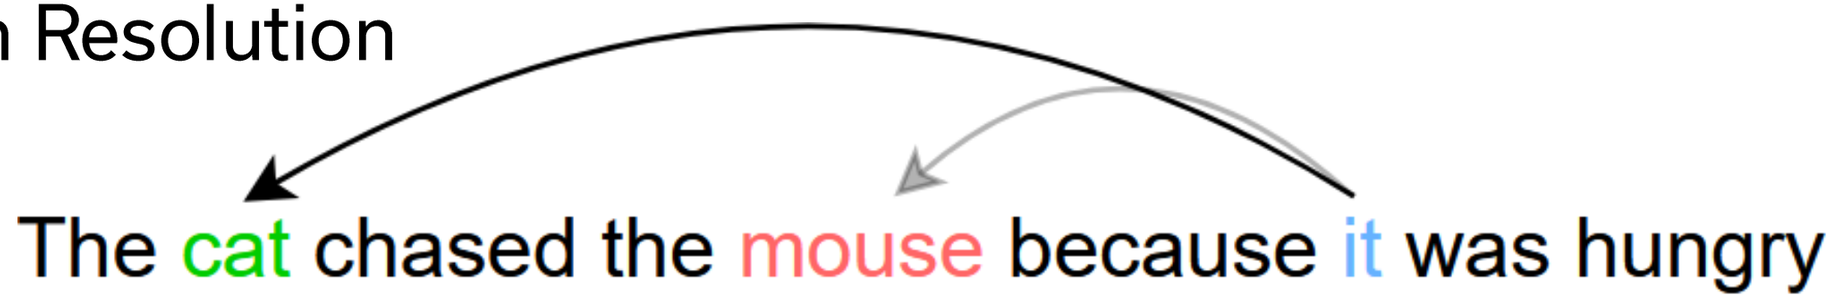




Attention Is All You Need

- 2017 paper introducing the concept of **attention** and **transformer models**
- Through seeing vast amounts of examples, model begins to understand relationships in text
- Model looks at entire sequence at once
- Each word “attends” to each other word
  - Stronger attention indicates stronger relationship

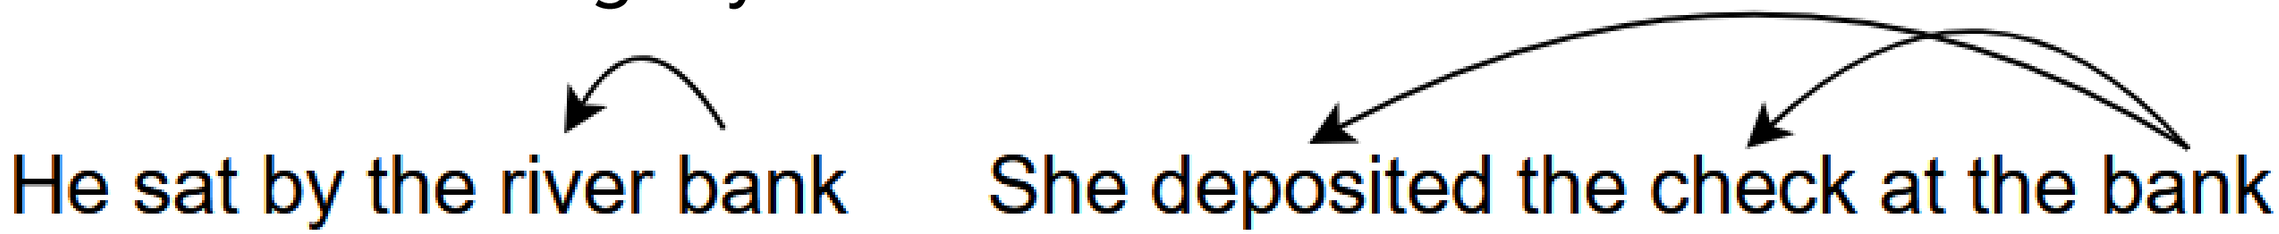
Pronoun Resolution



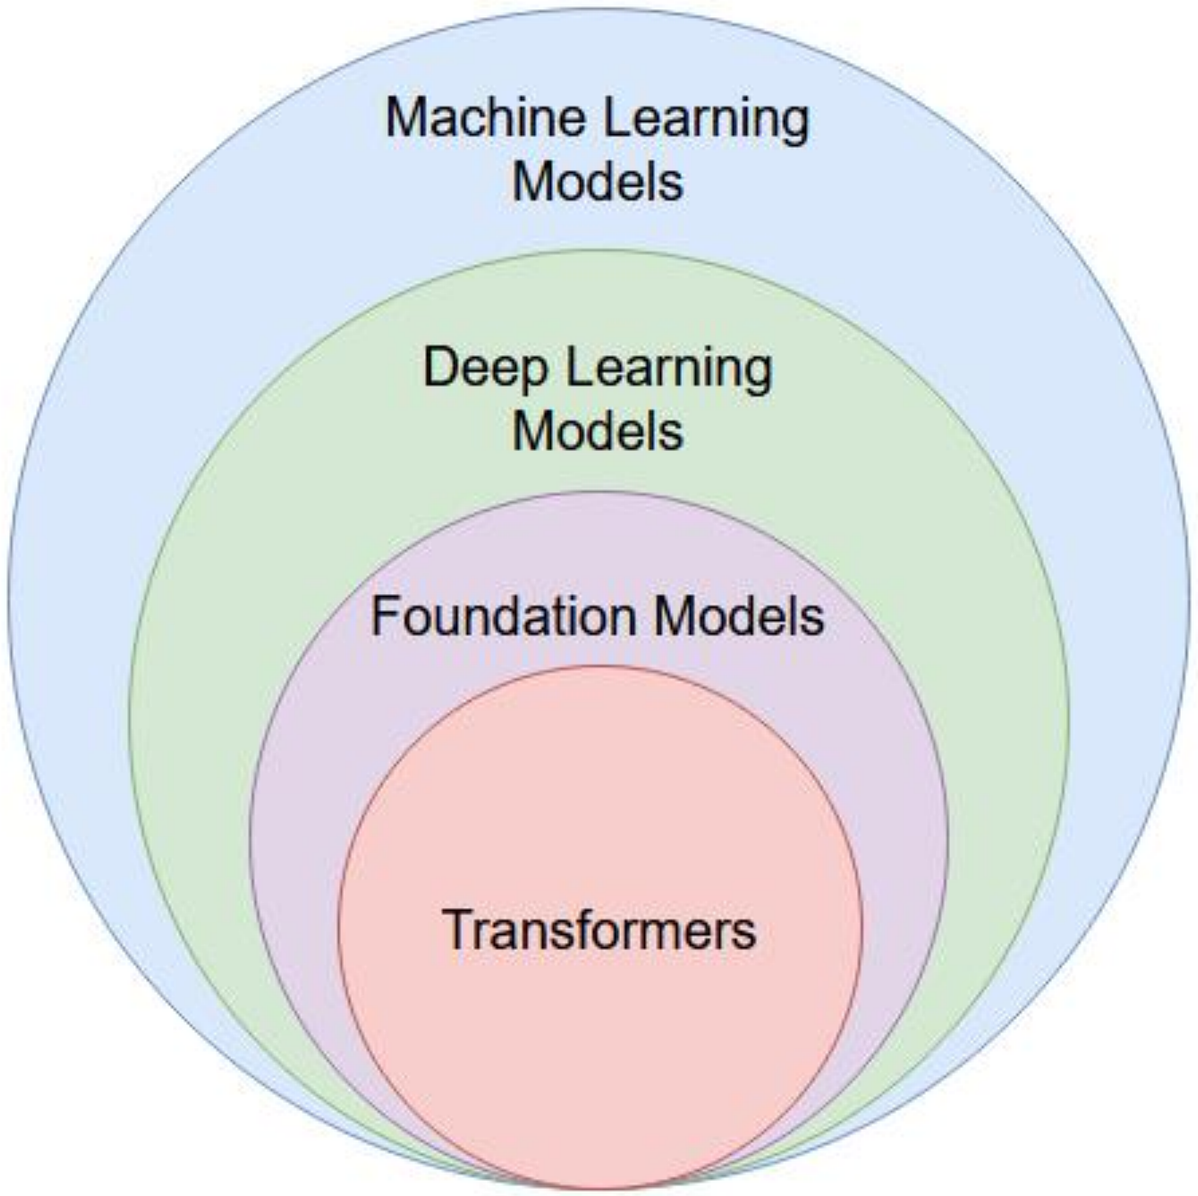
Attention Is All You Need

- 2017 paper introducing the concept of **attention** and **transformer models**
- Through seeing vast amounts of examples, model begins to understand relationships in text
- Model looks at entire sequence at once
- Each word “attends” to each other word
  - Stronger attention indicates stronger relationship

Remove word ambiguity



Transformers as Foundation Models

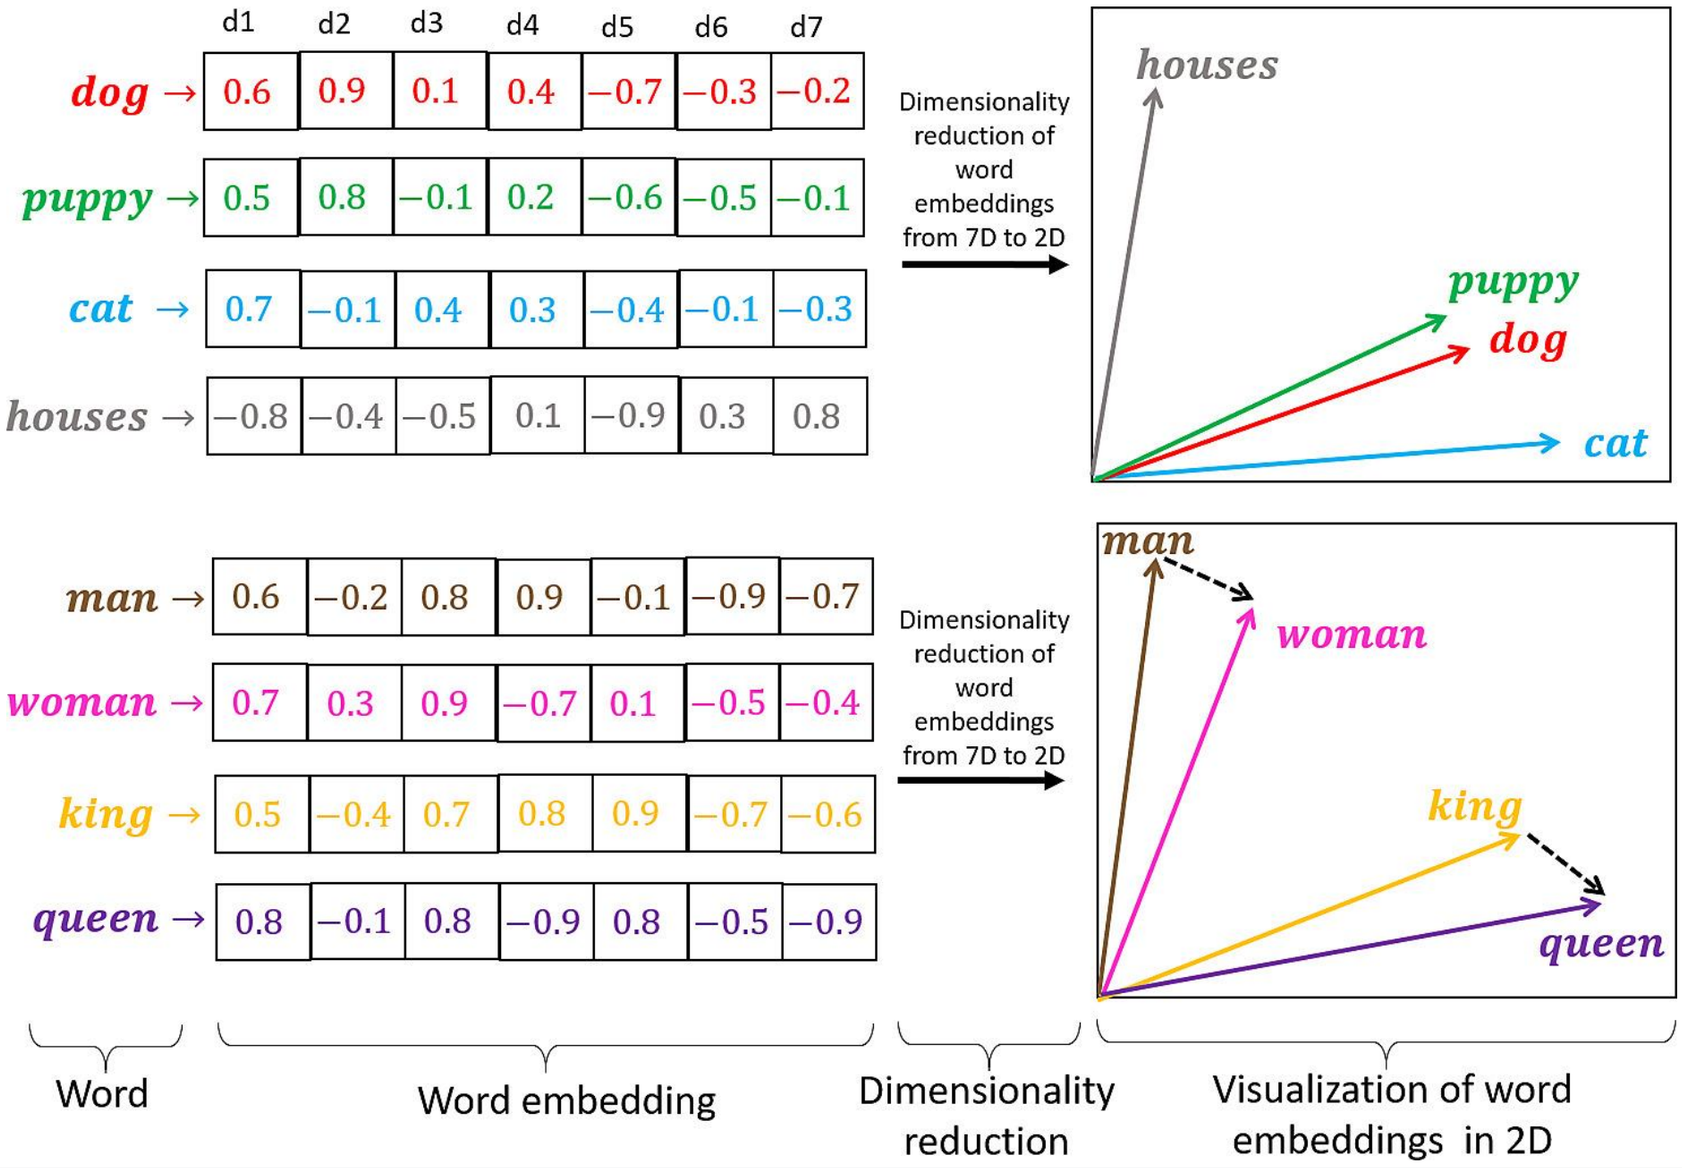


What is a Foundation Model?

- Originally, AI systems were narrow-purposed and task-specific
- **Foundation models** (FMs) are general-purpose
  - Trained once, applied to many tasks
- Most FMs are transformer-based
  - **Self-supervised** training
    - As it sees more data,
    - Learn broader features of the data rather than labels
    - Create **embeddings** of data

What are Embeddings?

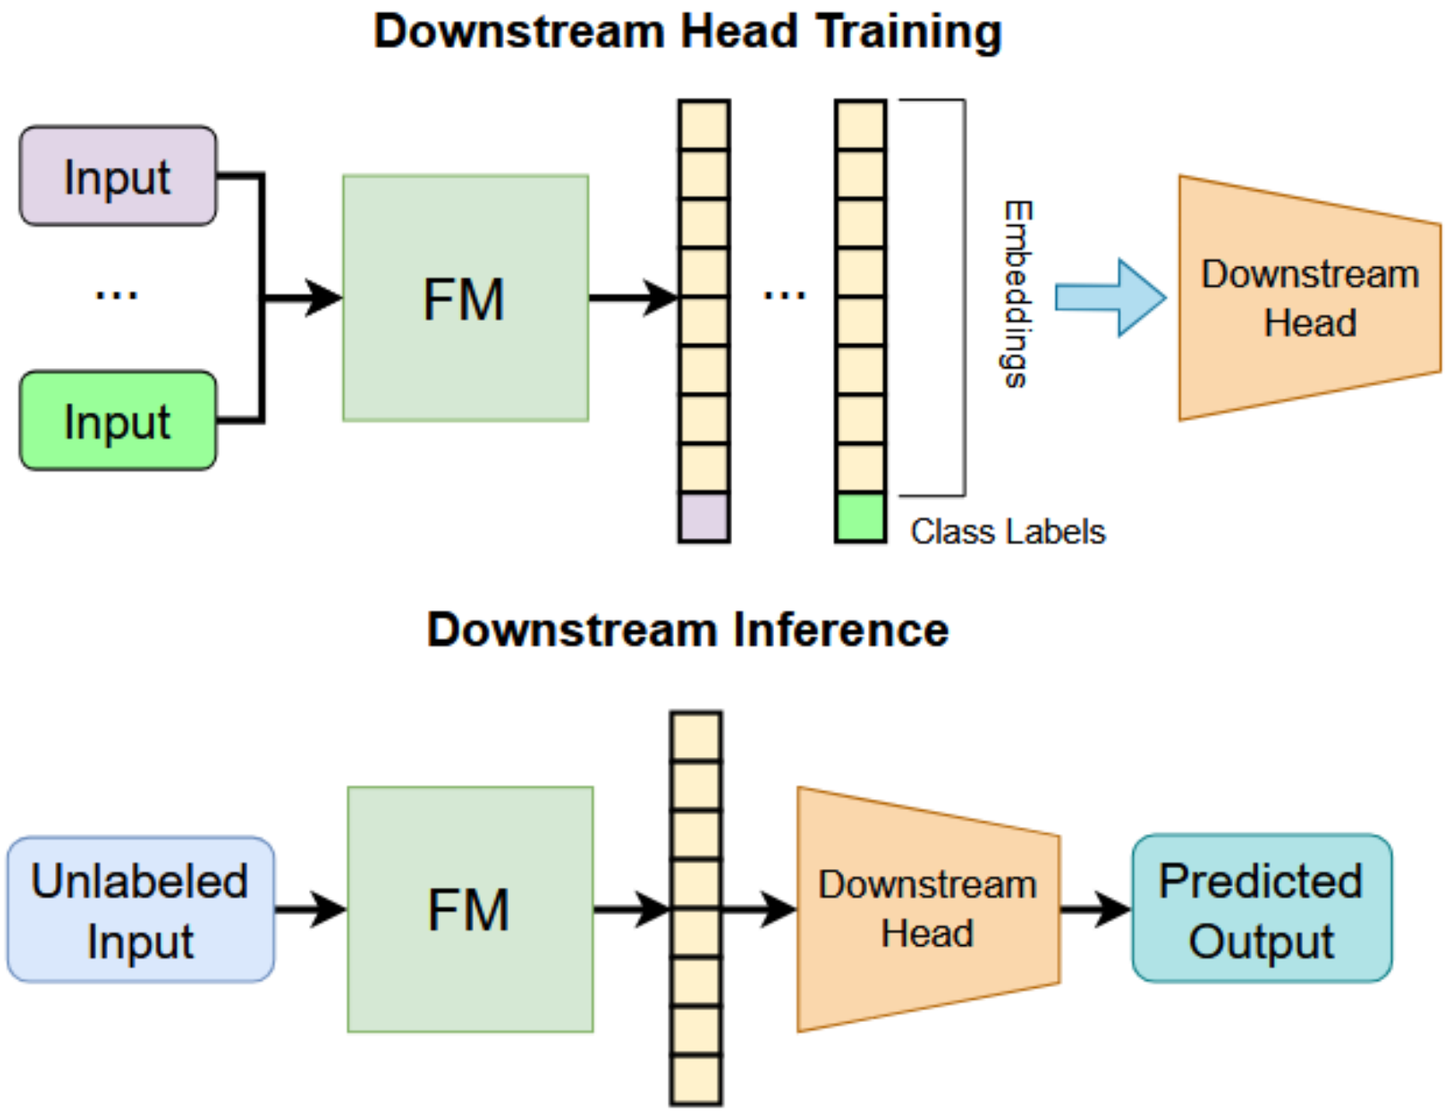
- **Embeddings** are a high-dimensional vector (set of numbers) that represent the features of input data based on FM's knowledge



Applying Foundation Models

- Task-specific **heads**
  - **Transfer Learning** from the FM to heads for downstream tasks
  - Allows for model versatility

Embeddings used as features in traditional ML model architectures



Tasks include:

- Classification of images or text
- Image Segmentation
- Linear Regression
- Next token prediction

Training a Large Language Model

1. Collect and Preprocess Data
2. Tokenization
3. Pre-Train the Transformer
4. Fine-tune

Training a Large Language Model –
Collect and Preprocess Data

- Scrape sources from the web
- Existing datasets
  - HuggingFace
- Personalized data

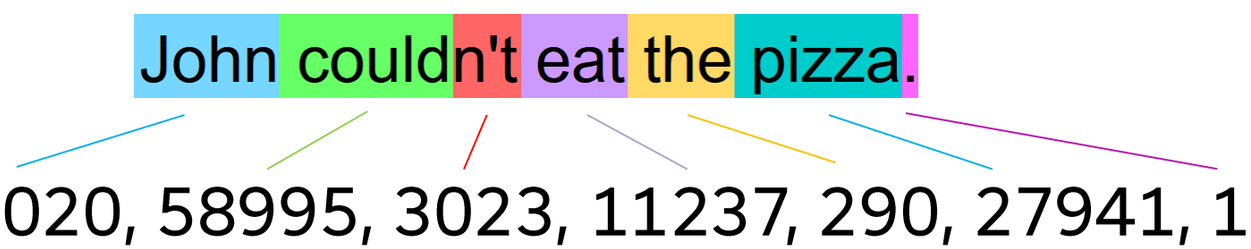


Training a Large Language Model – *Tokenization*

What is a token?

- Word/word-like unit
- Units of the input sequence for an LLM

John couldn't eat the pizza.

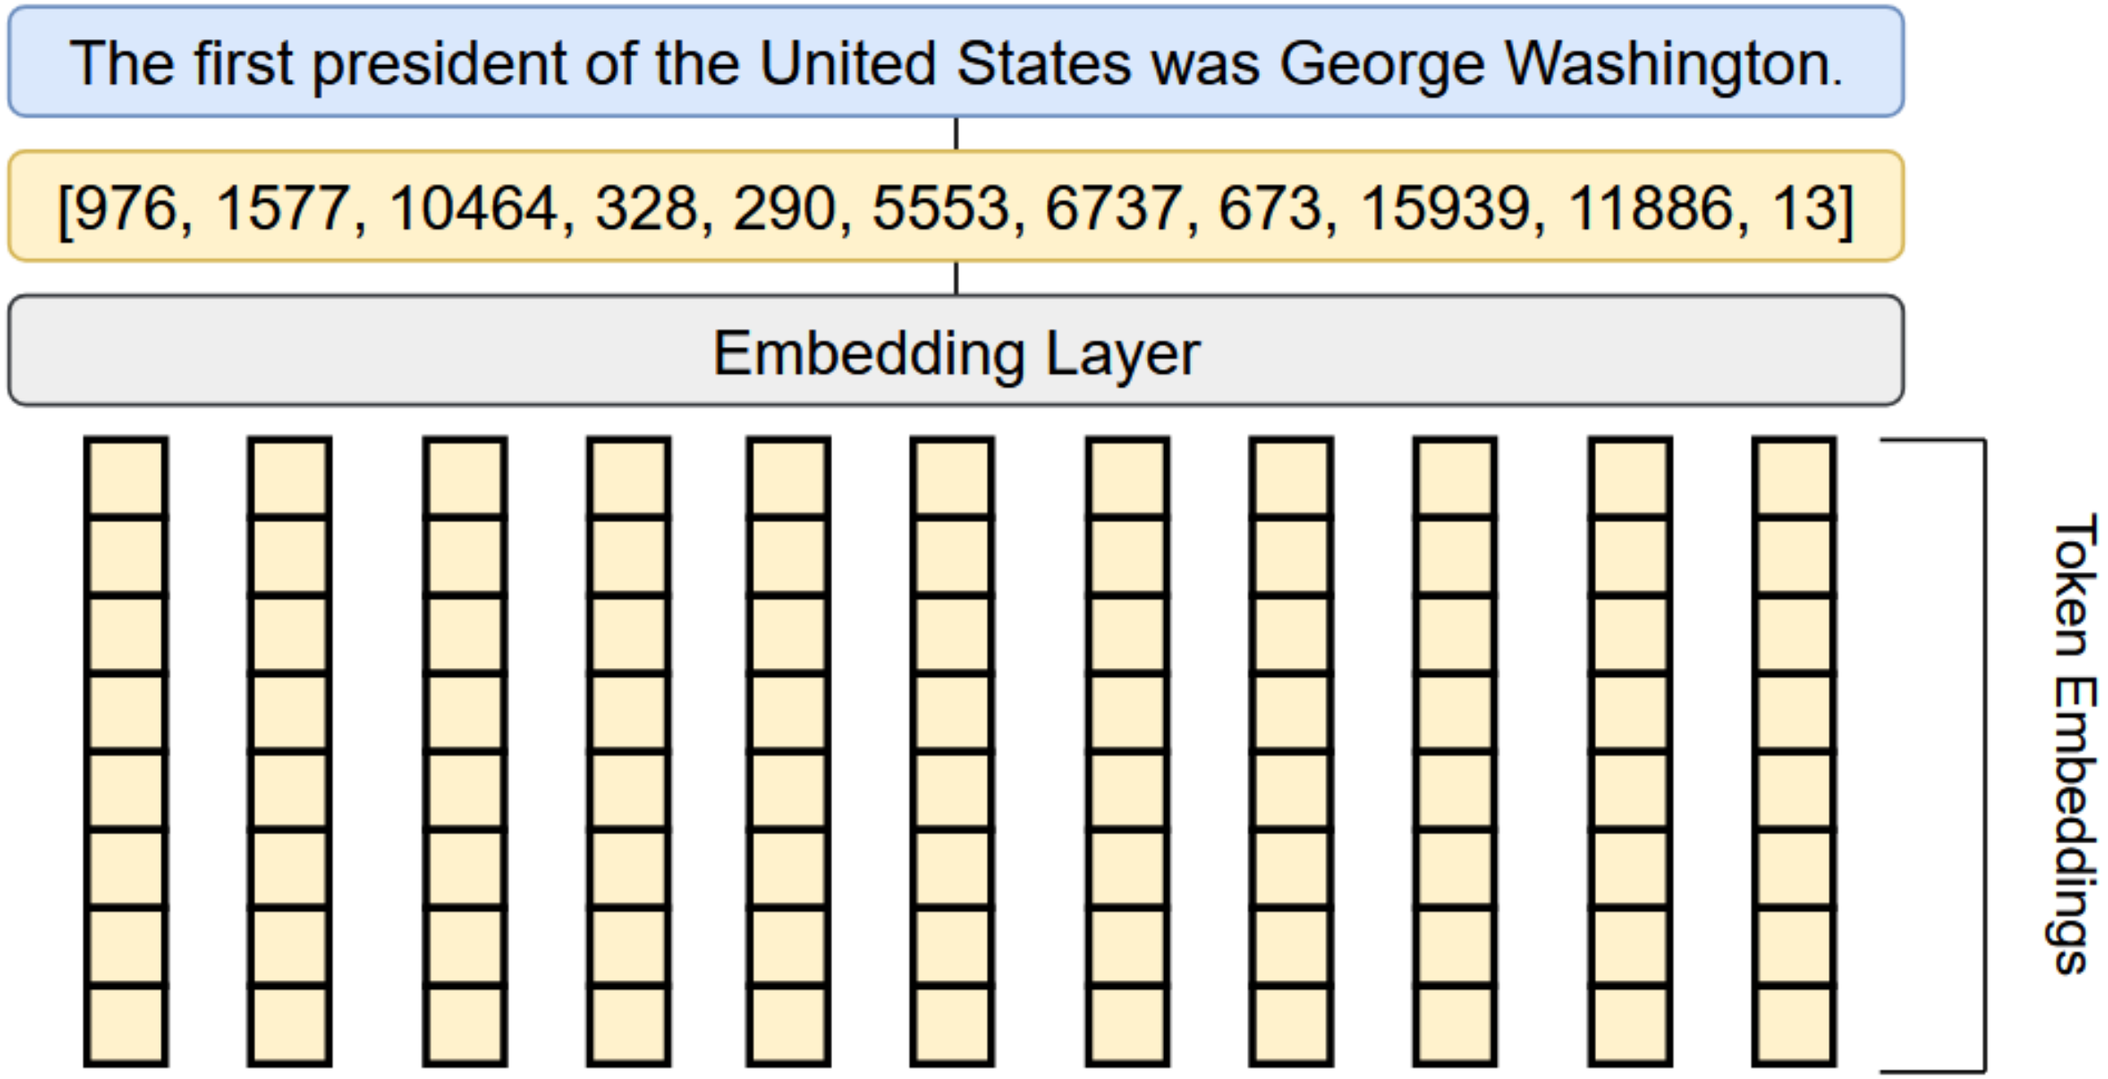


[21020, 58995, 3023, 11237, 290, 27941, 13] ← Token IDs

- Common sequences within the training dataset
  - Often turns out to be words, but not always
  - Punctuation, contractions, etc.
- The set of all tokens is the model's **vocabulary**
- Each token associated with an ID that maps to its embedding
- This ID is what the model sees

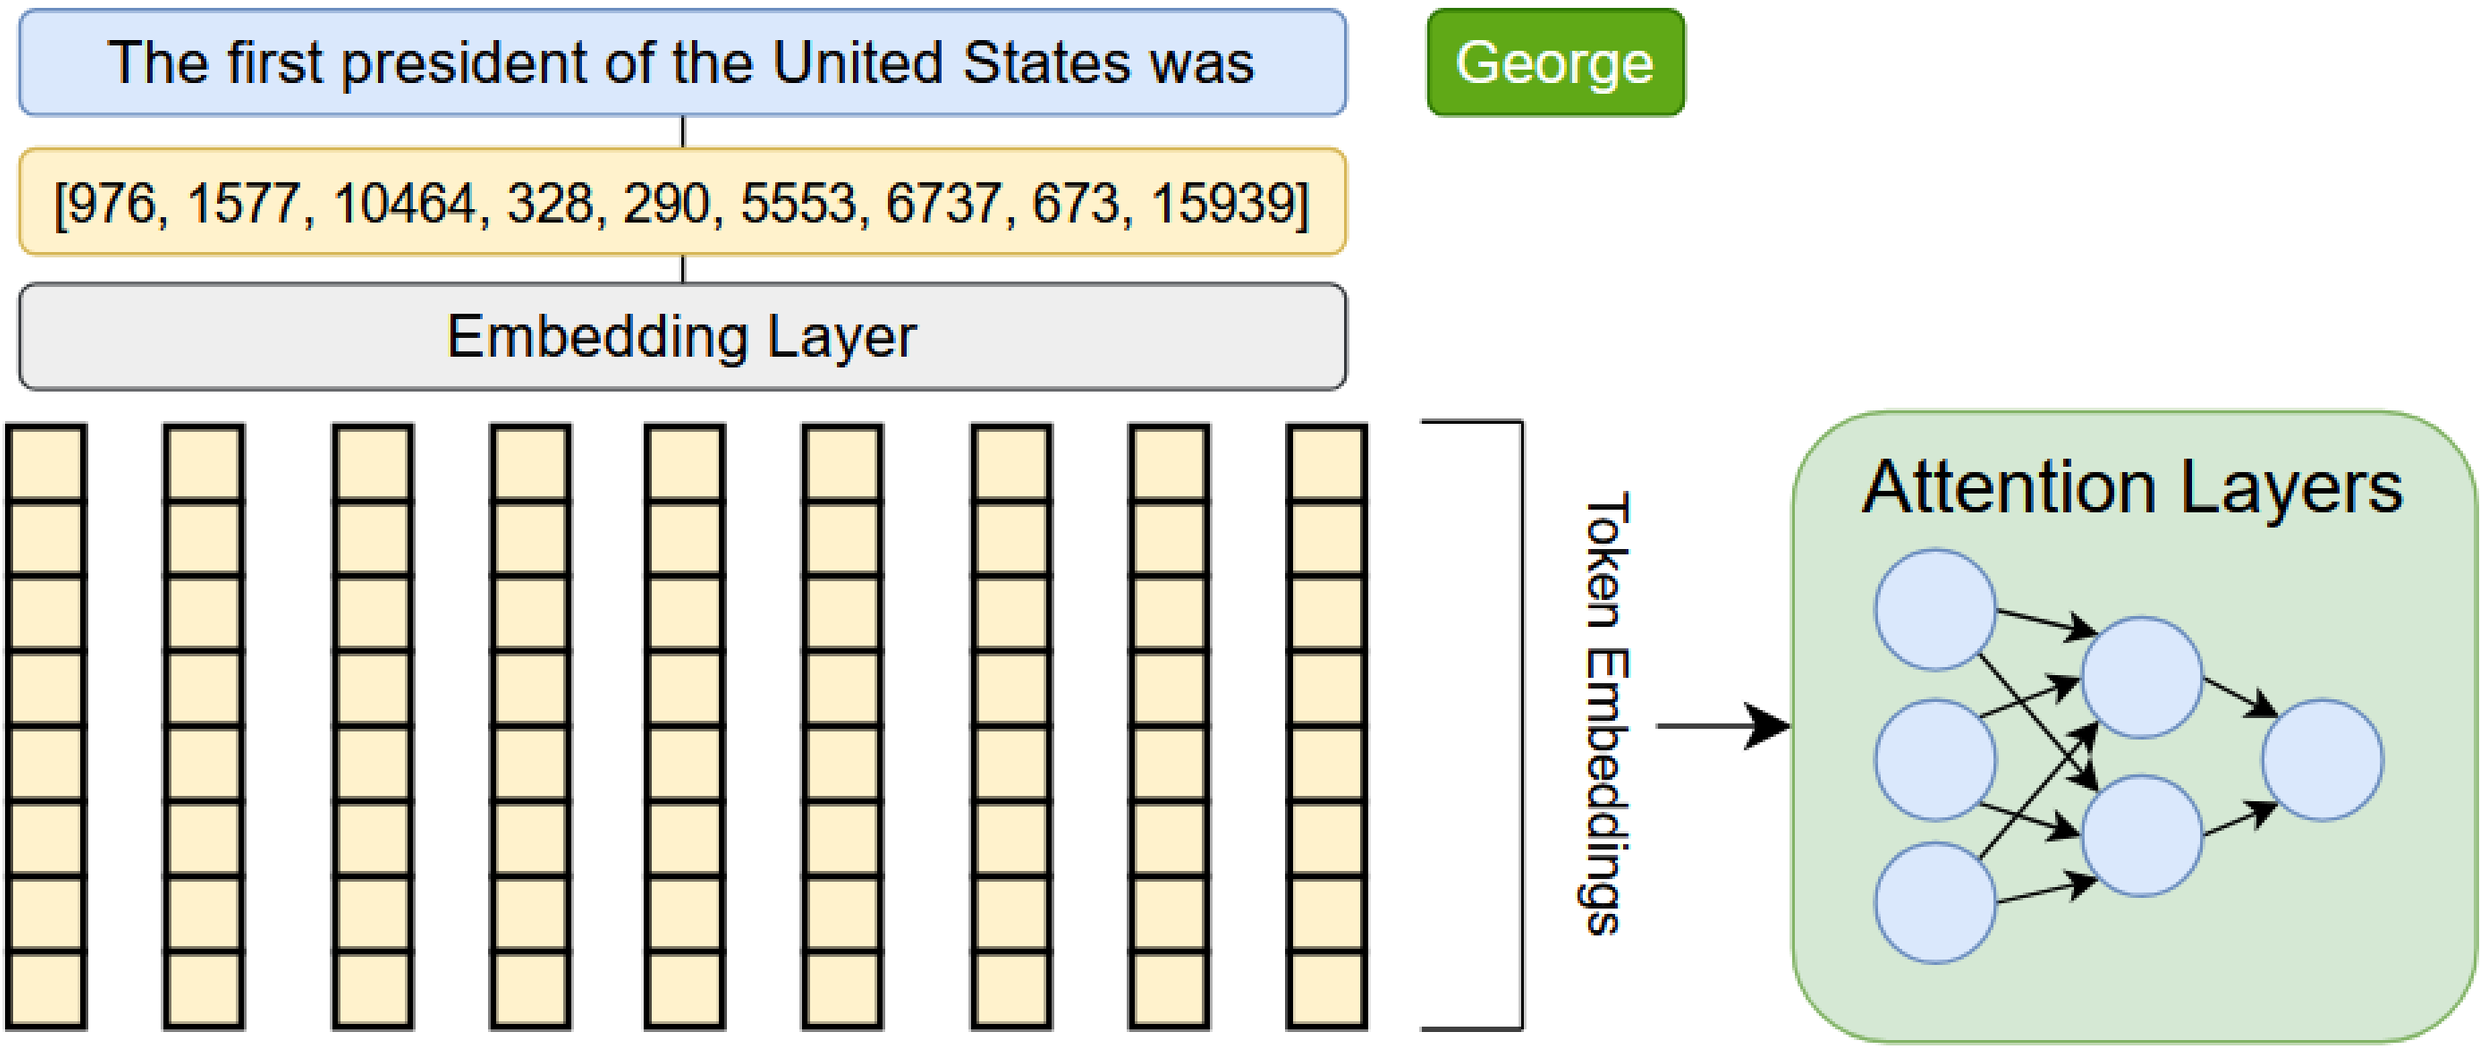
Training a Large Language Model –
Pre-Training the Transformer

- Sequence has existing embeddings for each token



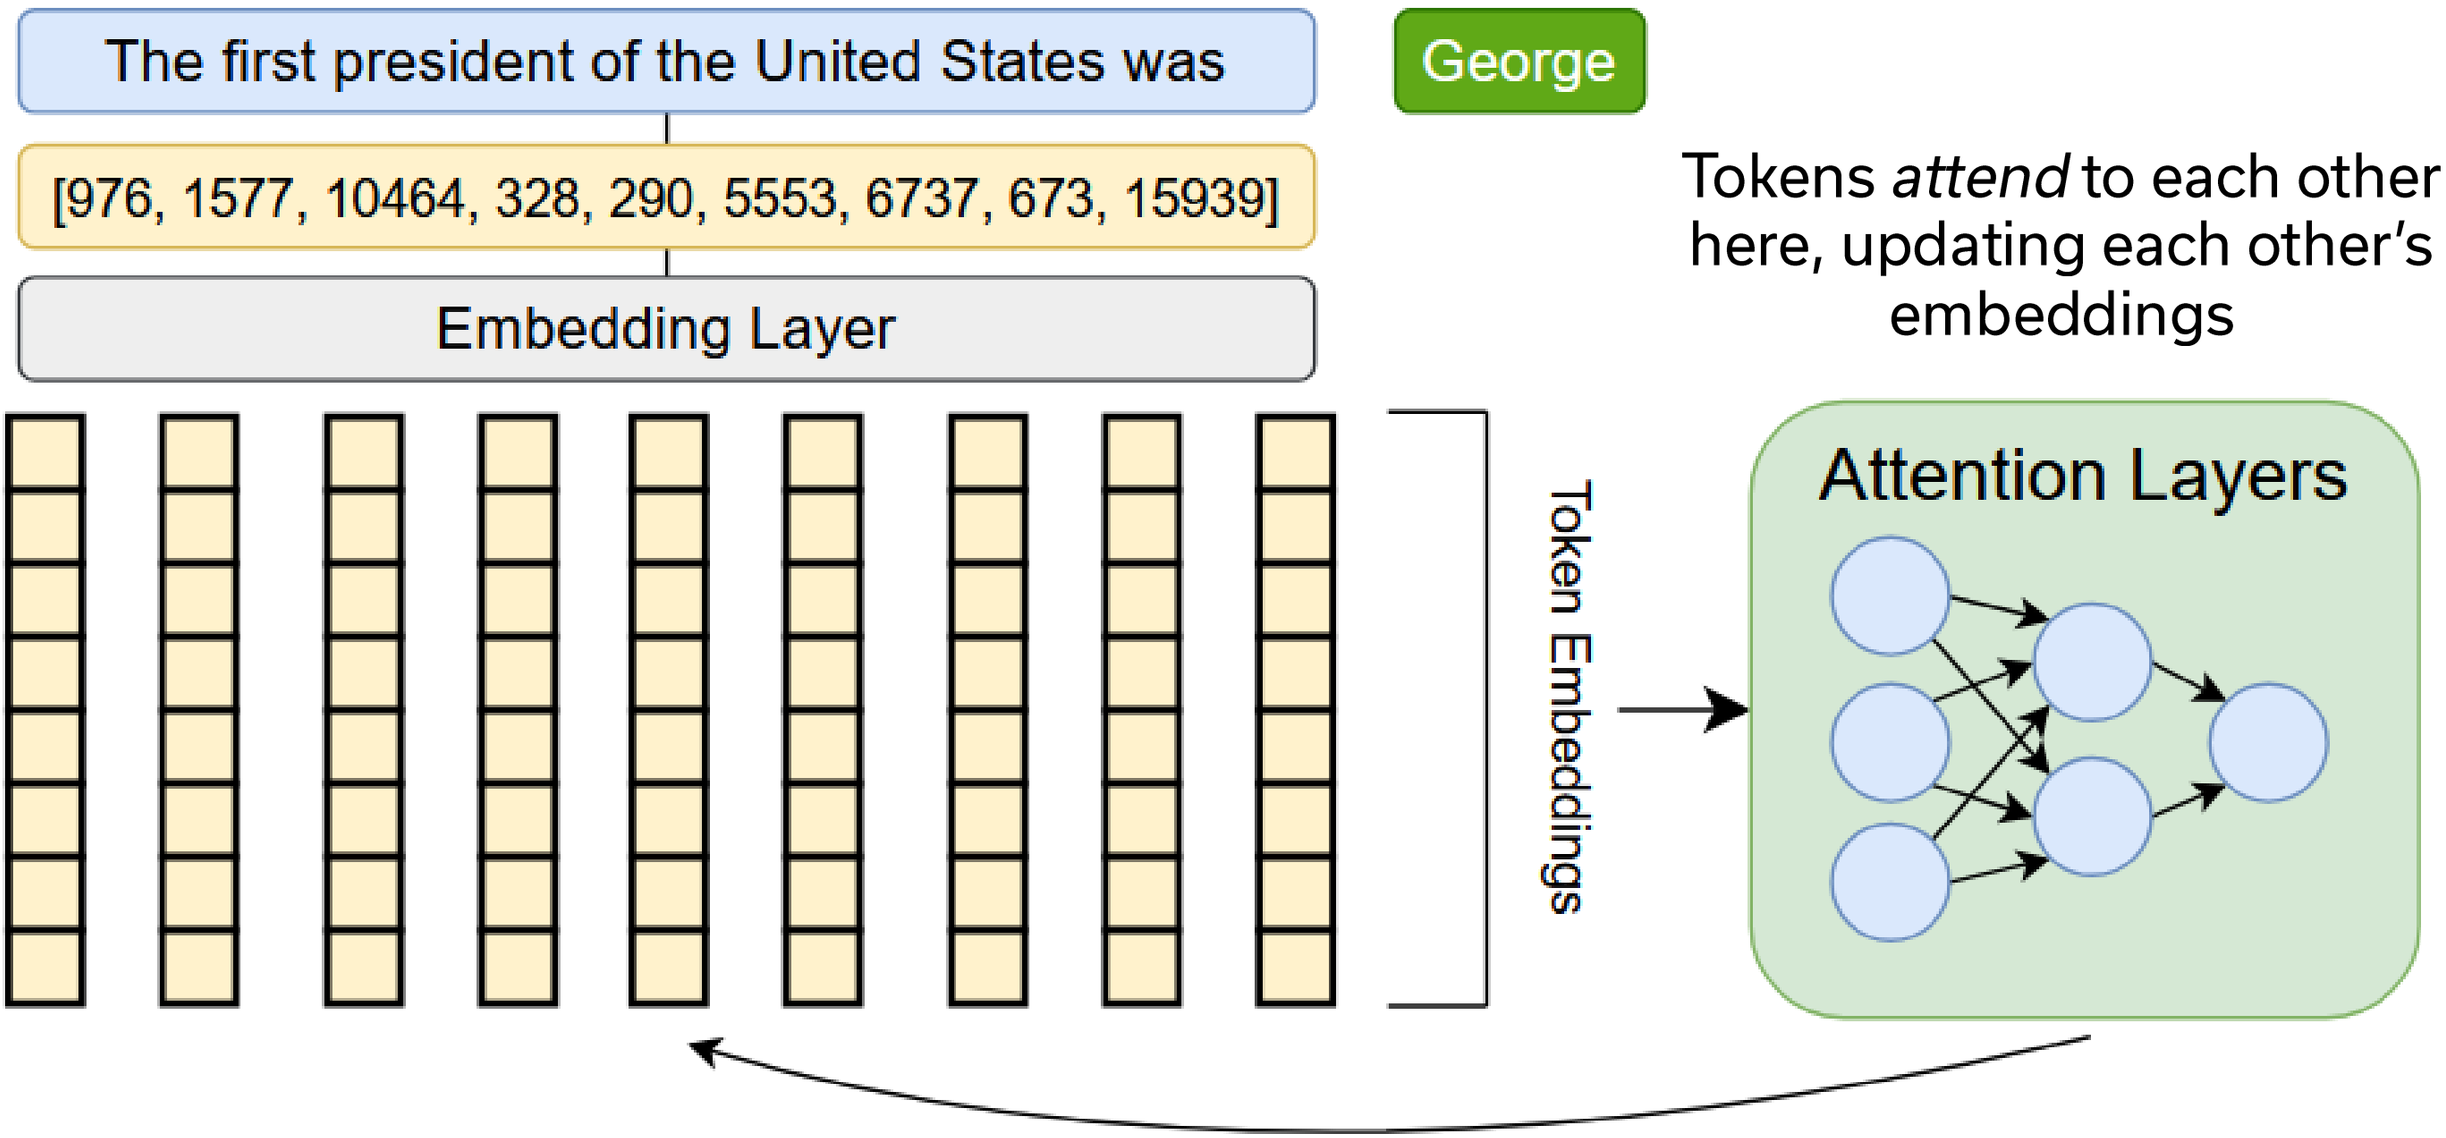
Training a Large Language Model –
Pre-Training the Transformer

- Provide an unfinished sequence
- Model must accurately predict next token



Training a Large Language Model –
Pre-Training the Transformer

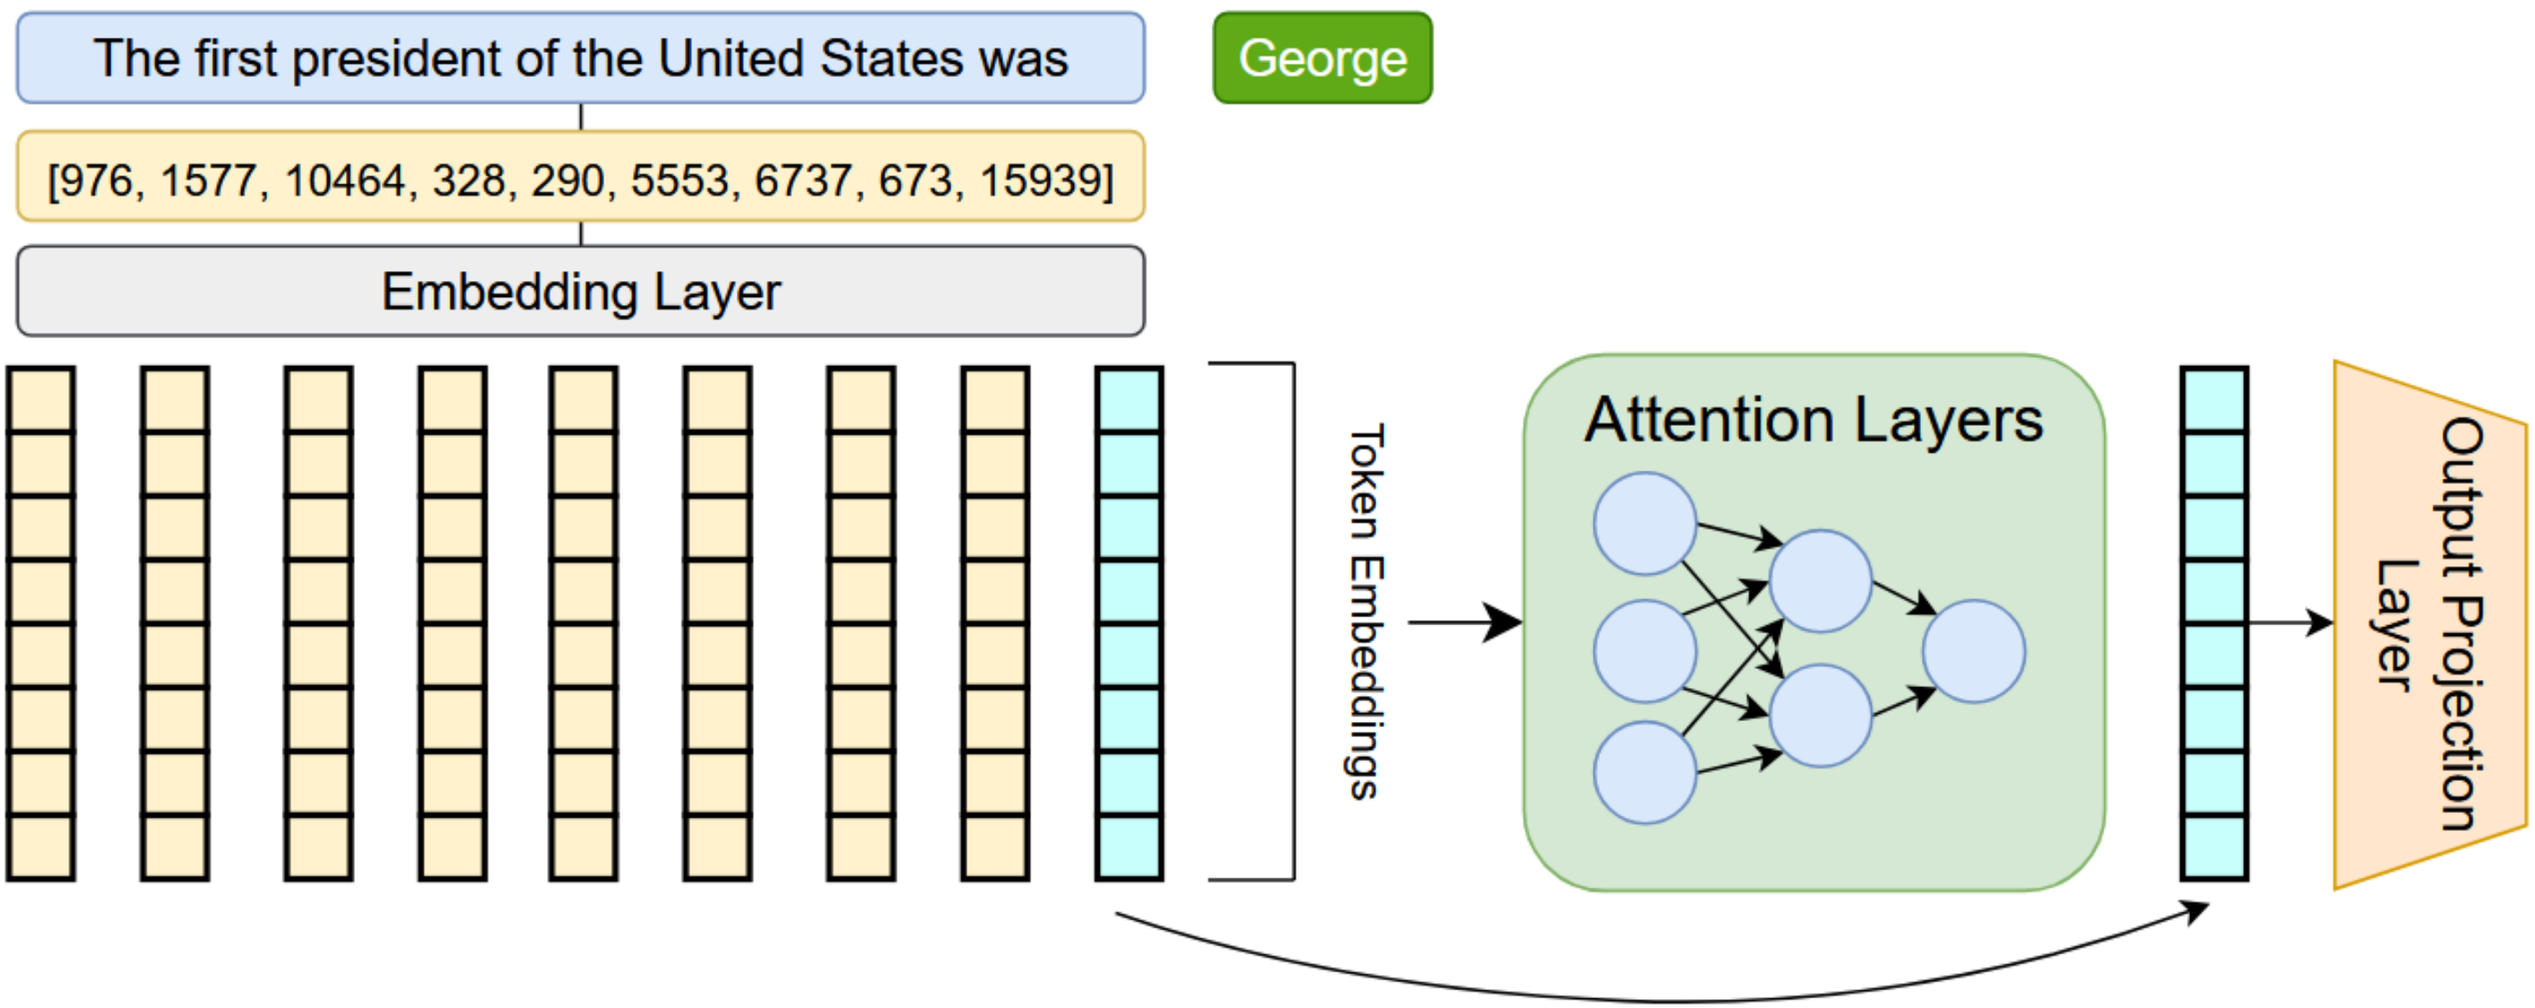
- Provide an unfinished sequence
- Model must accurately predict next token



Tokens *attend* to each other here, updating each other's embeddings

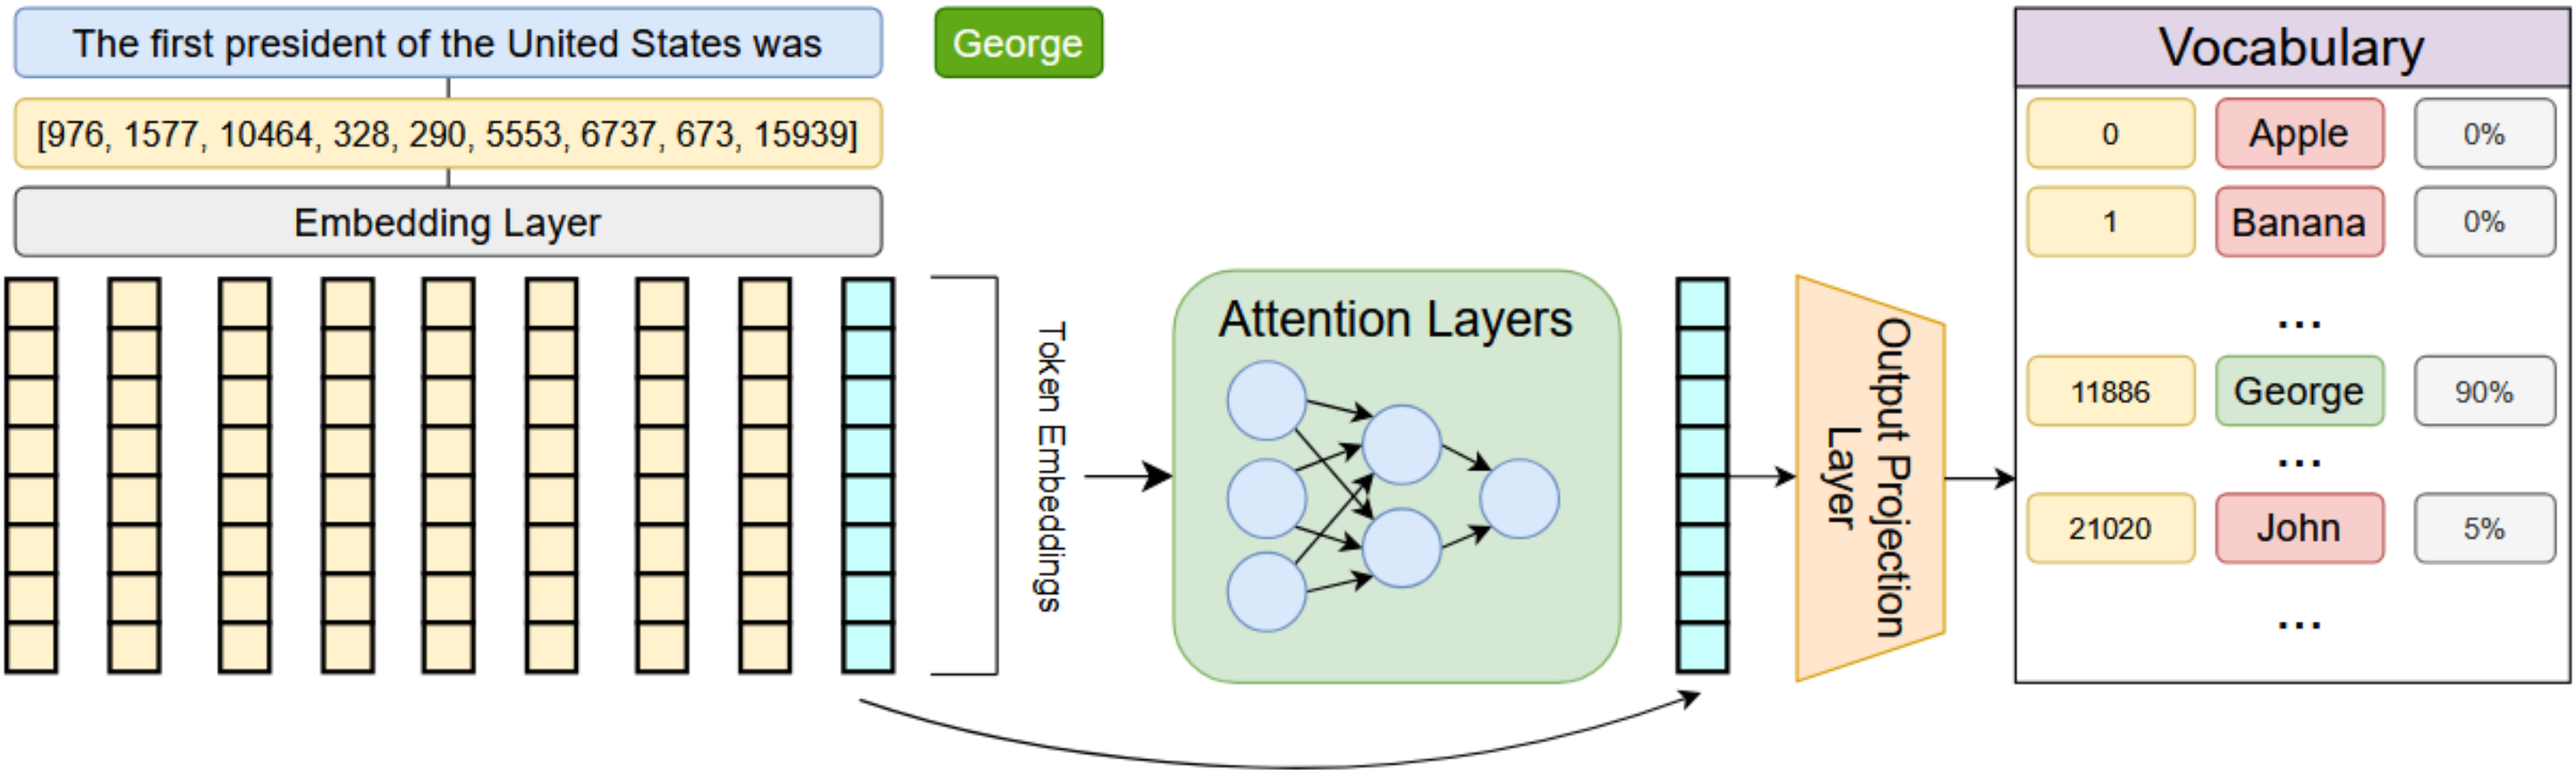
Training a Large Language Model – *Pre-Training the Transformer*

- Attended embedding of the last token passed to the **output projection layer**
  - Also called **Language Modeling (LM) head**



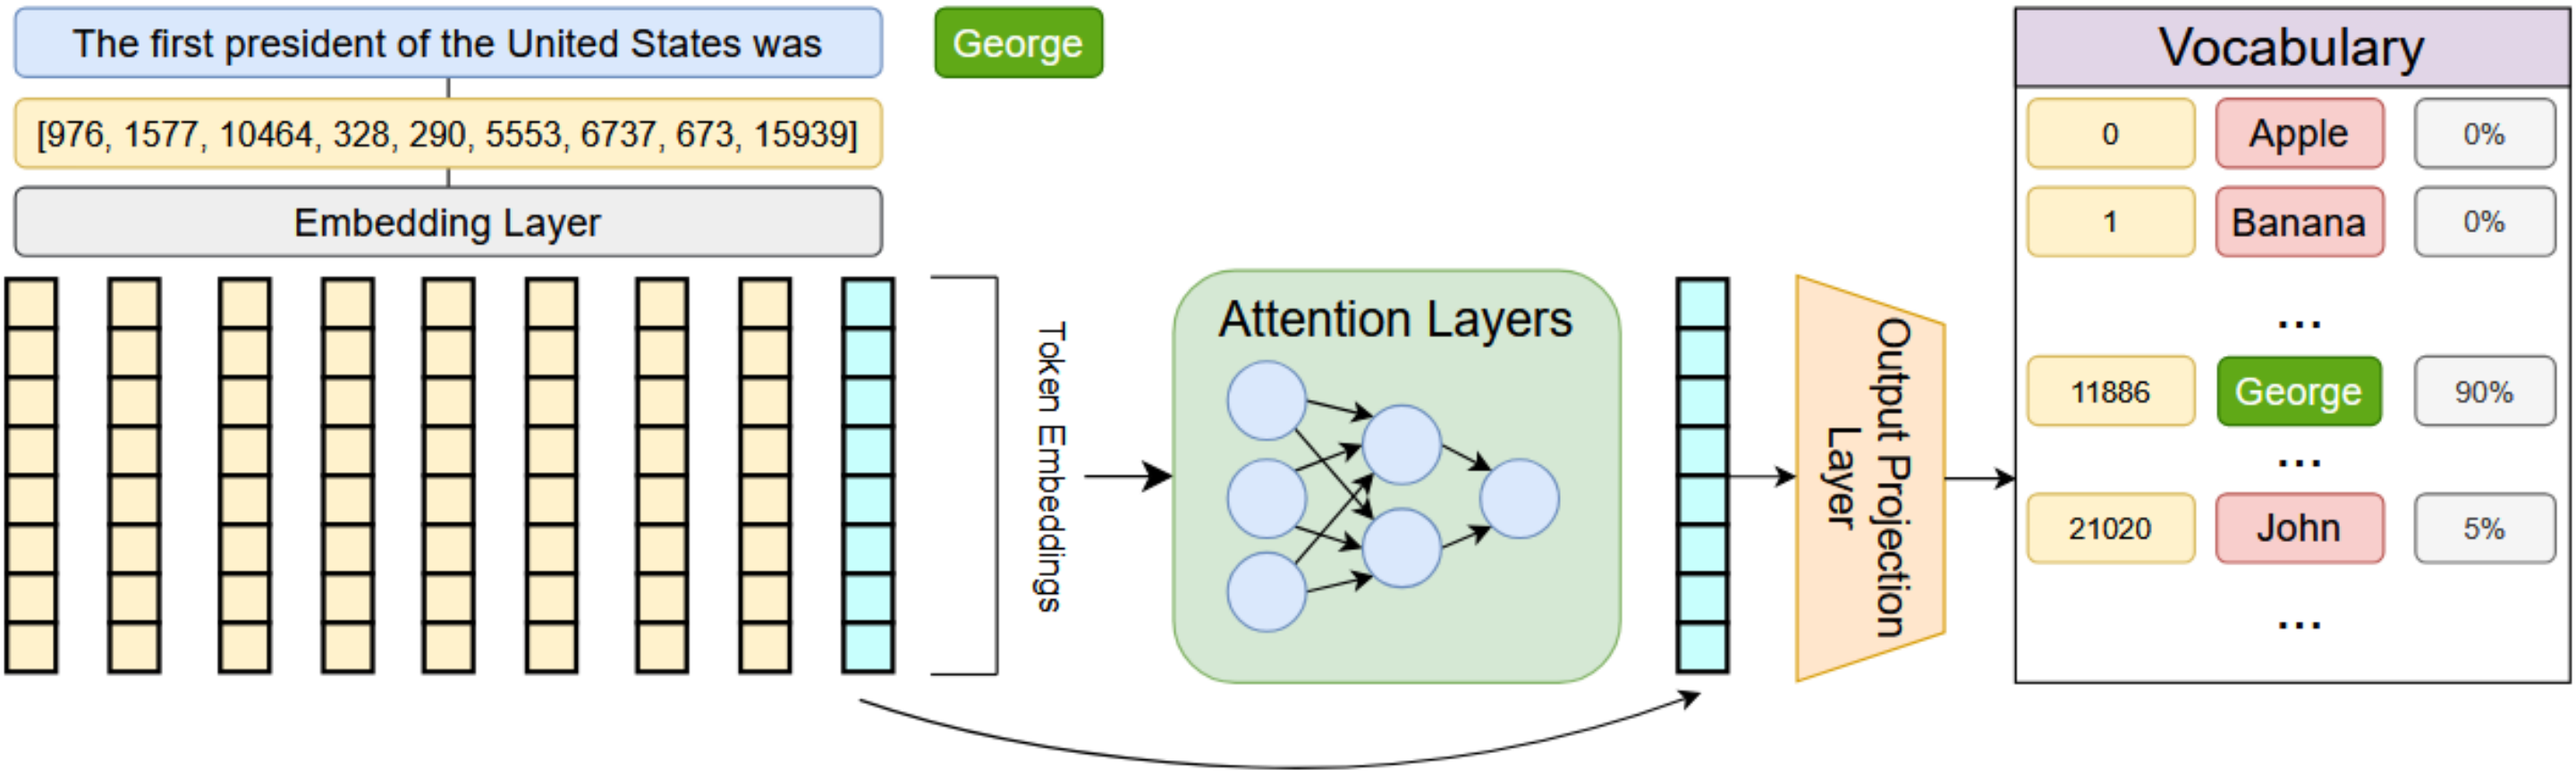
Training a Large Language Model –
Pre-Training the Transformer

- Creates a probability distribution for entire vocabulary
- Highest probability token is selected



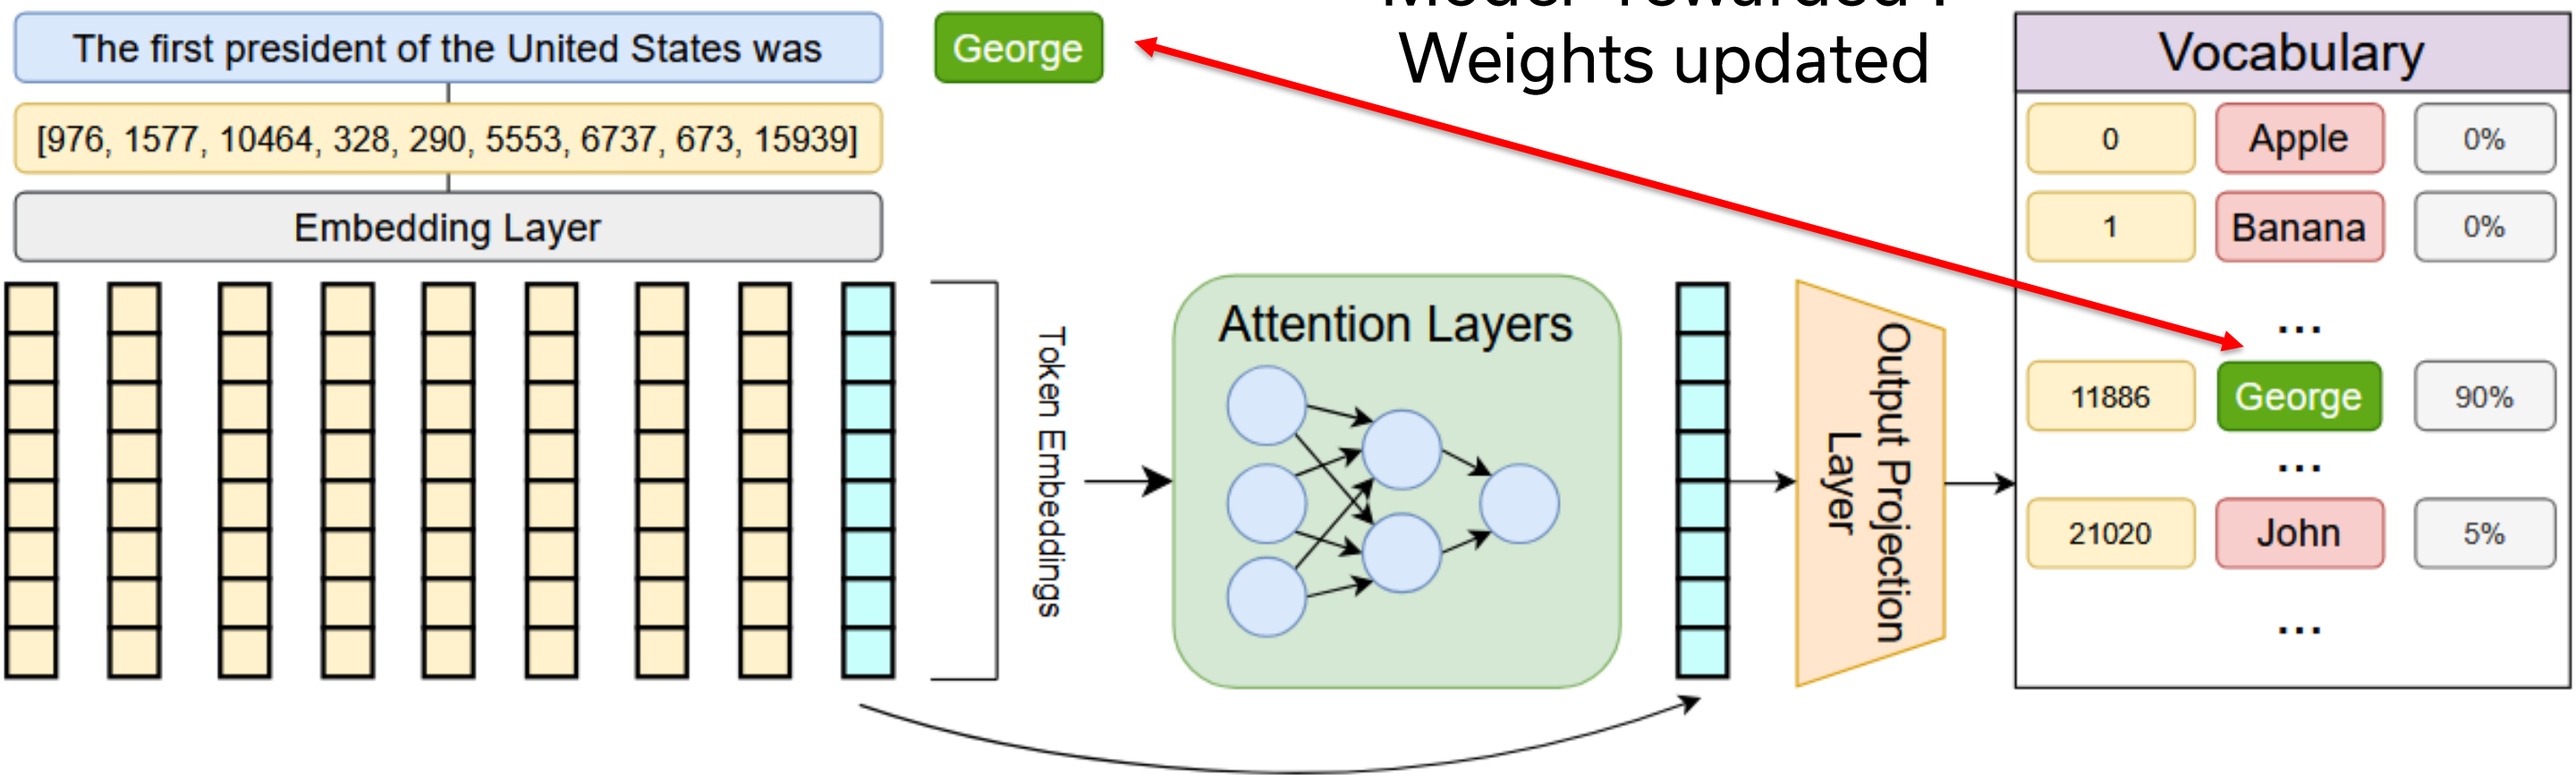
Training a Large Language Model –
Pre-Training the Transformer

- Creates a probability distribution for entire vocabulary
- Highest probability token is selected



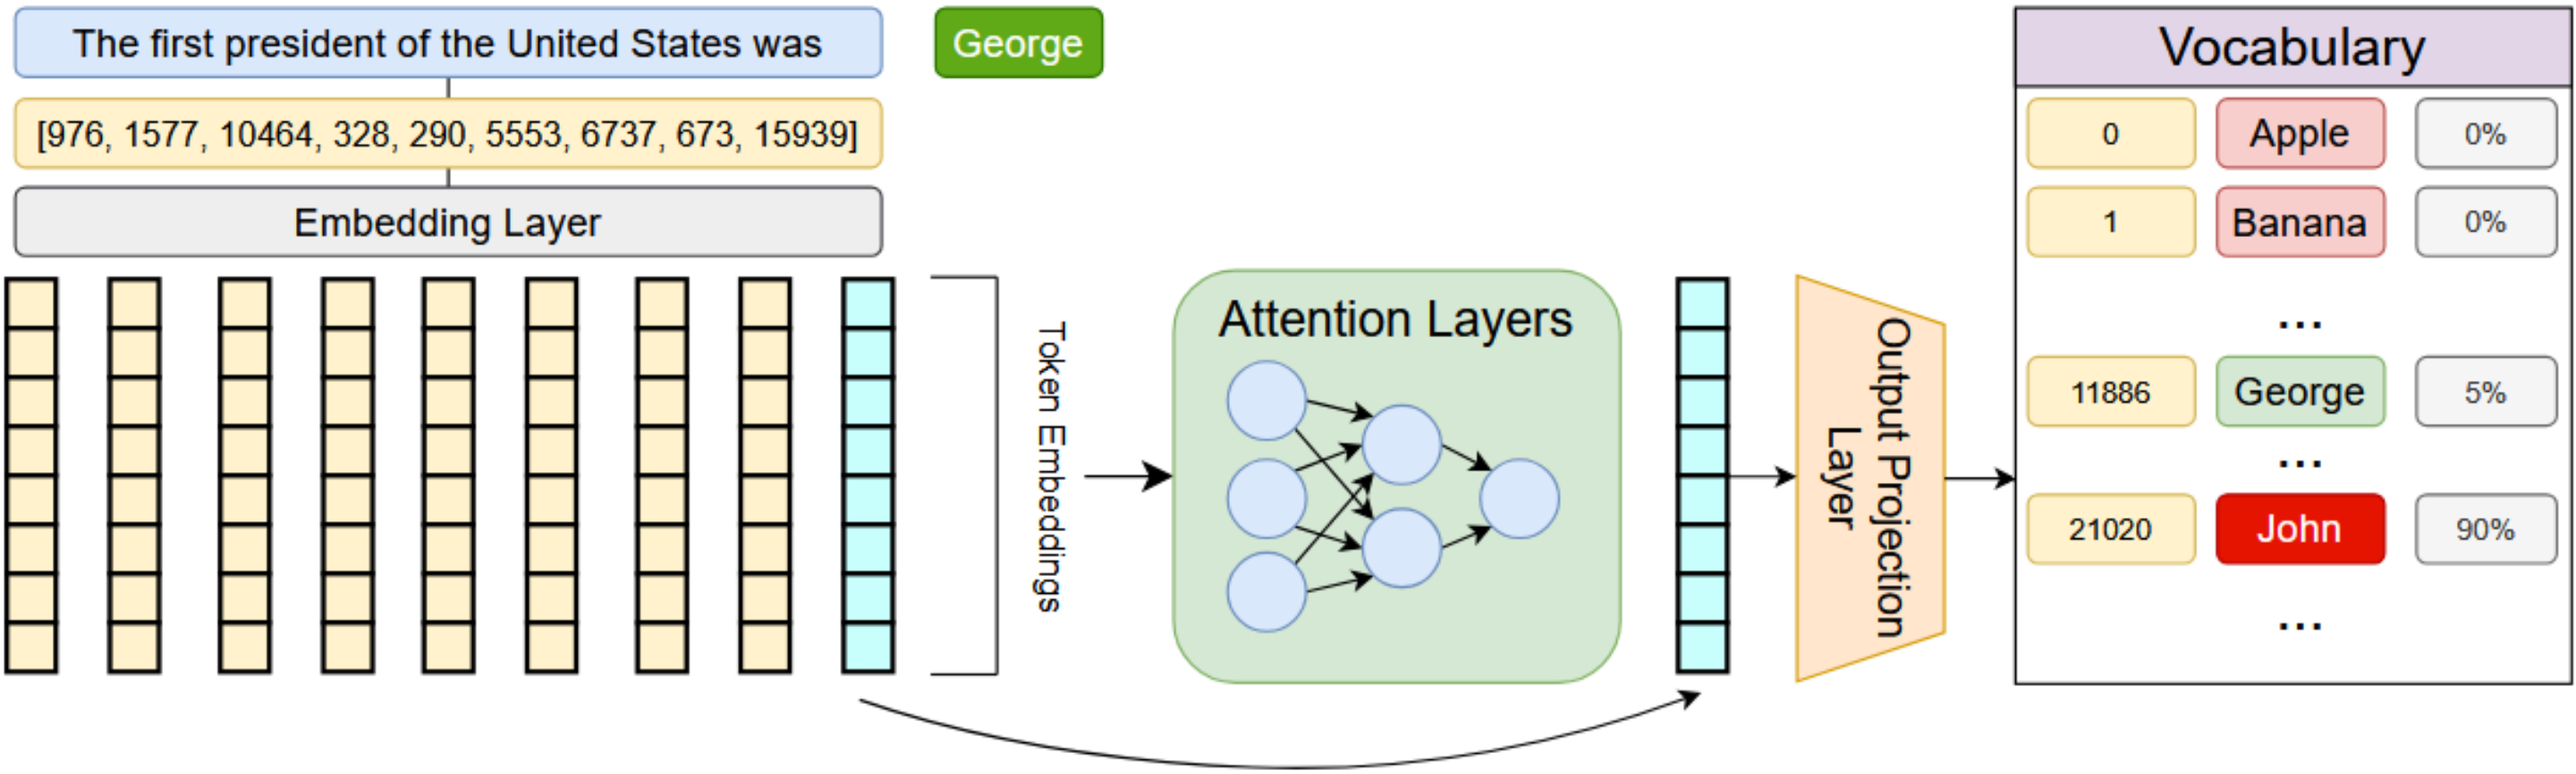
Training a Large Language Model –
Pre-Training the Transformer

- Creates a probability distribution for entire vocabulary
- Highest probability token is selected



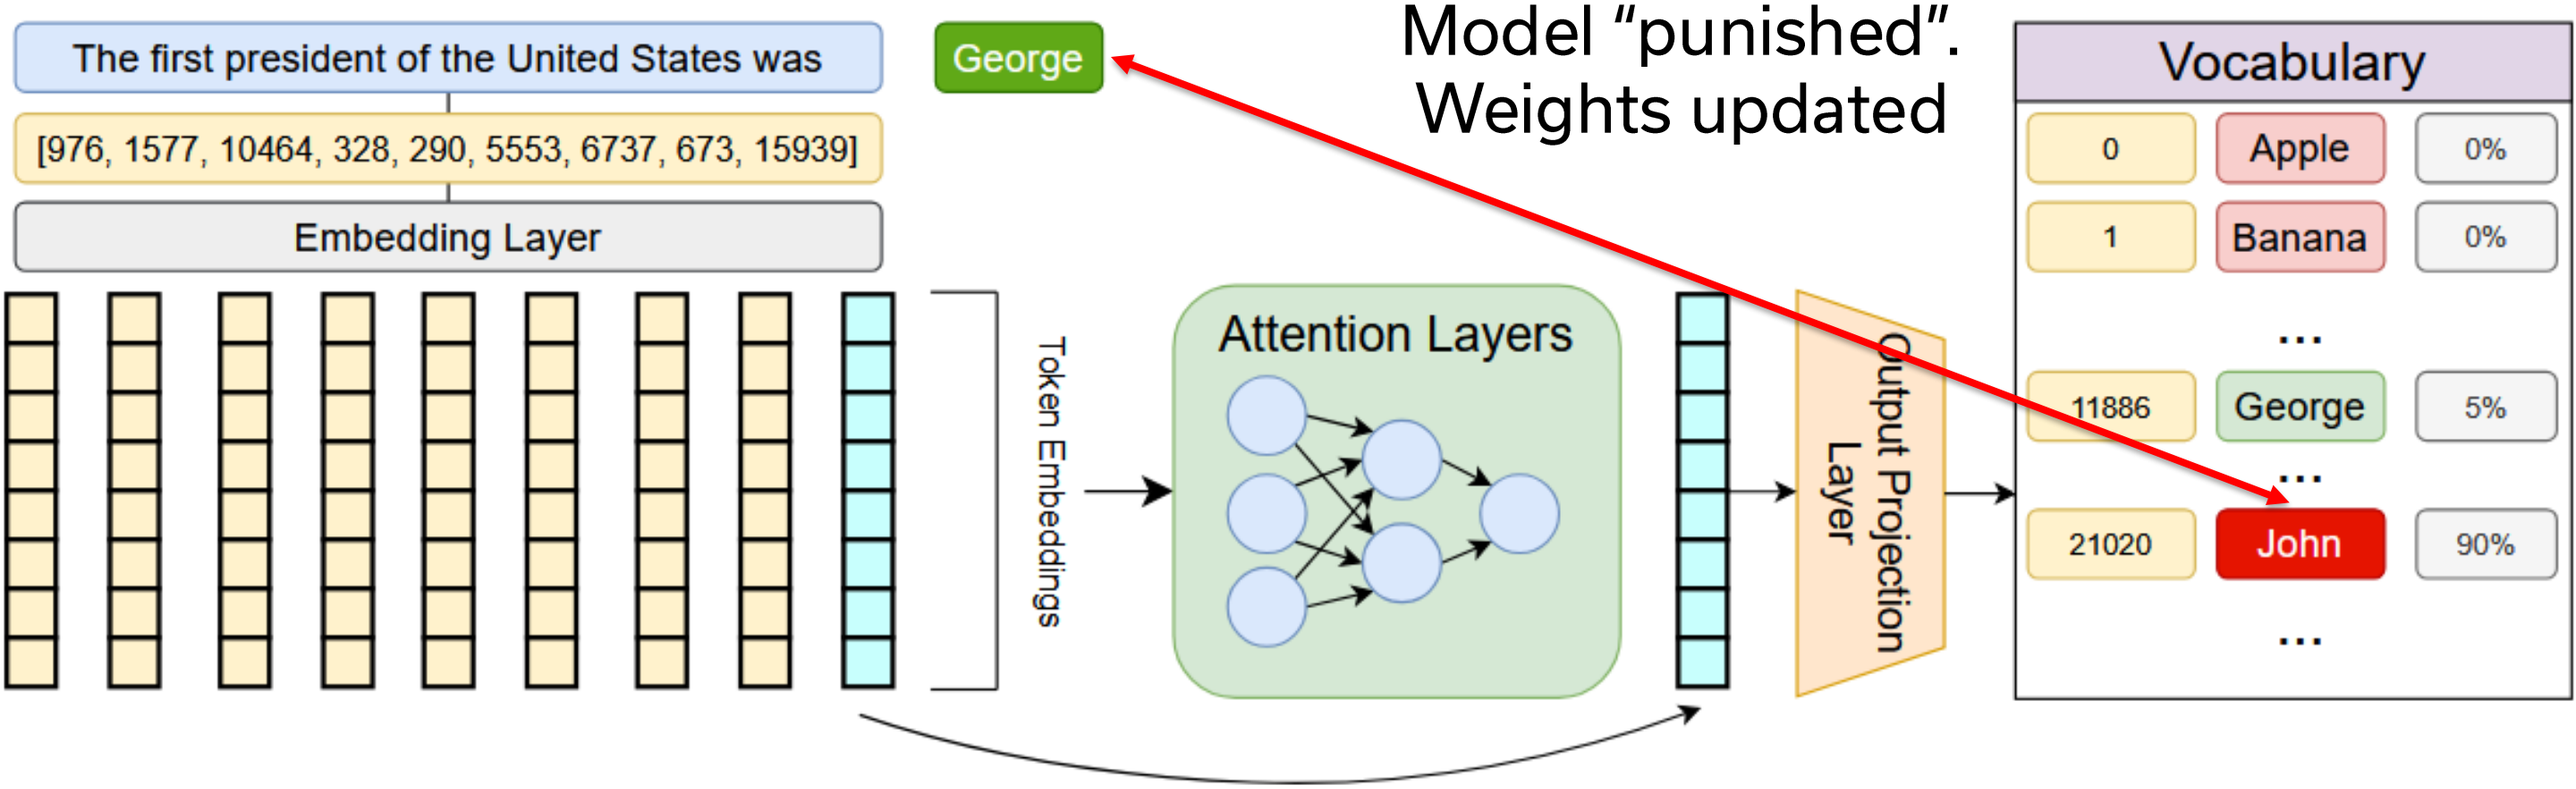
Training a Large Language Model –
Pre-Training the Transformer

- Creates a probability distribution for entire vocabulary
- Highest probability token is selected



Training a Large Language Model –
Pre-Training the Transformer

- Creates a probability distribution for entire vocabulary
- Highest probability token is selected

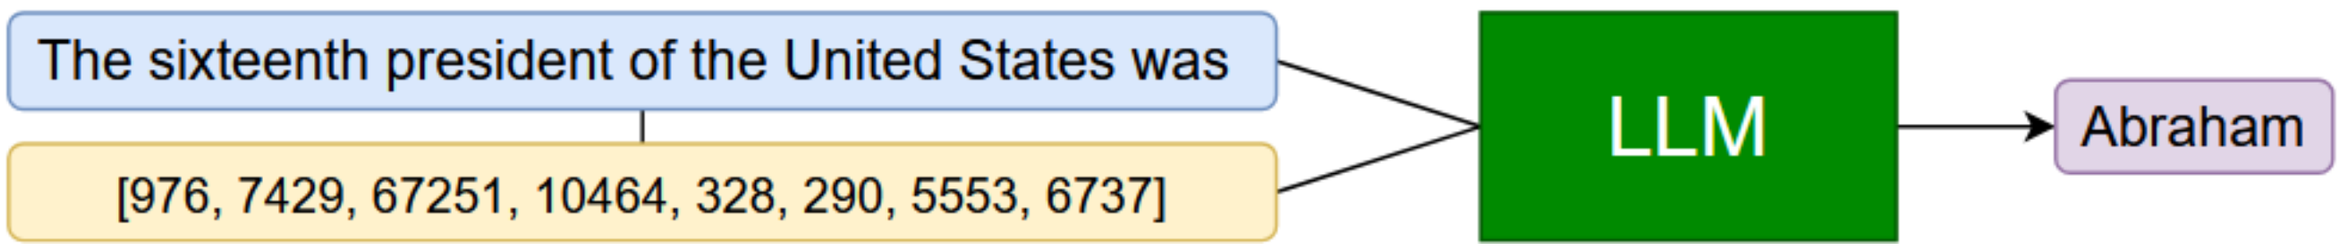


Training a Large Language Model –
What do we have now?

- For an input sequence, the model can now produce the next token until we tell it to stop!
- **Generative**
  - For any given input, mimic training data to complete sequence
- After enough training on accurate data, model gains “knowledge”

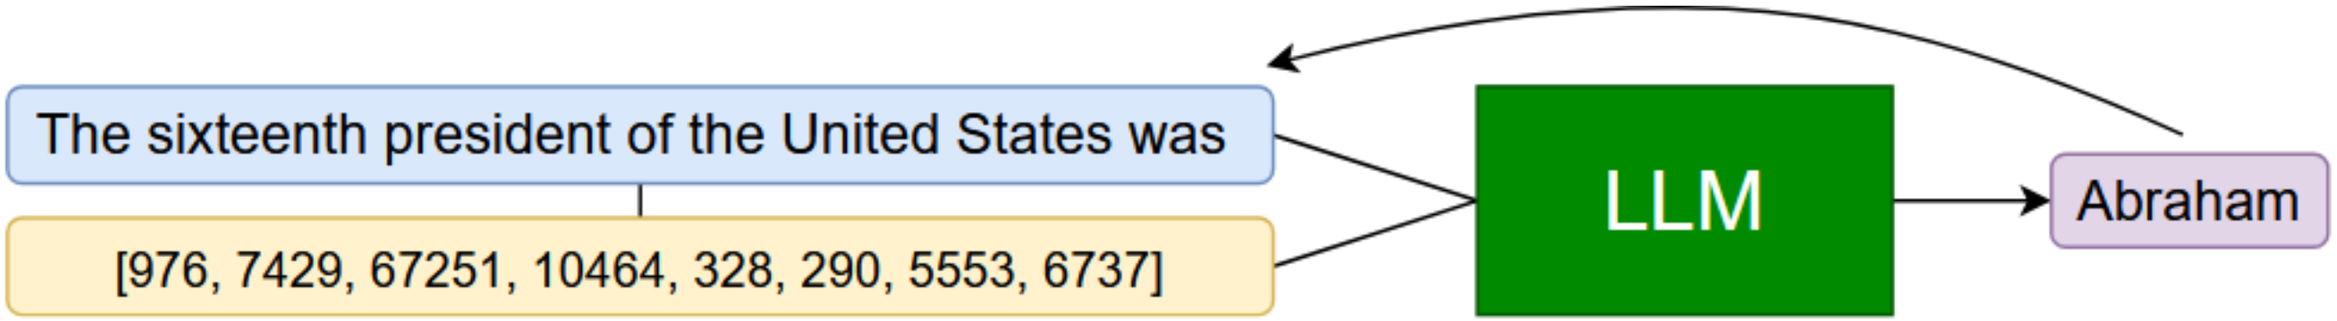
Training a Large Language Model –
What do we have now?

- For an input sequence, the model can now produce the next token until we tell it to stop!
- **Generative**
  - For any given input, mimic training data to complete sequence
- After enough training on accurate data, model gains “knowledge”



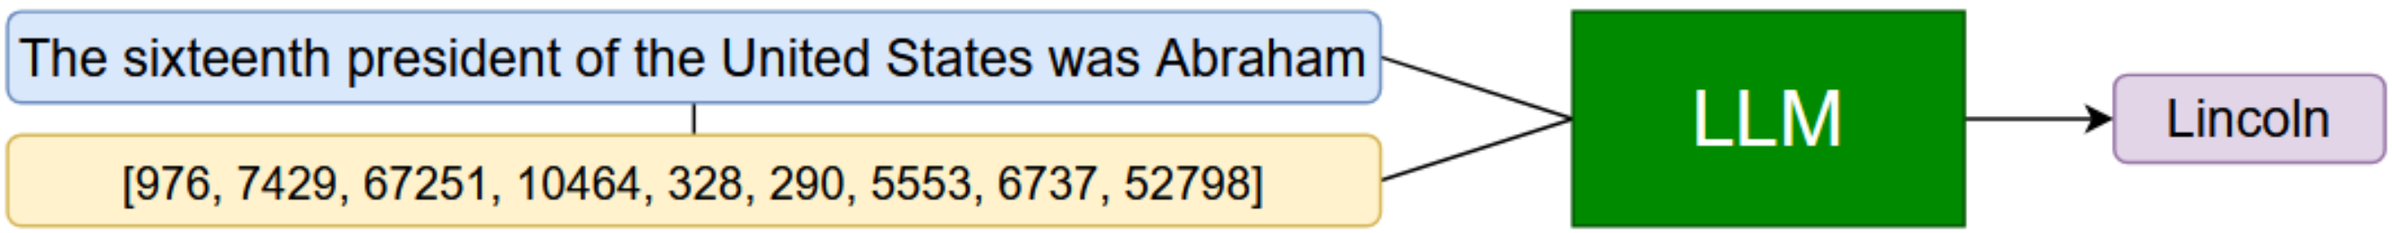
Training a Large Language Model –
What do we have now?

- For an input sequence, the model can now produce the next token until we tell it to stop!
- **Generative**
  - For any given input, mimic training data to complete sequence
- After enough training on accurate data, model gains “knowledge”



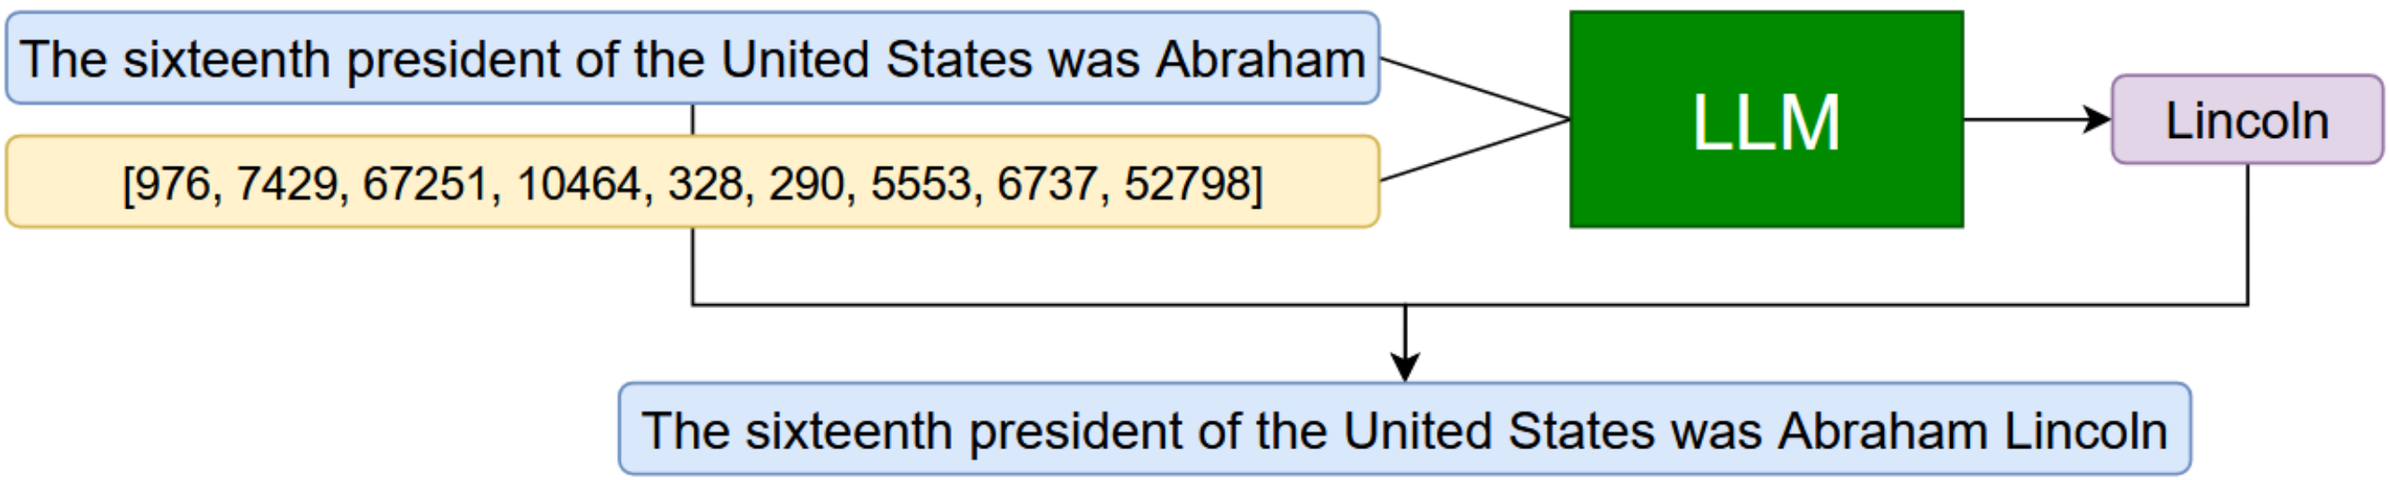
Training a Large Language Model –
What do we have now?

- For an input sequence, the model can now produce the next token until we tell it to stop!
- **Generative**
  - For any given input, mimic training data to complete sequence
- After enough training on accurate data, model gains “knowledge”



Training a Large Language Model –
What do we have now?

- For an input sequence, the model can now produce the next token until we tell it to stop!
- **Generative**
  - For any given input, mimic training data to complete sequence
- After enough training on accurate data, model gains “knowledge”



Training a Large Language Model – *What do we have now?*

Problem:

- Only performs sequence “completions”
- No conversational capacity
- Can’t serve as a chatbot/assistant

Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

- Following pre-training
- **Fine-tuning:** further training on a smaller, specialized dataset
- **Instruction Fine-tuning:** Fine-tune on dataset with "instruction" format
  - Turns of user input followed by model responses

Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

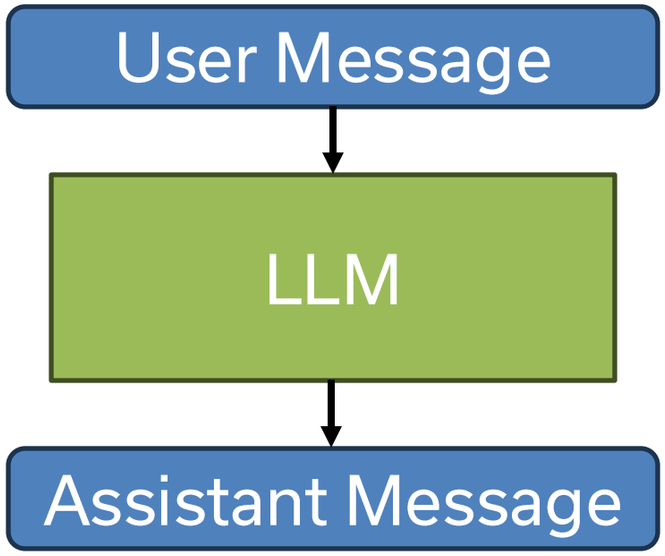
- Following pre-training
- **Fine-tuning:** further training on a smaller, specialized dataset
- **Instruction Fine-tuning:** Fine-tune on dataset with "instruction" format
  - Turns of user input followed by model responses

User Message

Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

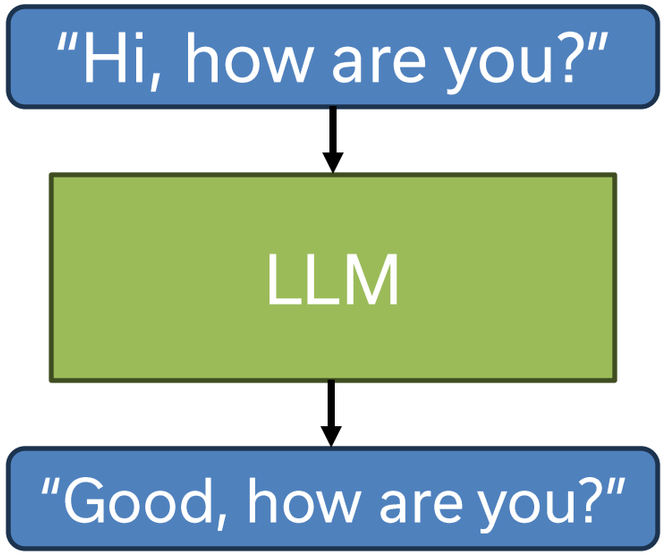
- Following pre-training
- Fine-tune on dataset with “instruction” format
  - Turns of user input followed by model responses
- **Fine-tuning**: further training on a smaller, specialized dataset



Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

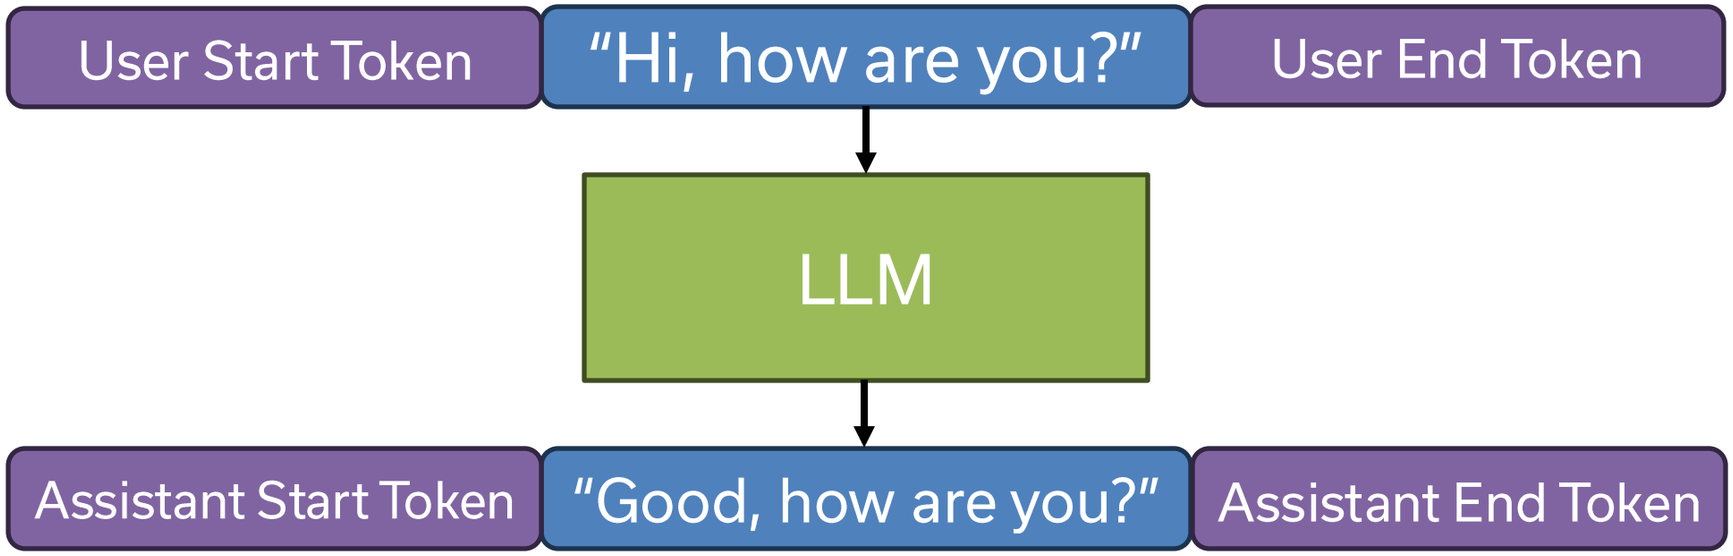
- Following pre-training
- Fine-tune on dataset with “instruction” format
  - Turns of user input followed by model responses
- **Fine-tuning**: further training on a smaller, specialized dataset



Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

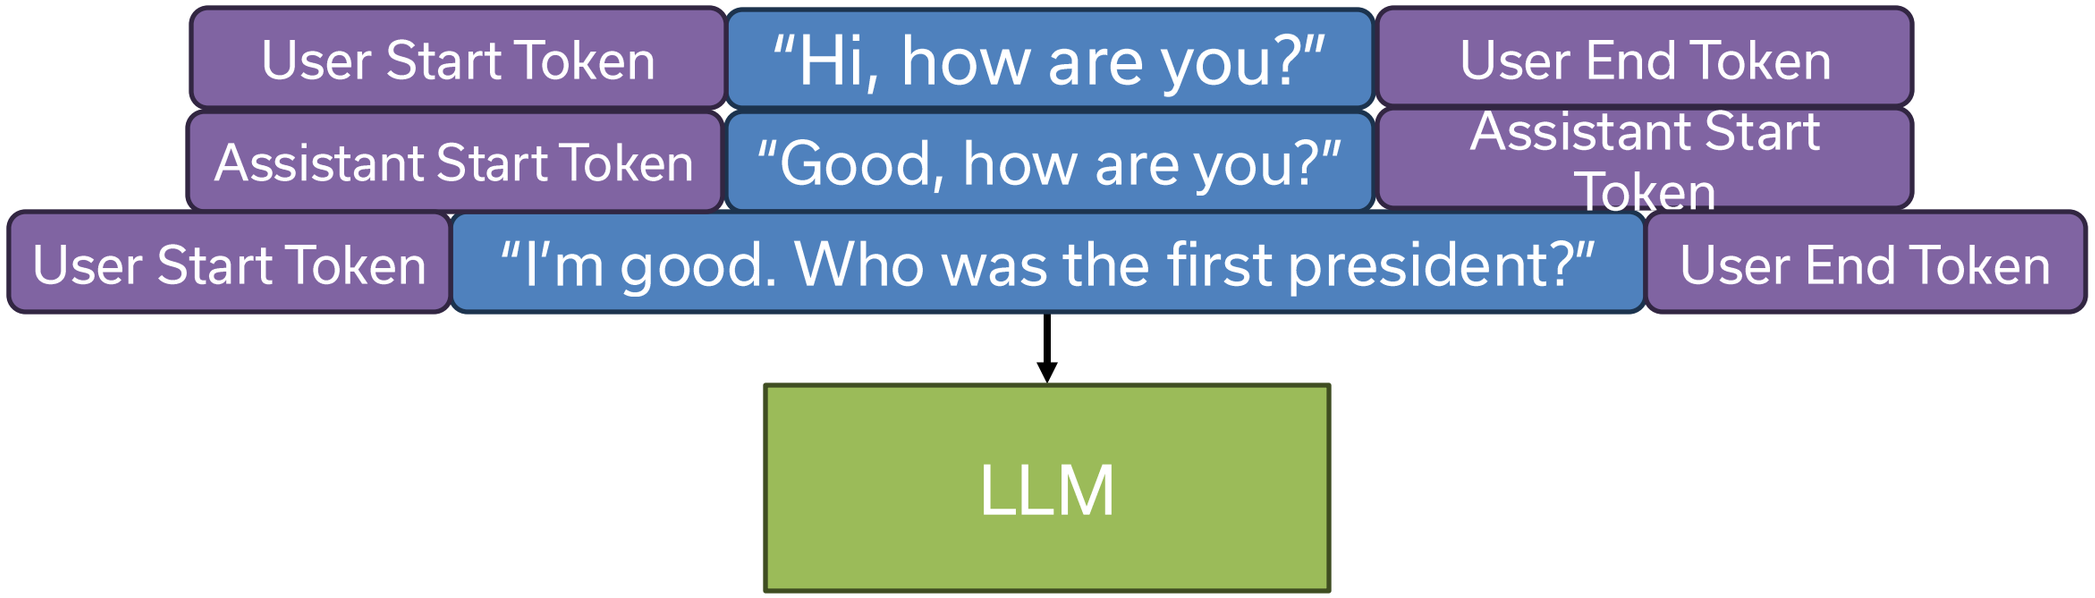
- Following pre-training
- Fine-tune on dataset with “instruction” format
  - Turns of user input followed by model responses
- **Fine-tuning**: further training on a smaller, specialized dataset



Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

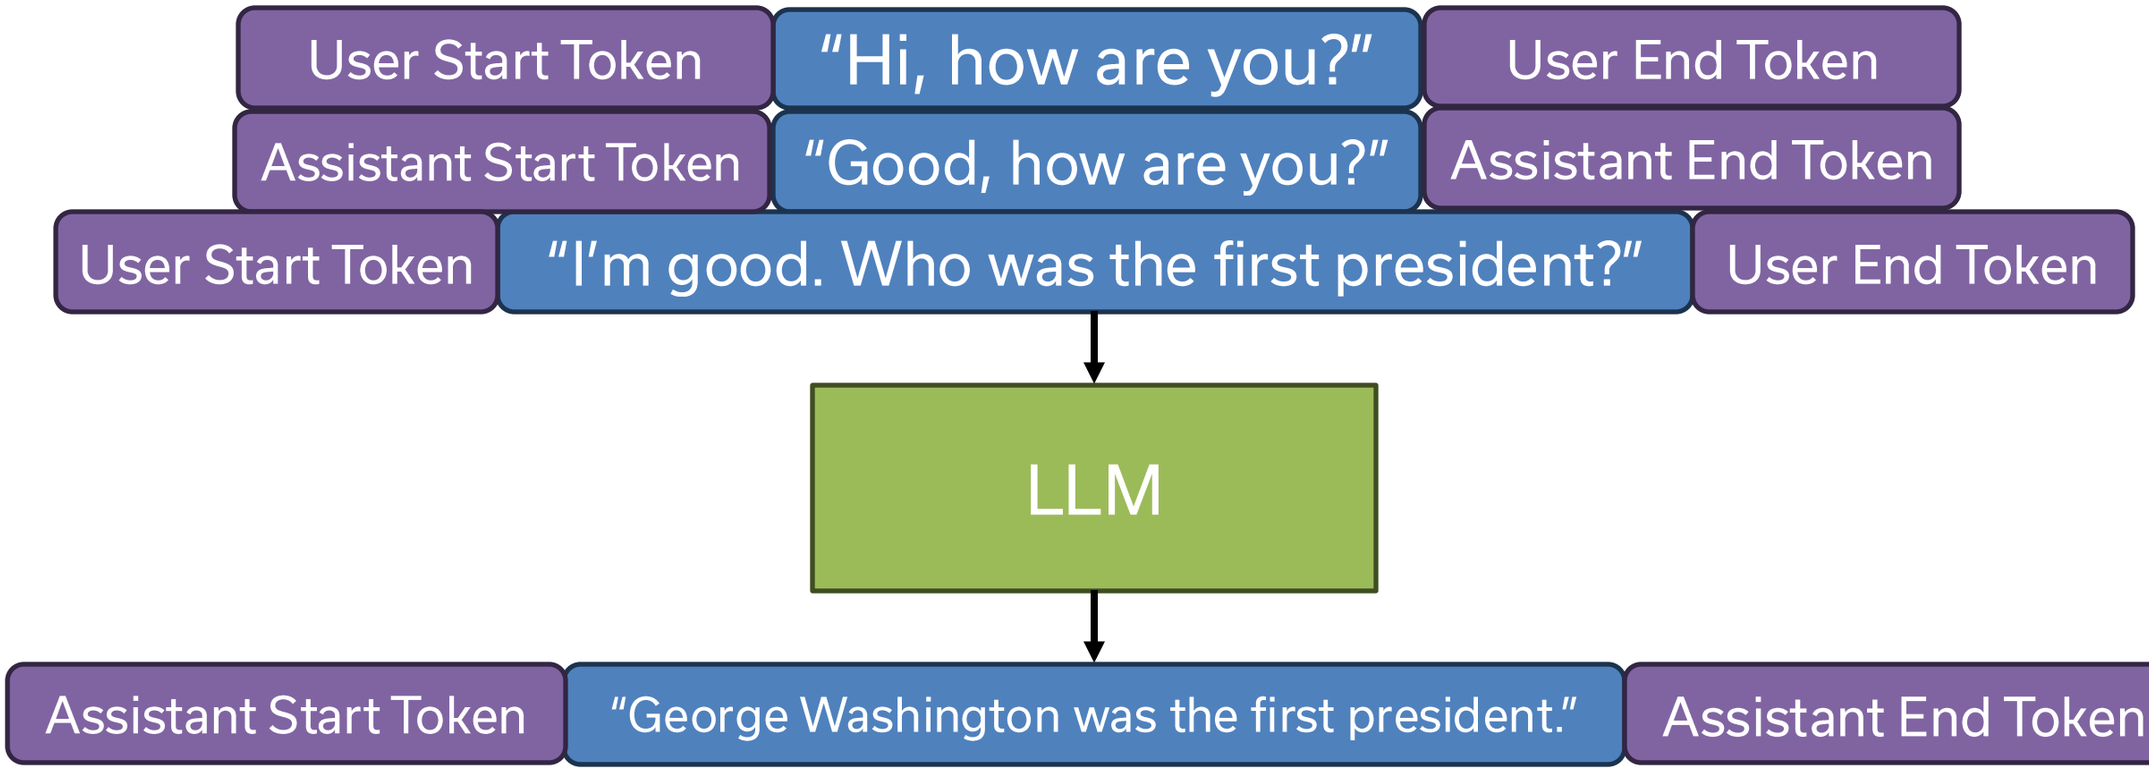
- Following pre-training
- Fine-tune on dataset with “instruction” format
  - Turns of user input followed by model responses



Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

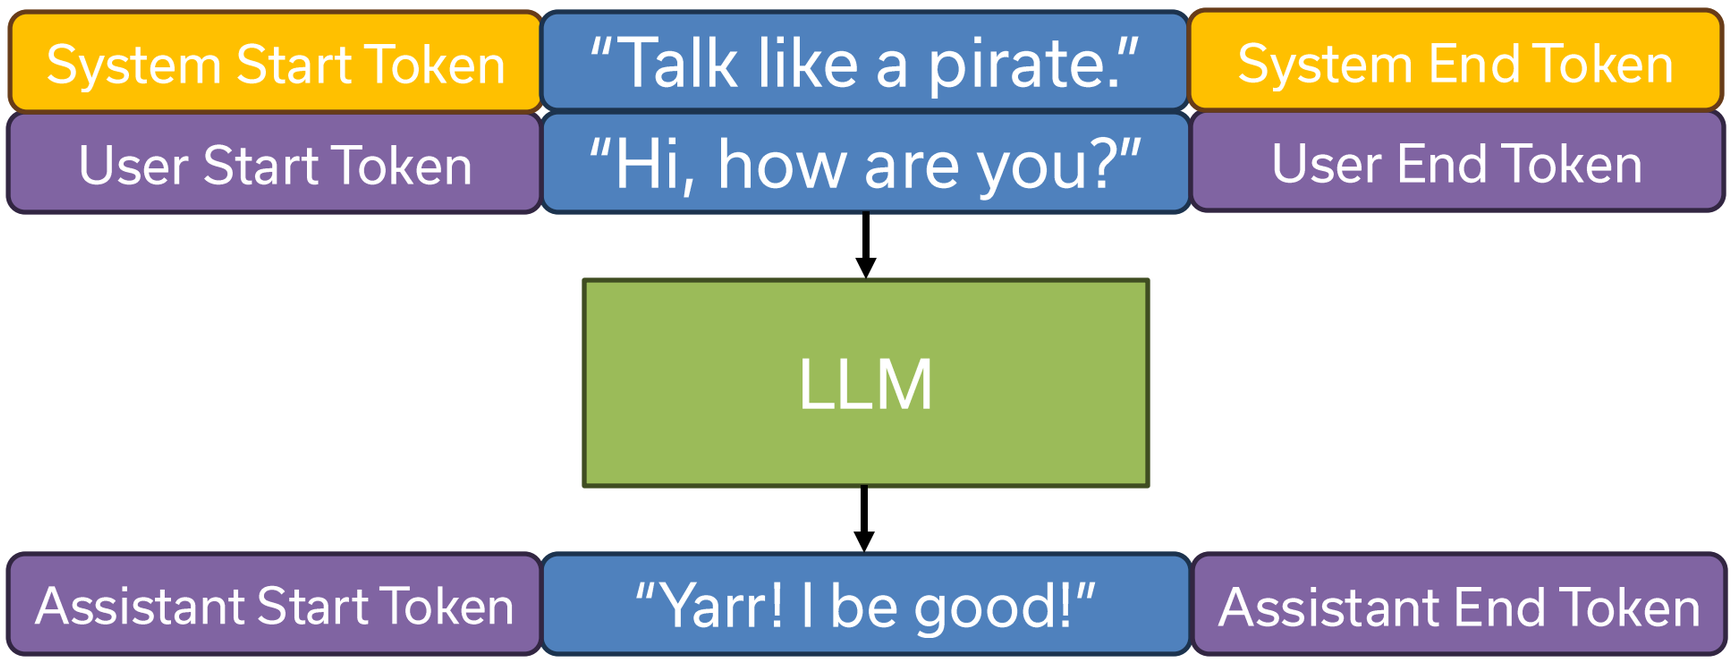
- Following pre-training
- Fine-tune on dataset with “instruction” format
  - Turns of user input followed by model responses



Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

- Following pre-training
- Fine-tune on dataset with “instruction” format
  - Turns of user input followed by model responses



Training a Large Language Model – *Fine-Tuning*

Instruction Fine-Tuning

- Following pre-training
- Fine-tune on dataset with “instruction” format
  - Turns of user input followed by model responses
- Create a dataset with messages formatted like this
- LLM response set as answers to questions, responses to instructions

Training a Large Language Model – *Fine-Tuning*

stackoverflow

Delete an element from a dictionary

Asked 14 years, 8 months ago Modified 9 months ago Viewed 2.9m times

2245

1. How do I delete an item from a dictionary in Python?
2. **Without** modifying the original dictionary, how do I obtain **another** dictionary with the item removed?

See also How can I remove a key from a Python dictionary? for the specific issue of removing an item (by key) that may not already be present.

python dictionary del

⟵ Question

Training a Large Language Model – *Fine-Tuning*

stackoverflow

Delete an element from a dictionary

Asked 14 years, 8 months ago Modified 9 months ago Viewed 2.9m times

2245

1. How do I delete an item from a dictionary in Python?
2. **Without** modifying the original dictionary, how do I obtain **another** dictionary with the item removed?

See also How can I remove a key from a Python dictionary? for the specific issue of removing an item (by key) that may not already be present.

python dictionary del

← Question

Response →

2571

The `del` statement removes an element:

```
del d[key]
```

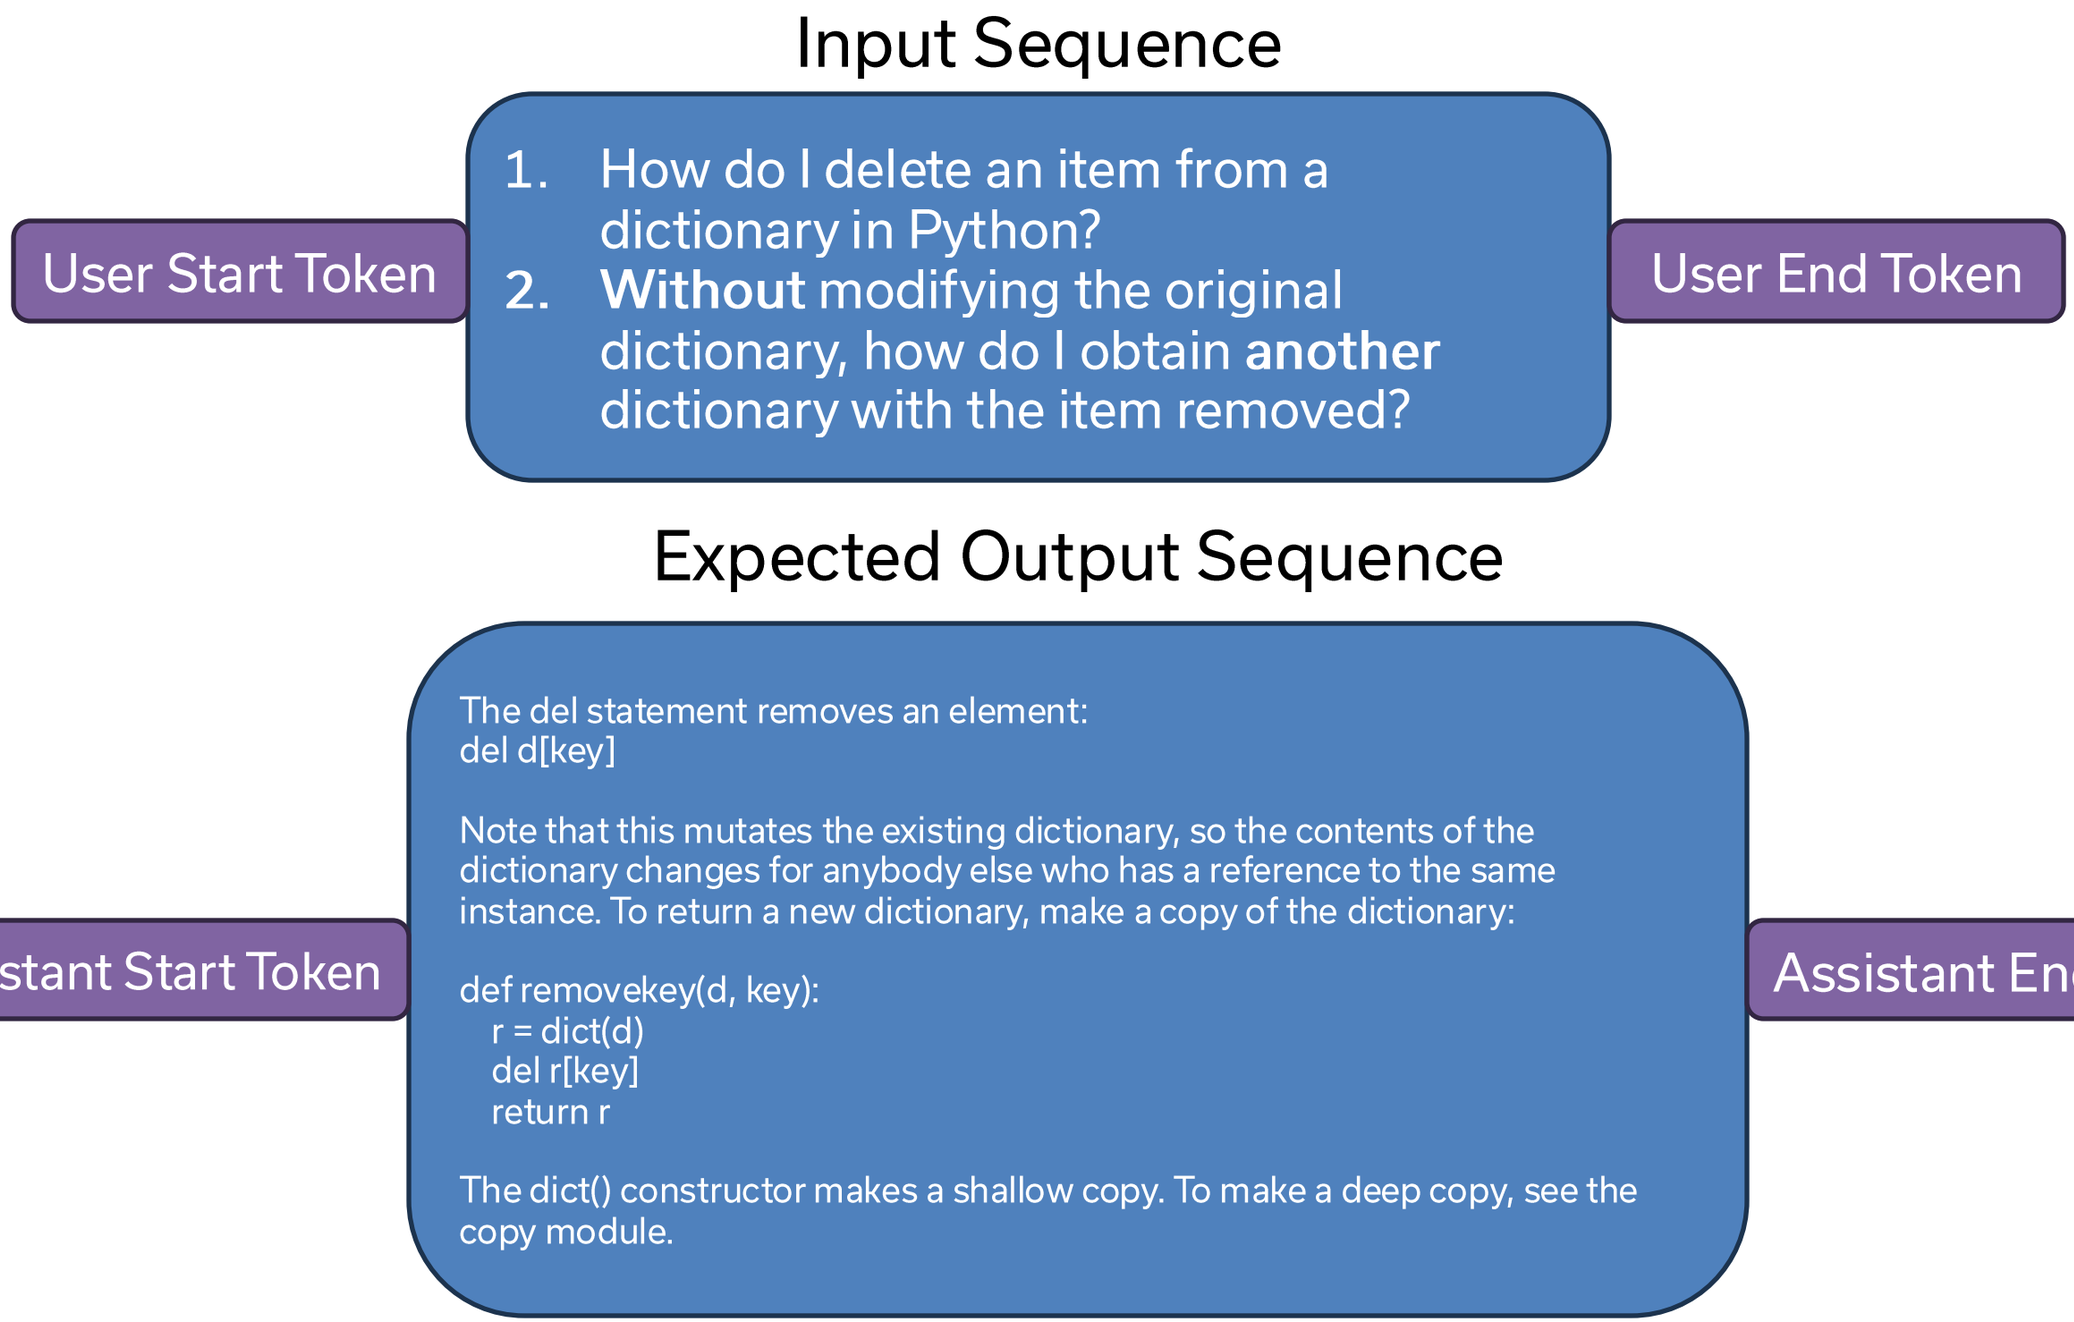
Note that this mutates the existing dictionary, so the contents of the dictionary changes for anybody else who has a reference to the same instance. To return a *new* dictionary, make a copy of the dictionary:

```
def removekey(d, key):
    r = dict(d)
    del r[key]
    return r
```

The `dict()` constructor makes a *shallow copy*. To make a deep copy, see the `copy` module.

NAIRR Pilot KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM ICPE Higher Education Matters University of Kentucky

Training a Large Language Model – *Fine-Tuning*

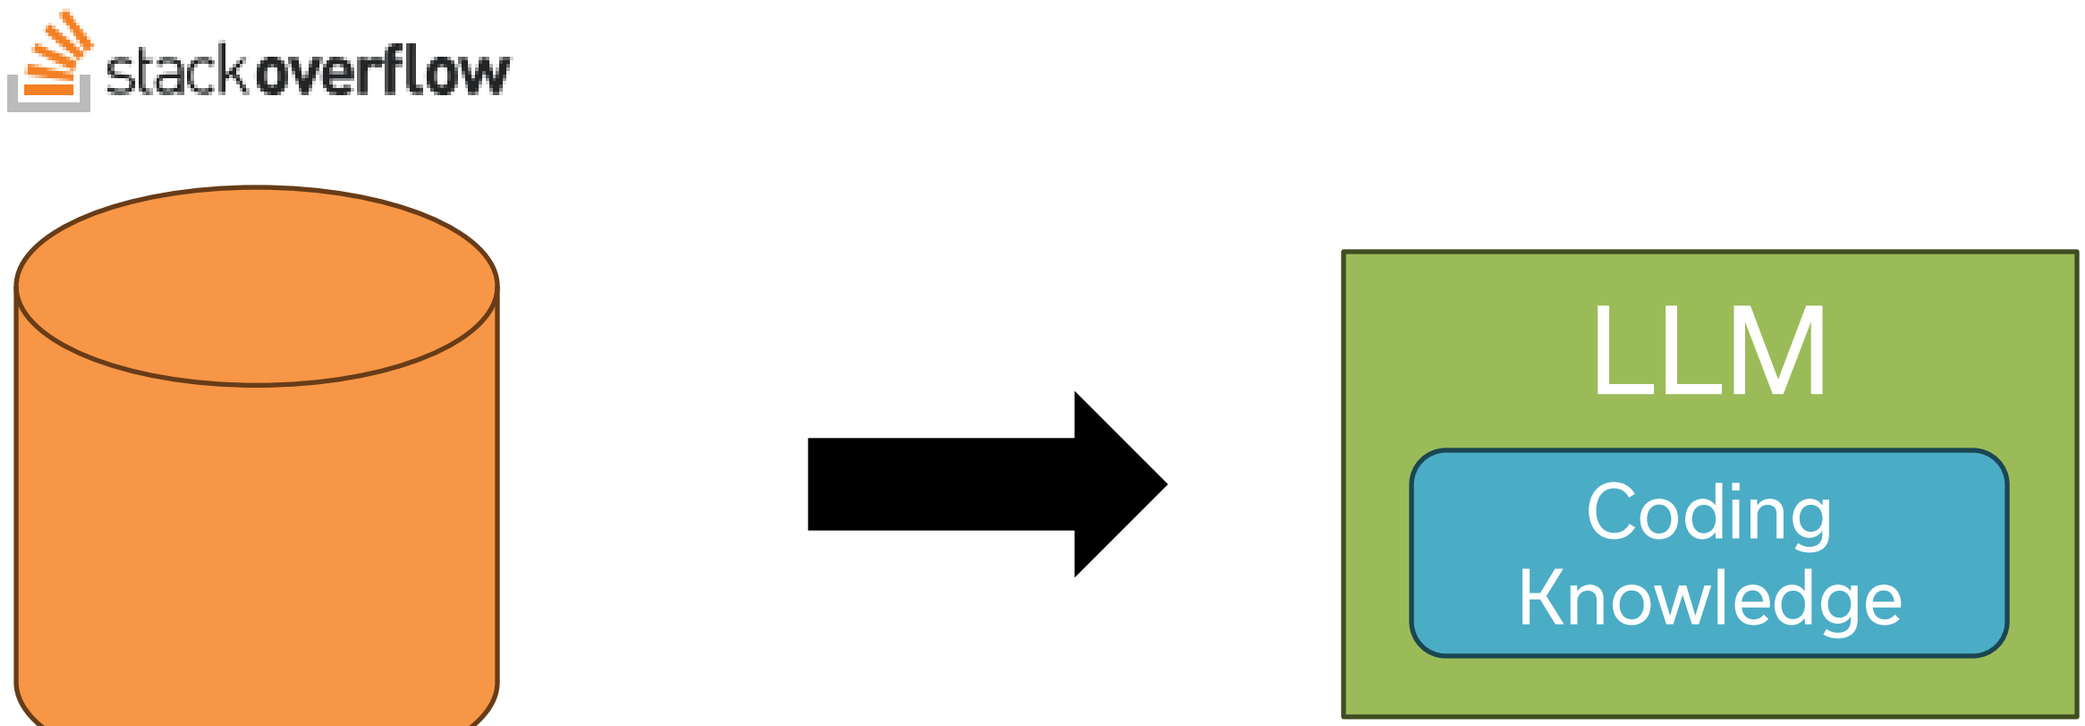


Training a Large Language Model

- After training at scale, you get modern LLMs

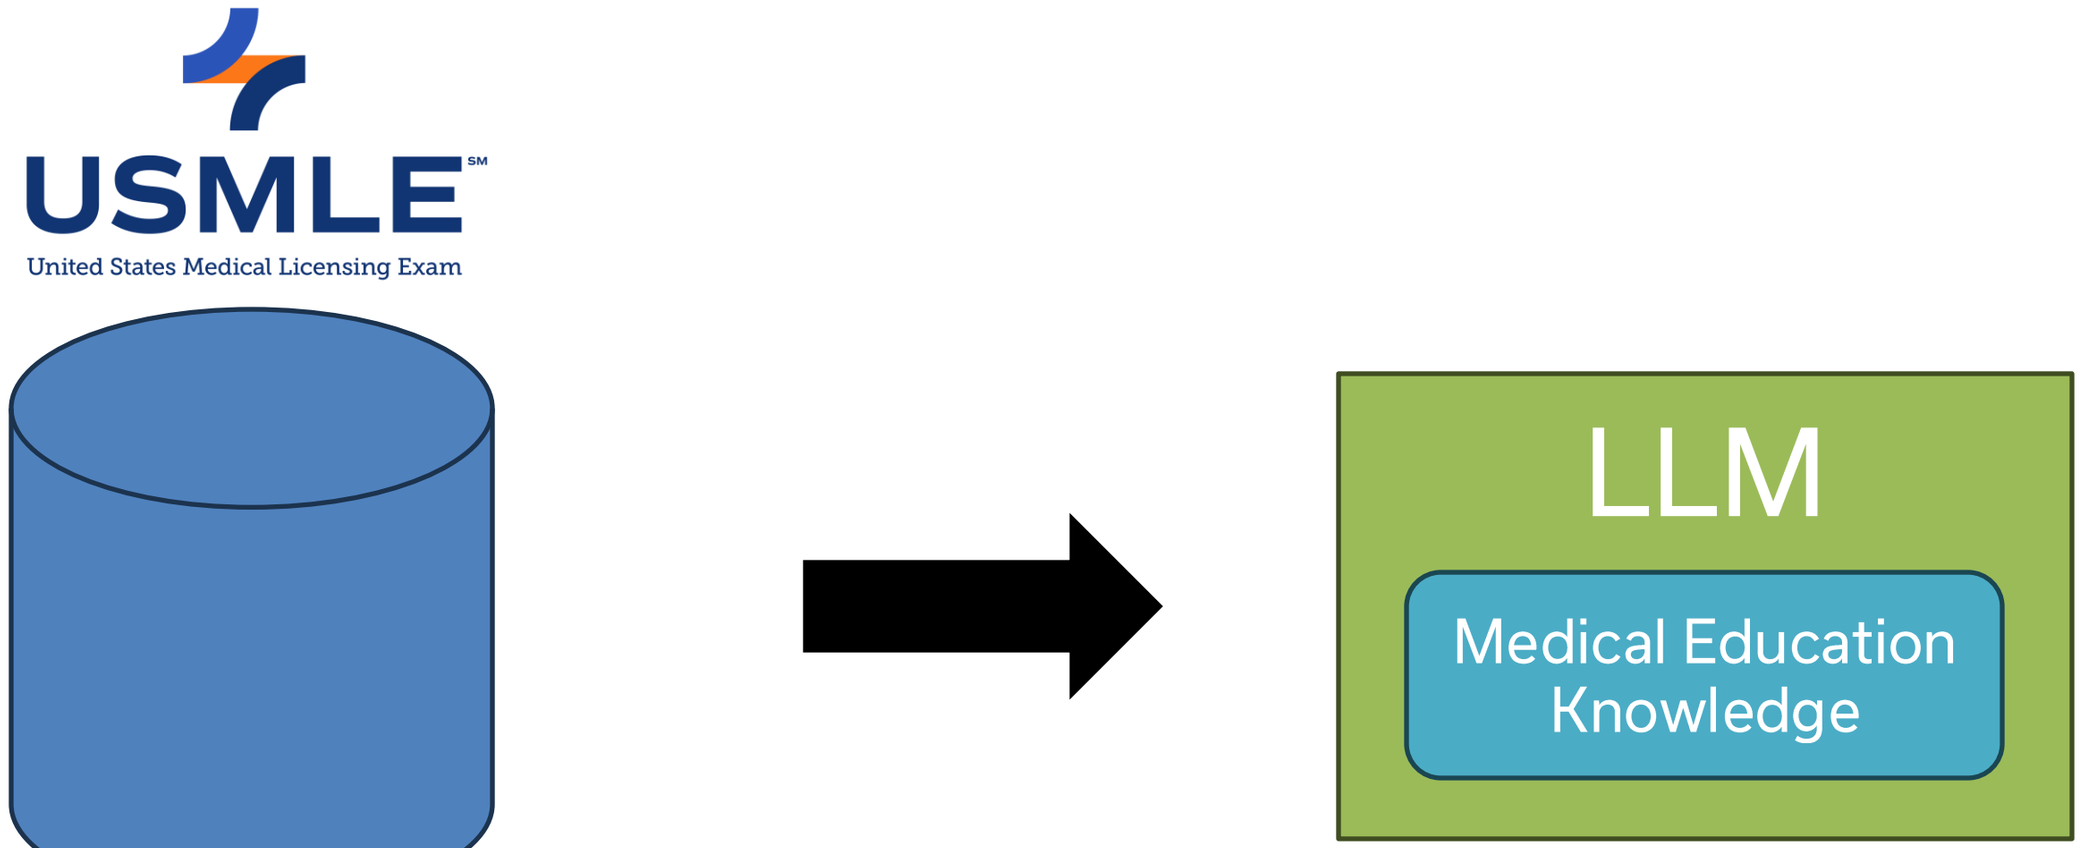
Training a Large Language Model

- After training at scale, you get modern LLMs



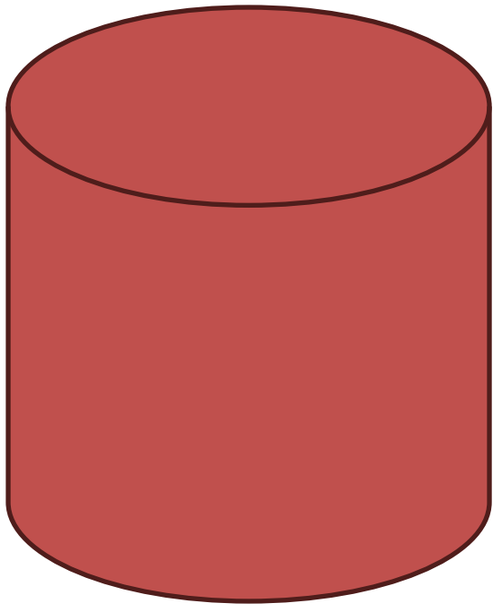
Training a Large Language Model

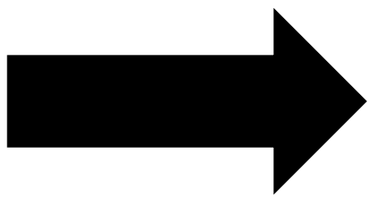
- After training at scale, you get modern LLMs



Training a Large Language Model

- After training at scale, you get modern LLMs

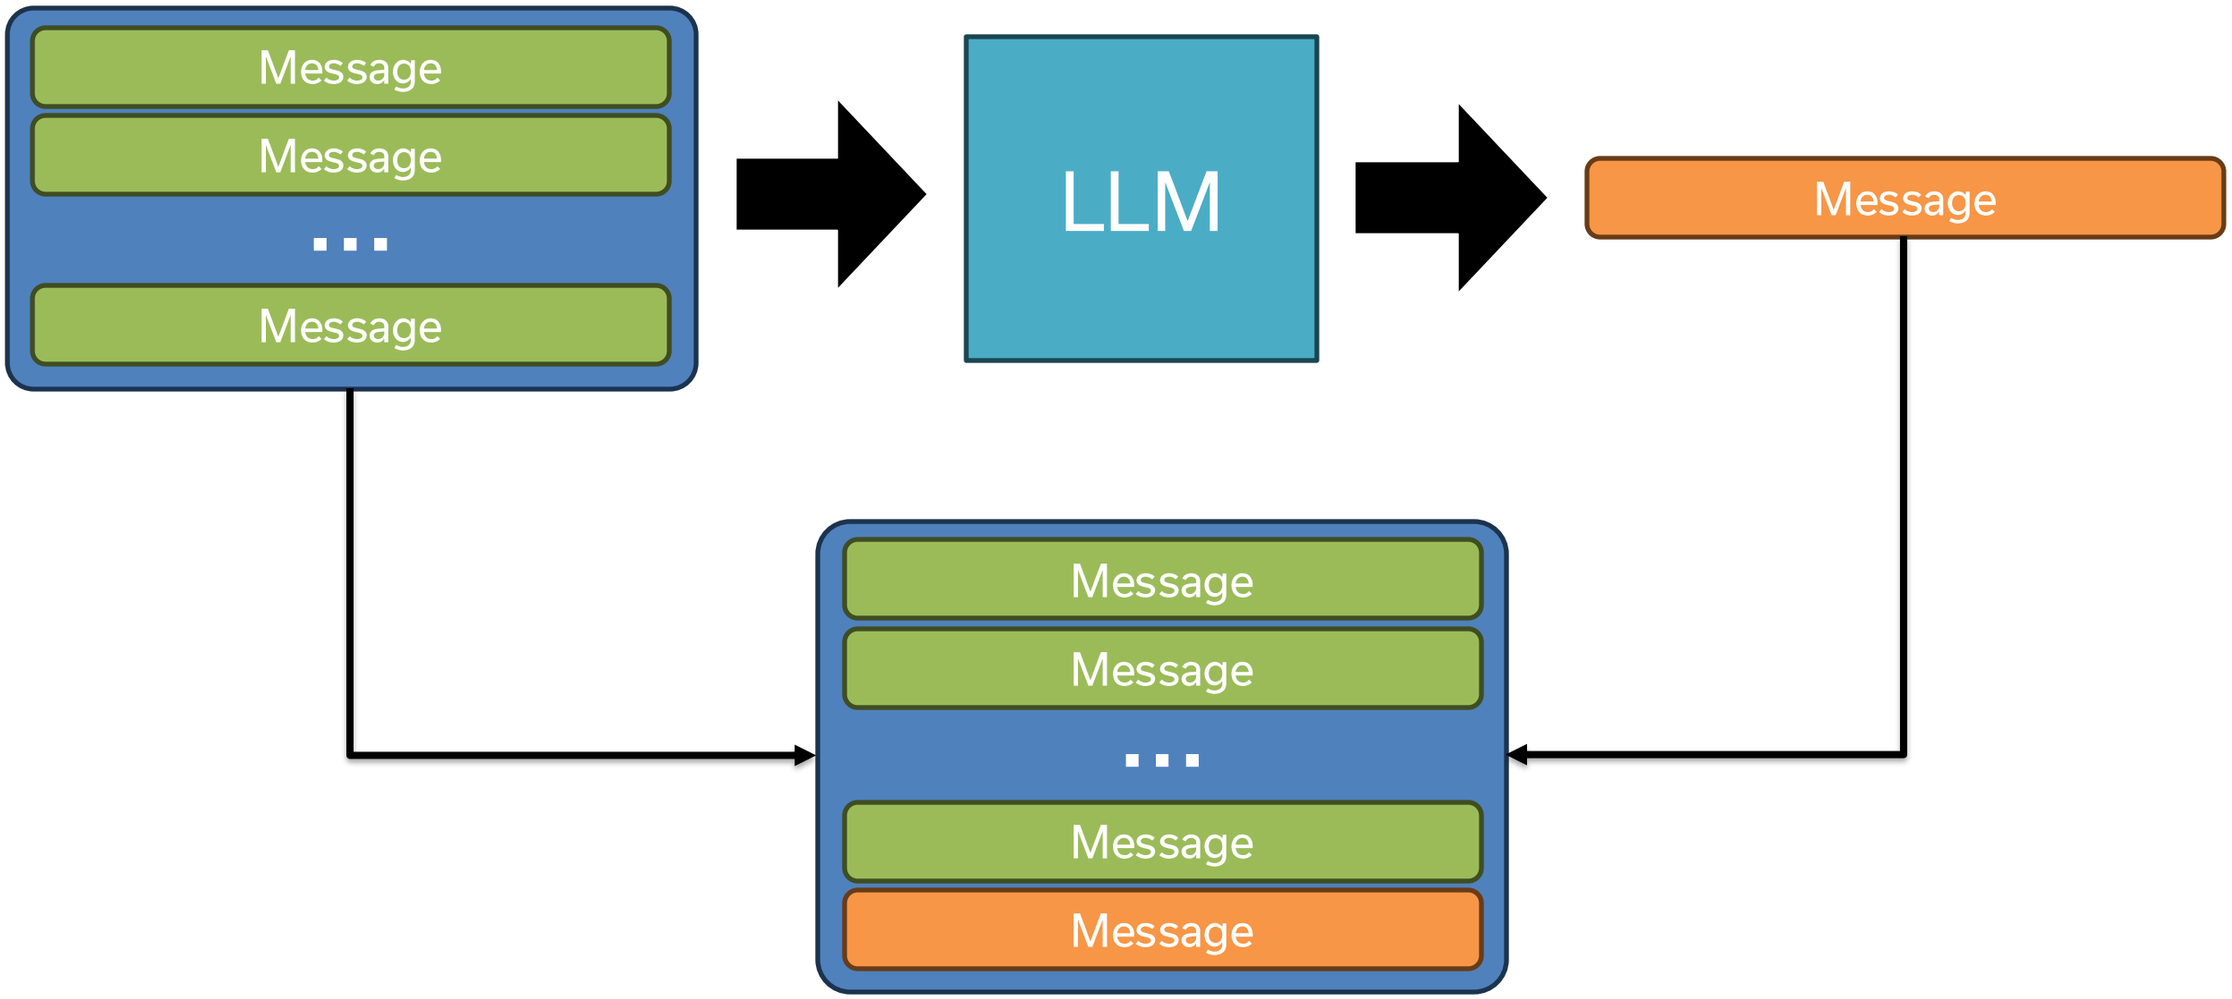




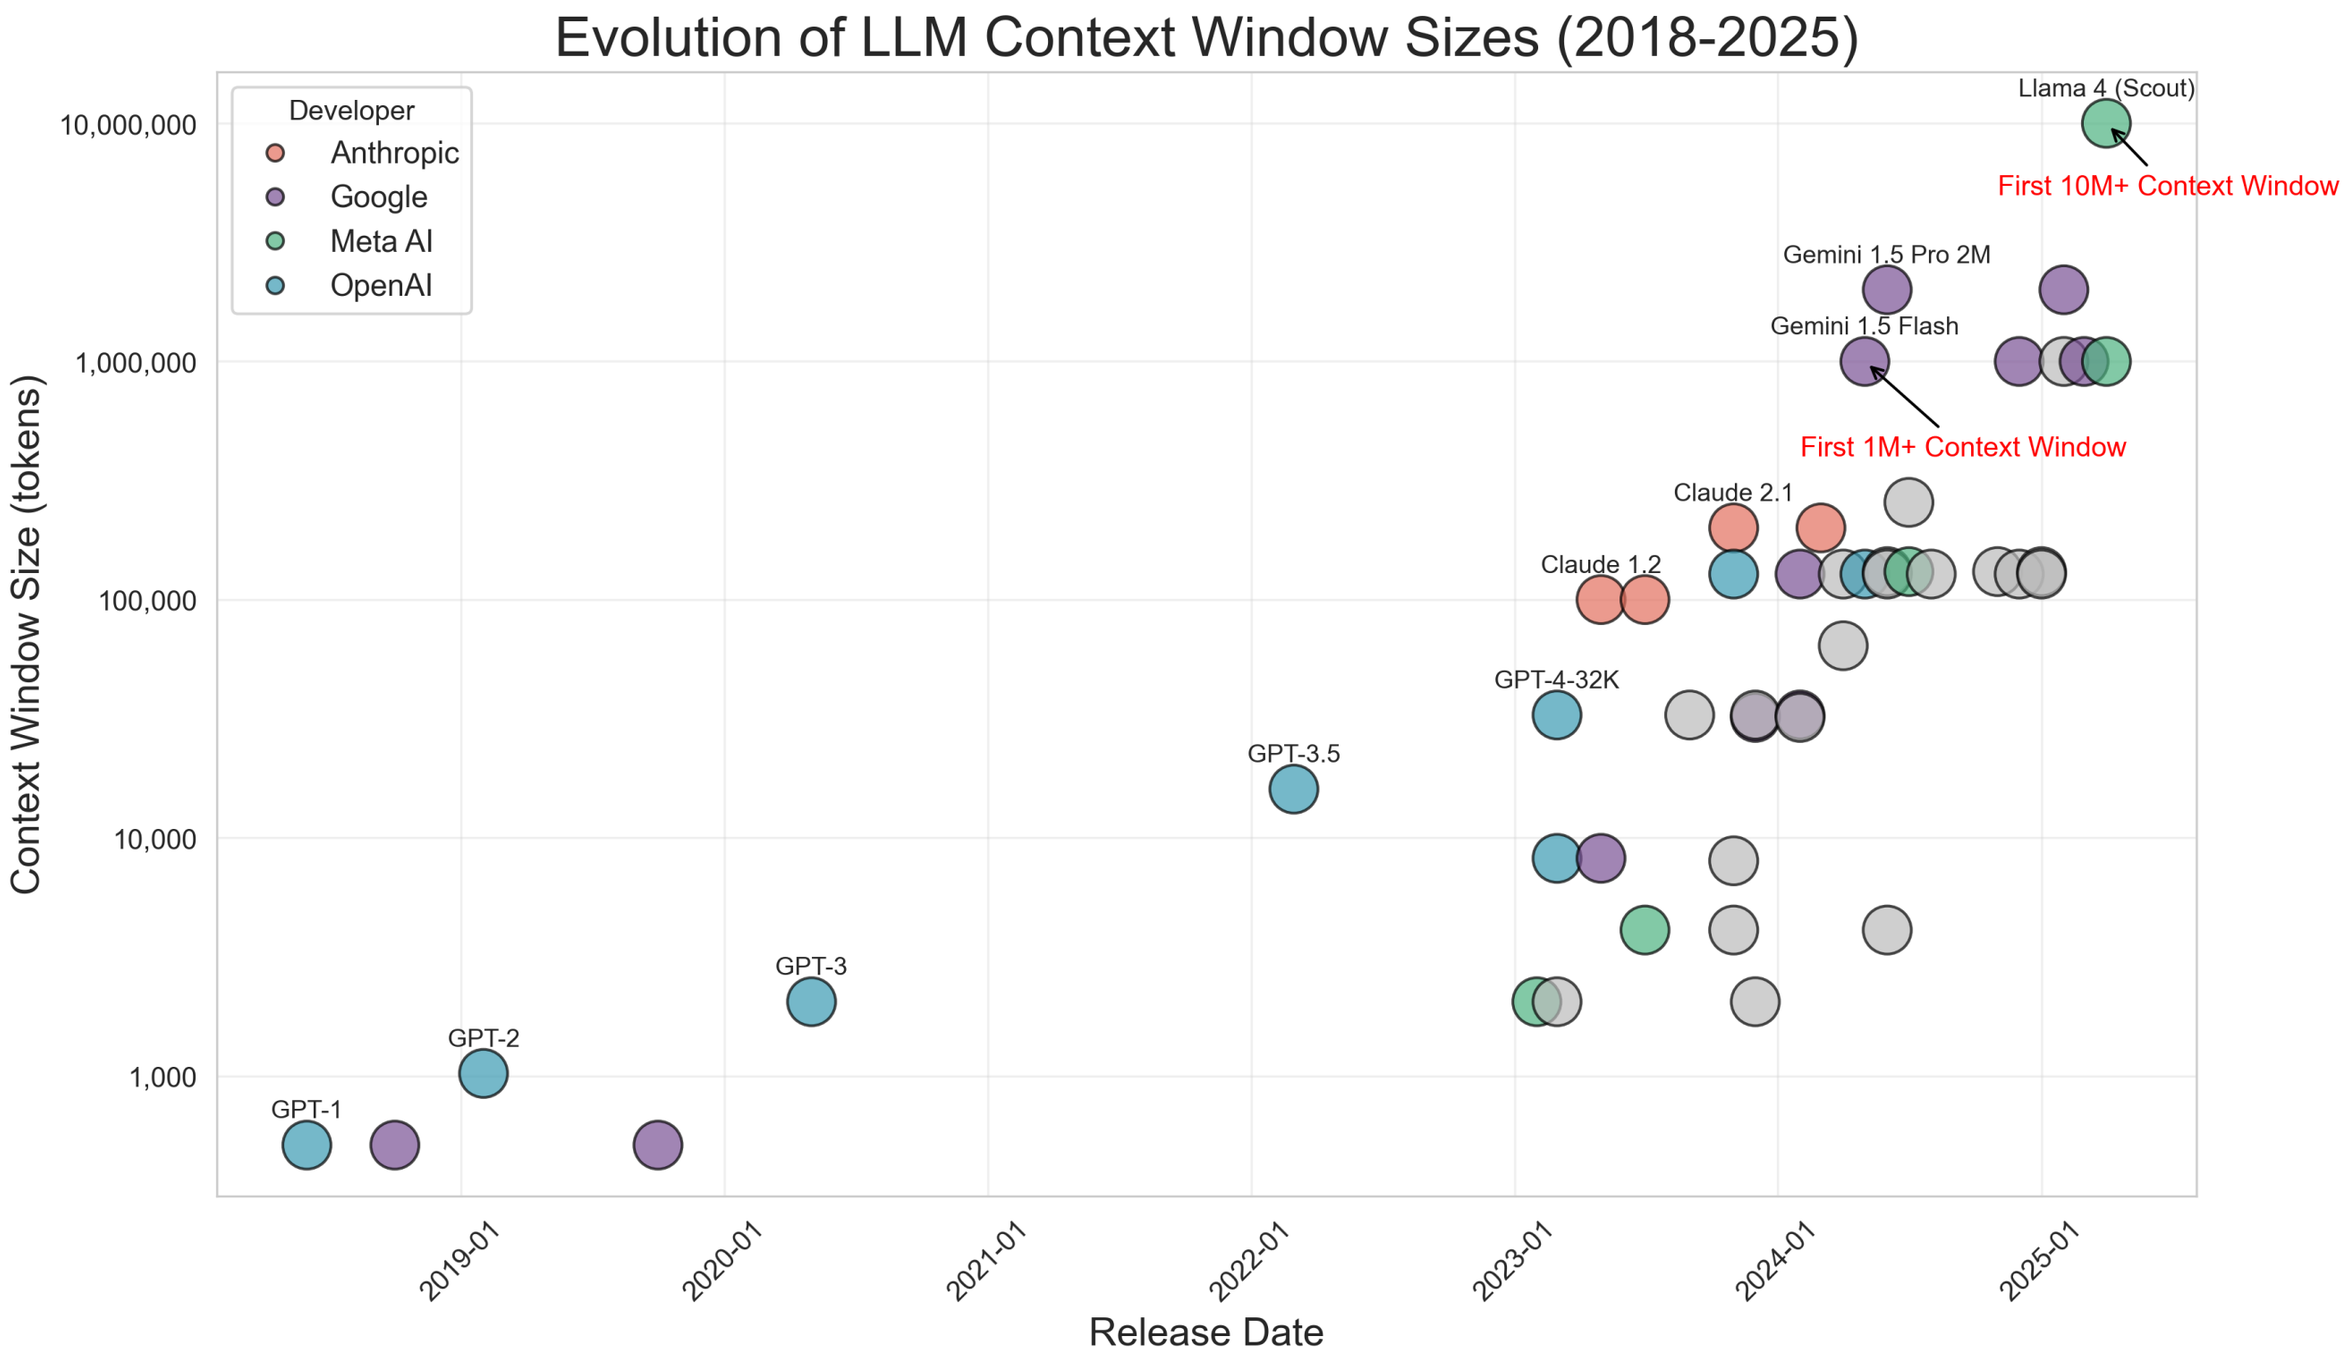


Context Length

- The maximum number of tokens a model can process at once
  - Total input and output tokens



Context Length



Training a Large Language Model

Understanding these principles makes LLMs less of a black box

What resources are required for this?



Training a Large Language Model

Requirements:

- Data
- Computational Power
- Time

Size	Use Case	Dataset Size	Computation	Time
Small	Basic chat, classification	Modest	Single GPU	Minutes-Hours
Medium	Reasoning, summarization	Large	Multiple GPUs	Hours-Days
Large	Translation, coding, agent	Massive	GPU cluster	Days
Frontier	General purpose AI	More massive	Data center	Months

Training a Large Language Model

- More parameters (weights) → Higher memory requirements
- More data → More training steps to learn
- More steps → More compute time and energy

Estimated time and cost to train ChatGPT 3.5

	1 NVIDIA V100 GPU	1024 NVIDIA A100 GPUs
Available Memory	32 GB	81,920 GB
Time	355 years	~34 days
Cost	$4.6-5 million	$2-4 million

Training a Large Language Model

Use and fine-tune existing LLMs rather than train from scratch

Lean on efforts from companies with resources to train these models

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Using a Large Language Model

What are the best ways to prompt an LLM?

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Prompting: Give Clear Context & Role

- Assign a role
  - *“You are a research assistant”*
  - *“You are an agricultural expert”*
- Provide relevant background context
- State your goal explicitly
- Define constraints
  - Audience
  - Format
  - Length

Prompting: Specify Output Structure and Expectations

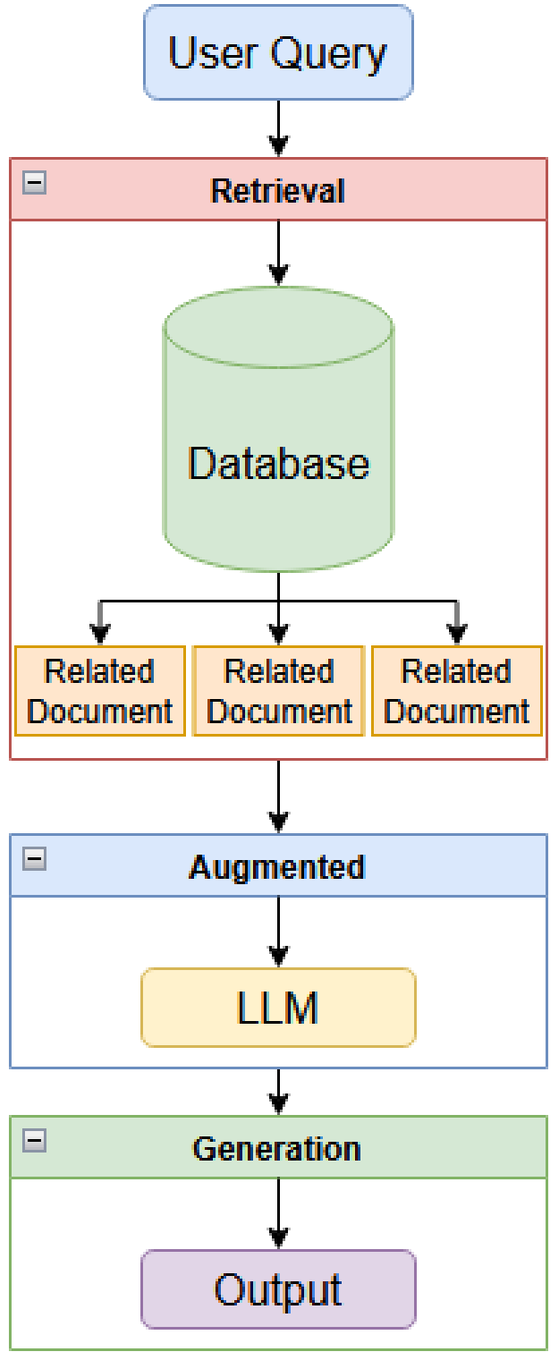
- Request a specific format
  - Bulleted list, table, JSON, markdown, LaTeX
- Provide examples when precision matters
  - Few-Shot prompting
  - Examples of what the output should be helps model be more accurate
- Tell the model “Think step-by-step” for solving problems
  - “Reasoning” models
- Define limitations and success criteria

Using a Large Language Model

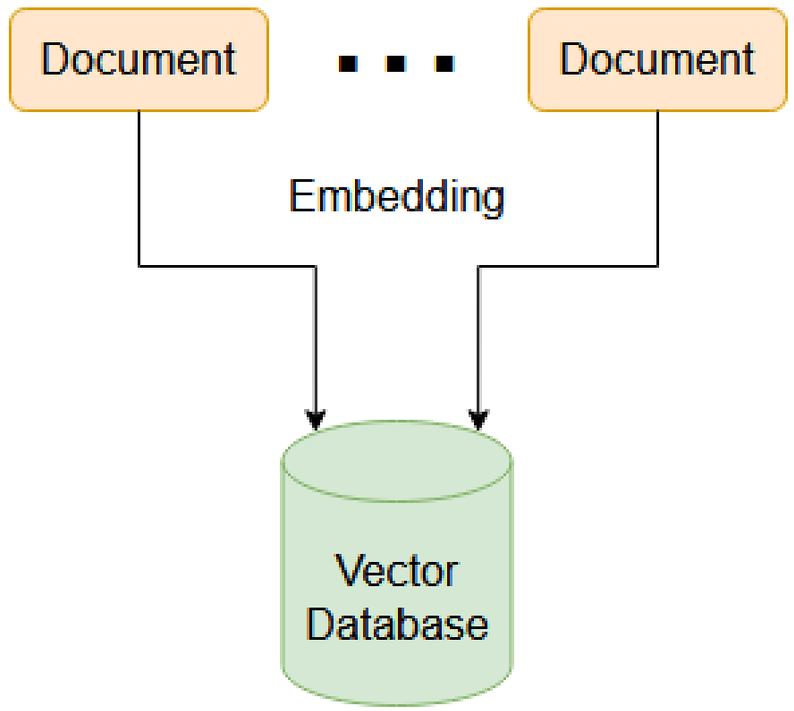
How do you customize models for your specific needs?

Retrieval-Augmented Generation (RAG)

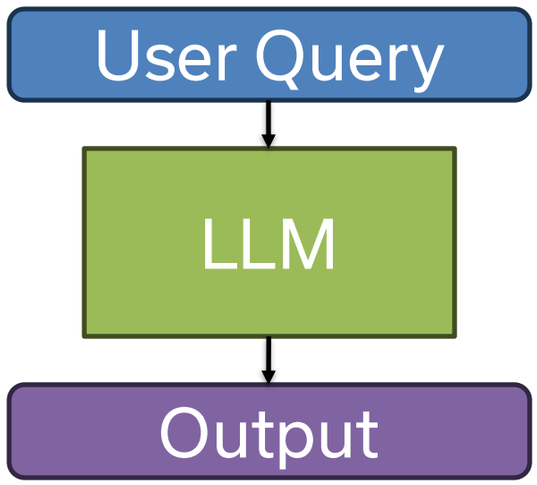
- LLMs often hallucinate
  - Provide inaccurate responses with high confidence
- **RAG**: Supplying specific information to LLM context
- Great for large document/data stores
  - All data would not fit within context window
  - Intelligently retrieve relevant data

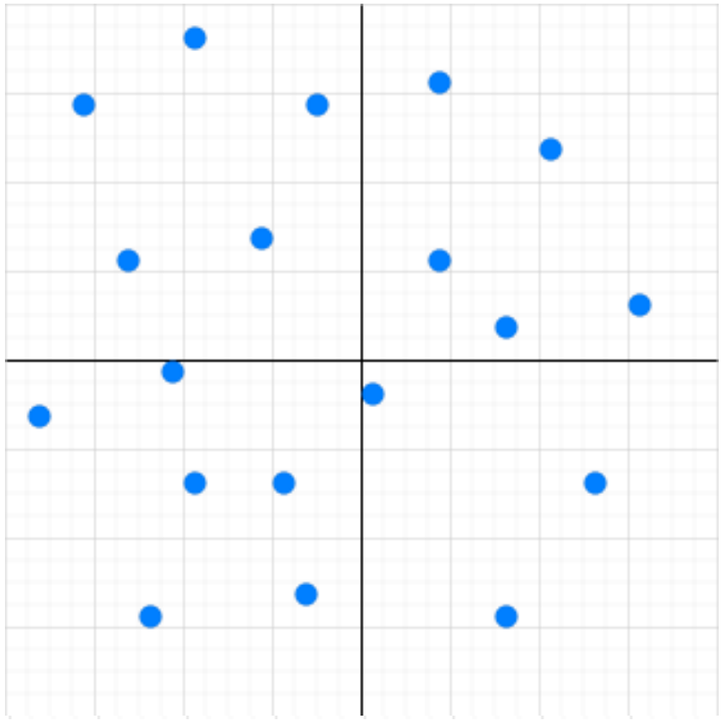


Vector Databases and Embedding Search

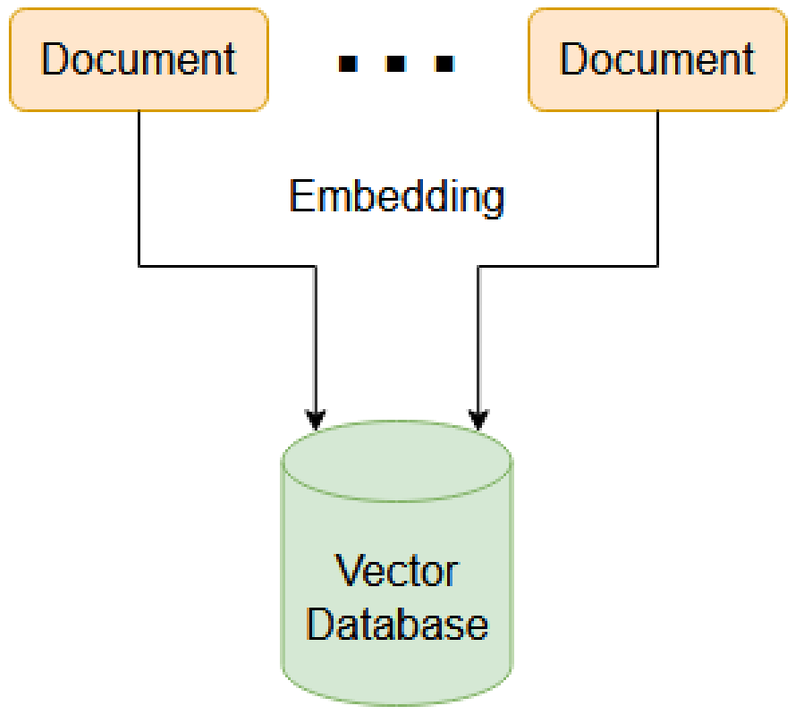


- Embed data with transformer model
- Store in embedding-data pair in **vector database**
- **Similarity Searching**
  - More like documents embed closer together

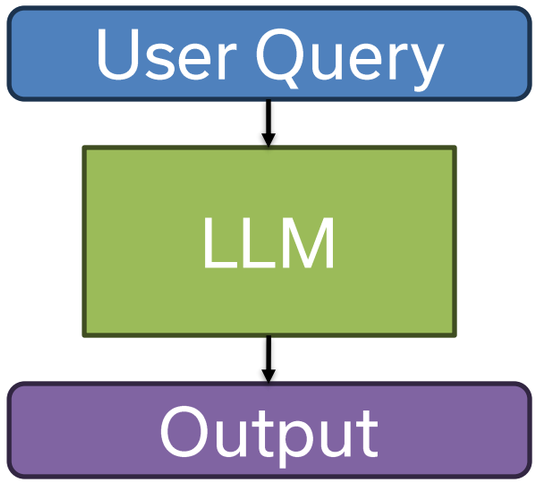


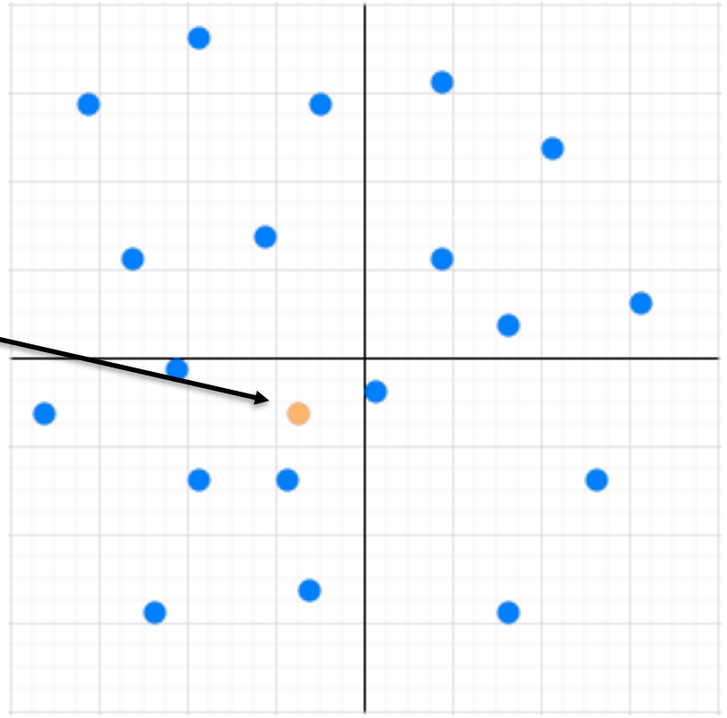


Vector Databases and Embedding Search

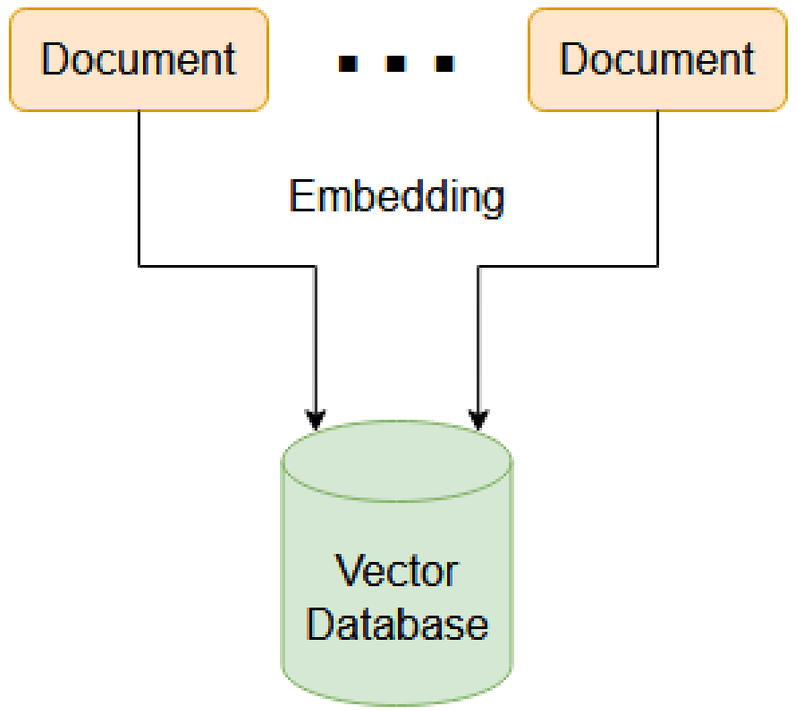


- Embed data with transformer model
- Store in embedding-data pair in **vector database**
- **Similarity Searching**
  - More like documents embed closer together
  - Convert query to strong search term
  - Embed search term in vector space

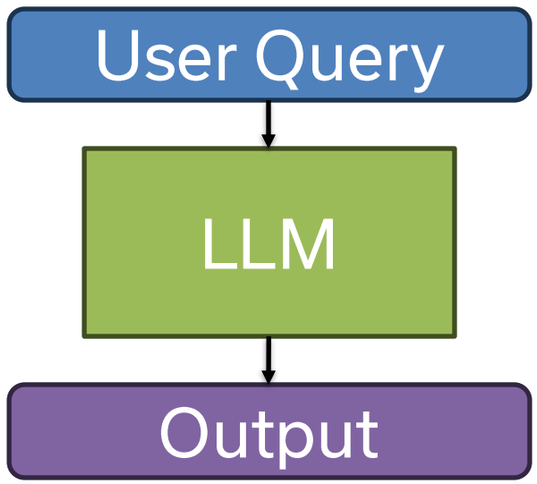


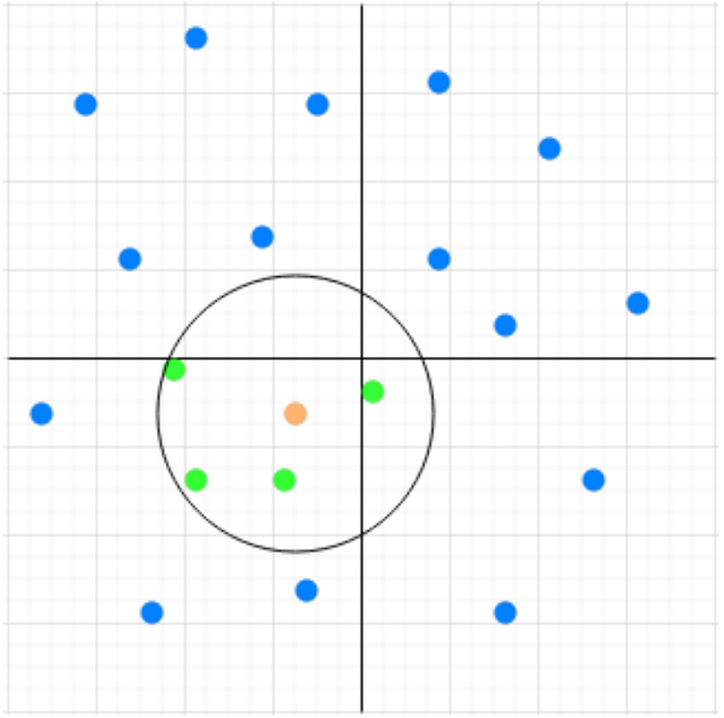


Vector Databases and Embedding Search



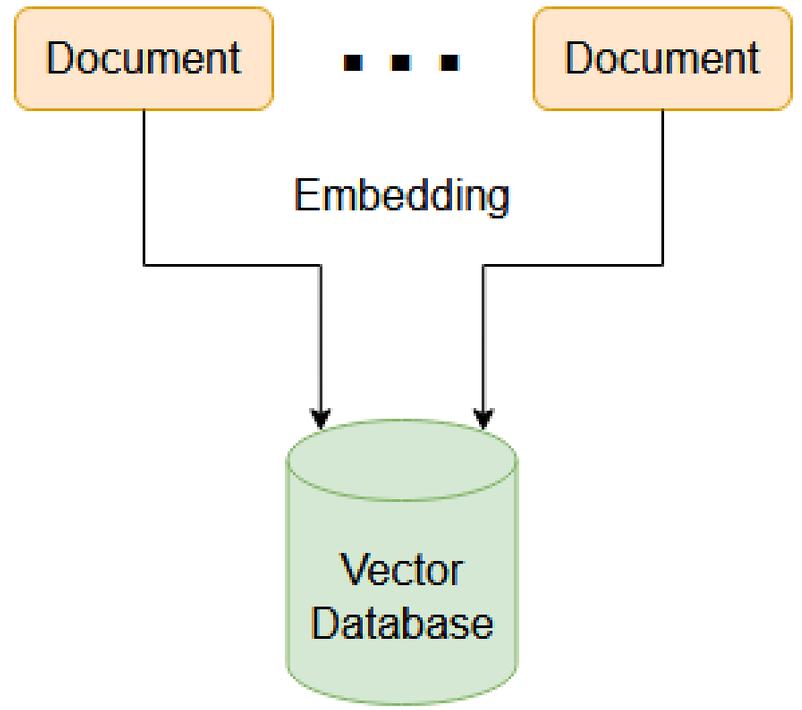
- Embed data with transformer model
- Store in embedding-data pair in **vector database**
- **Similarity Searching**
  - More like documents embed closer together
  - Convert query to strong search term
  - Embed search term in vector space
  - Find most similar docs

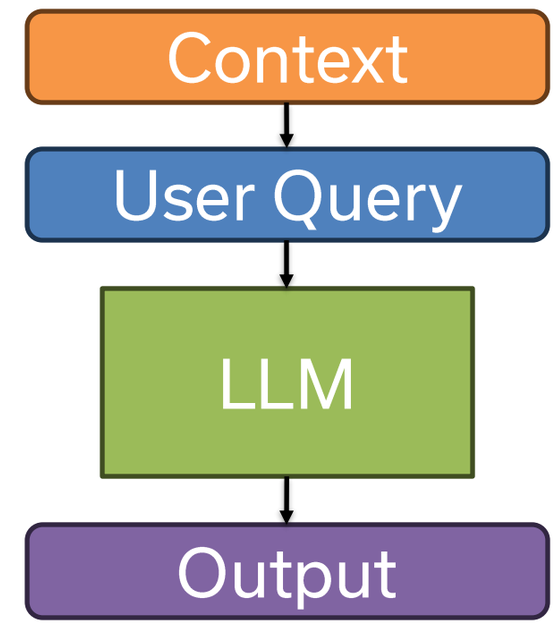


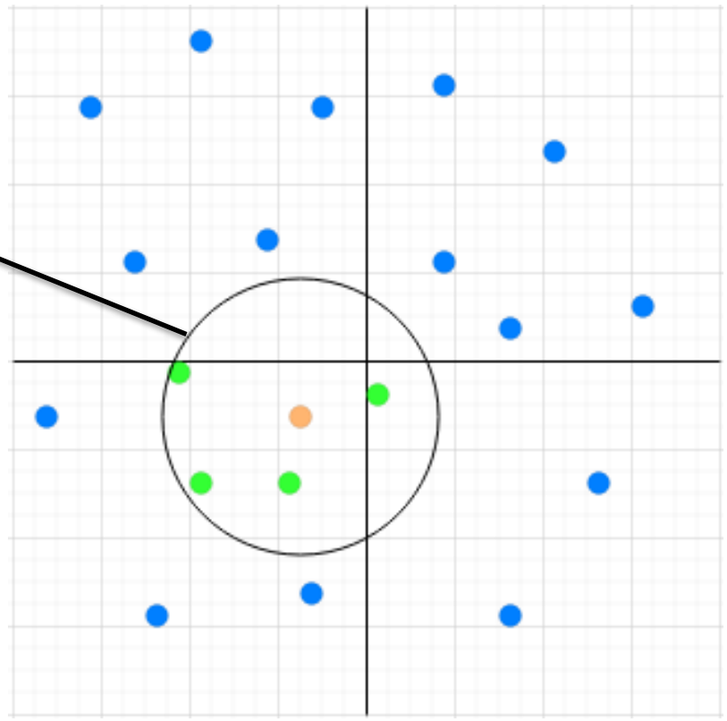


Vector Databases and Embedding Search

- Embed data with transformer model
- Store in embedding-data pair in **vector database**
- **Similarity Searching**
  - More like documents embed closer together
  - Convert query to strong search term
  - Embed search term in vector space
  - Find most similar docs
  - Add to context







Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities



Tool Calling and Agents

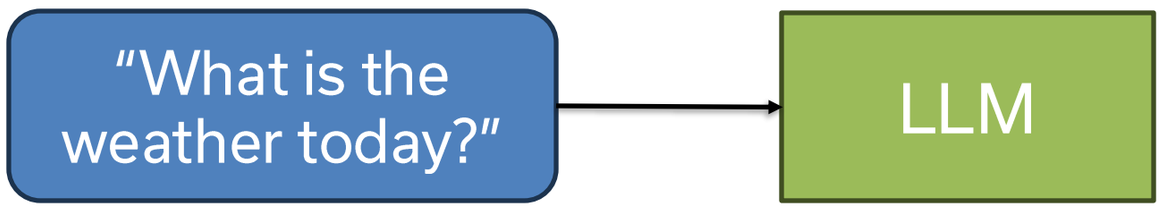
- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities



This is a hallucination

Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities

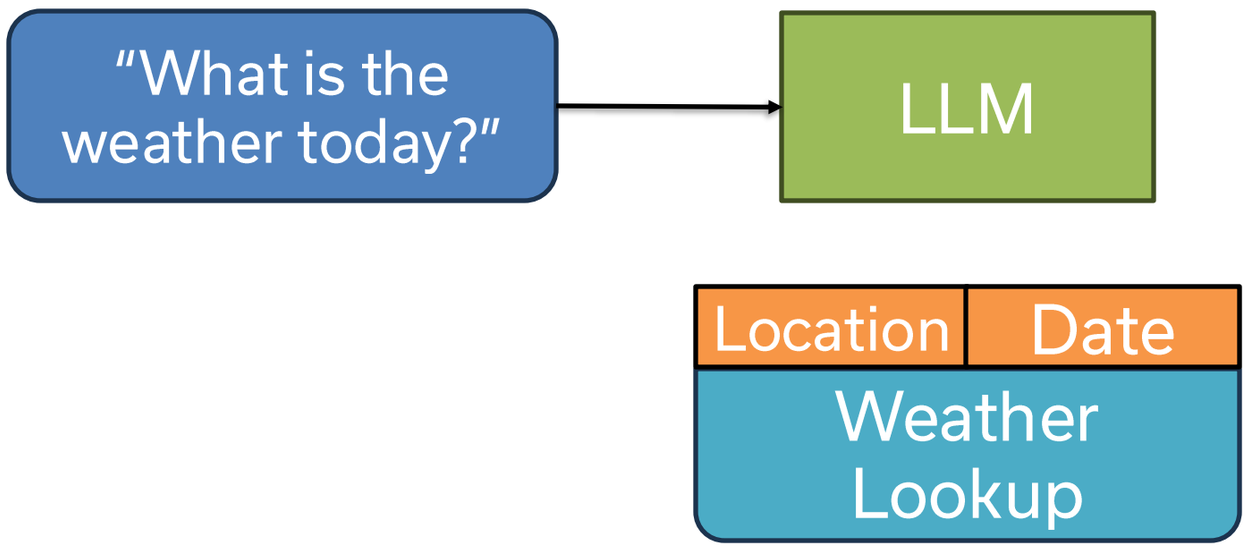


Need:

- Location
- Date
- Current weather for that date and location

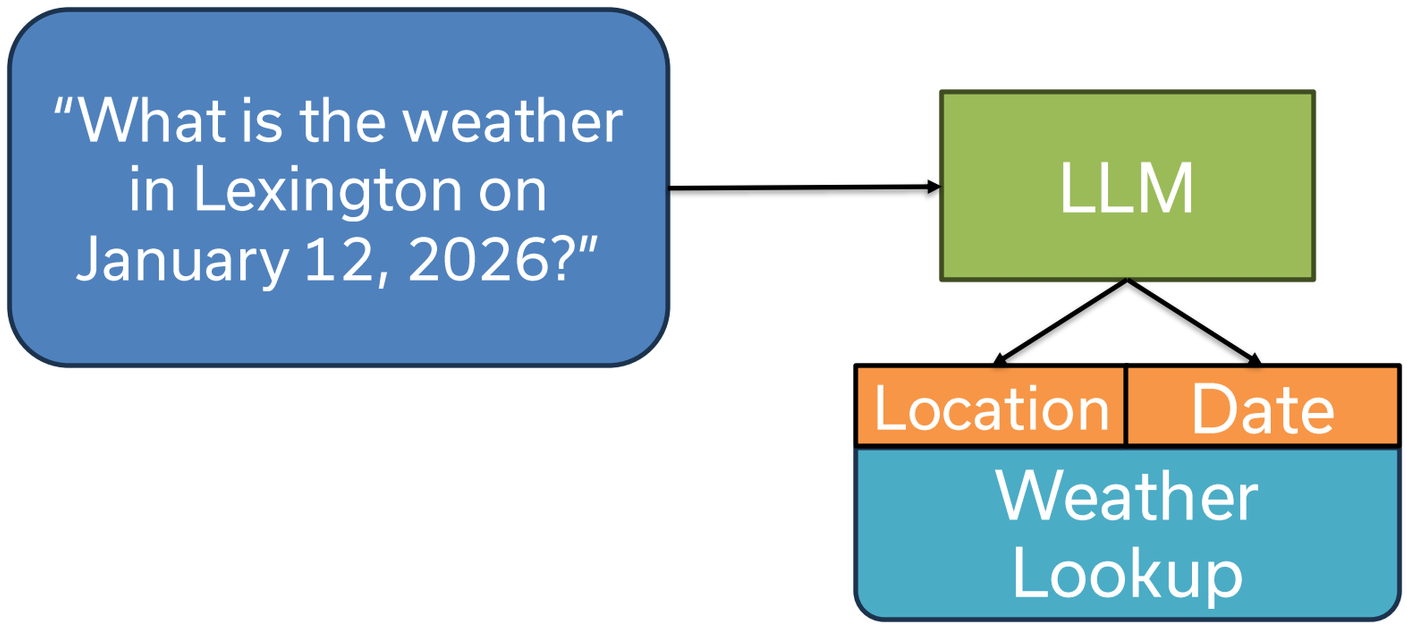
Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities



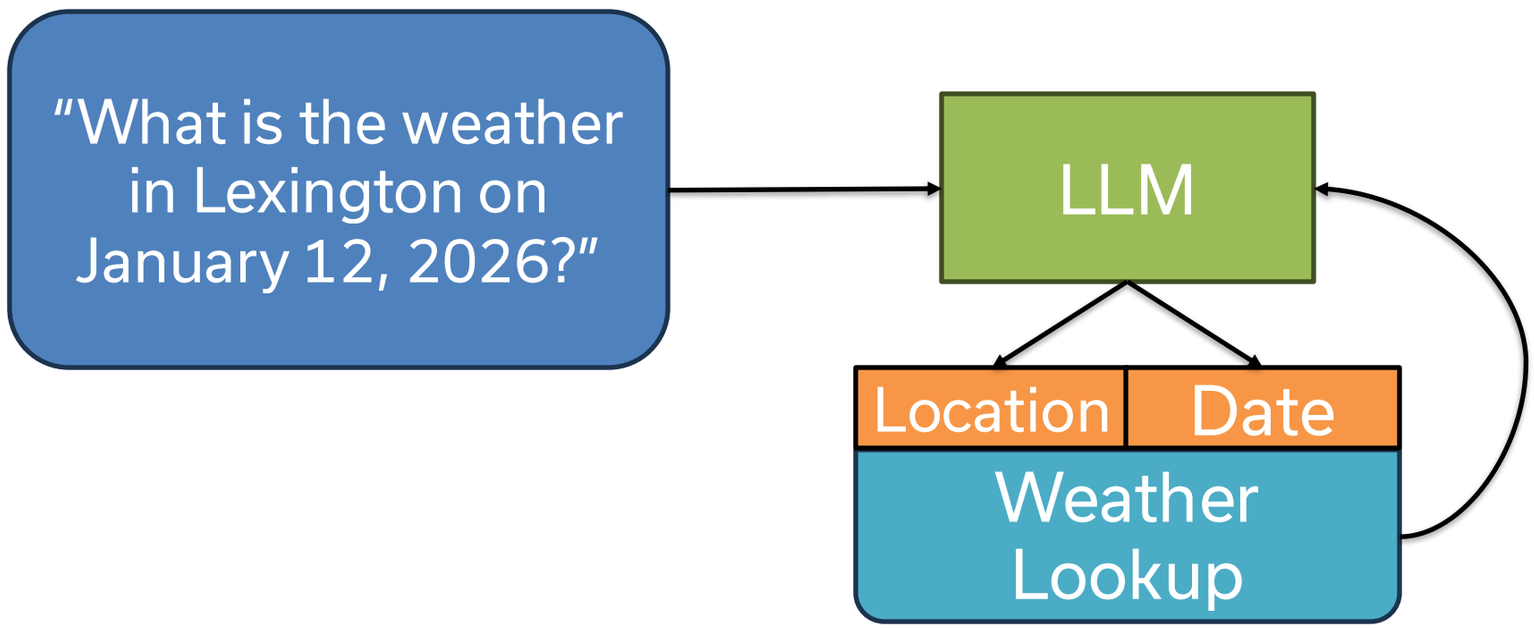
Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities



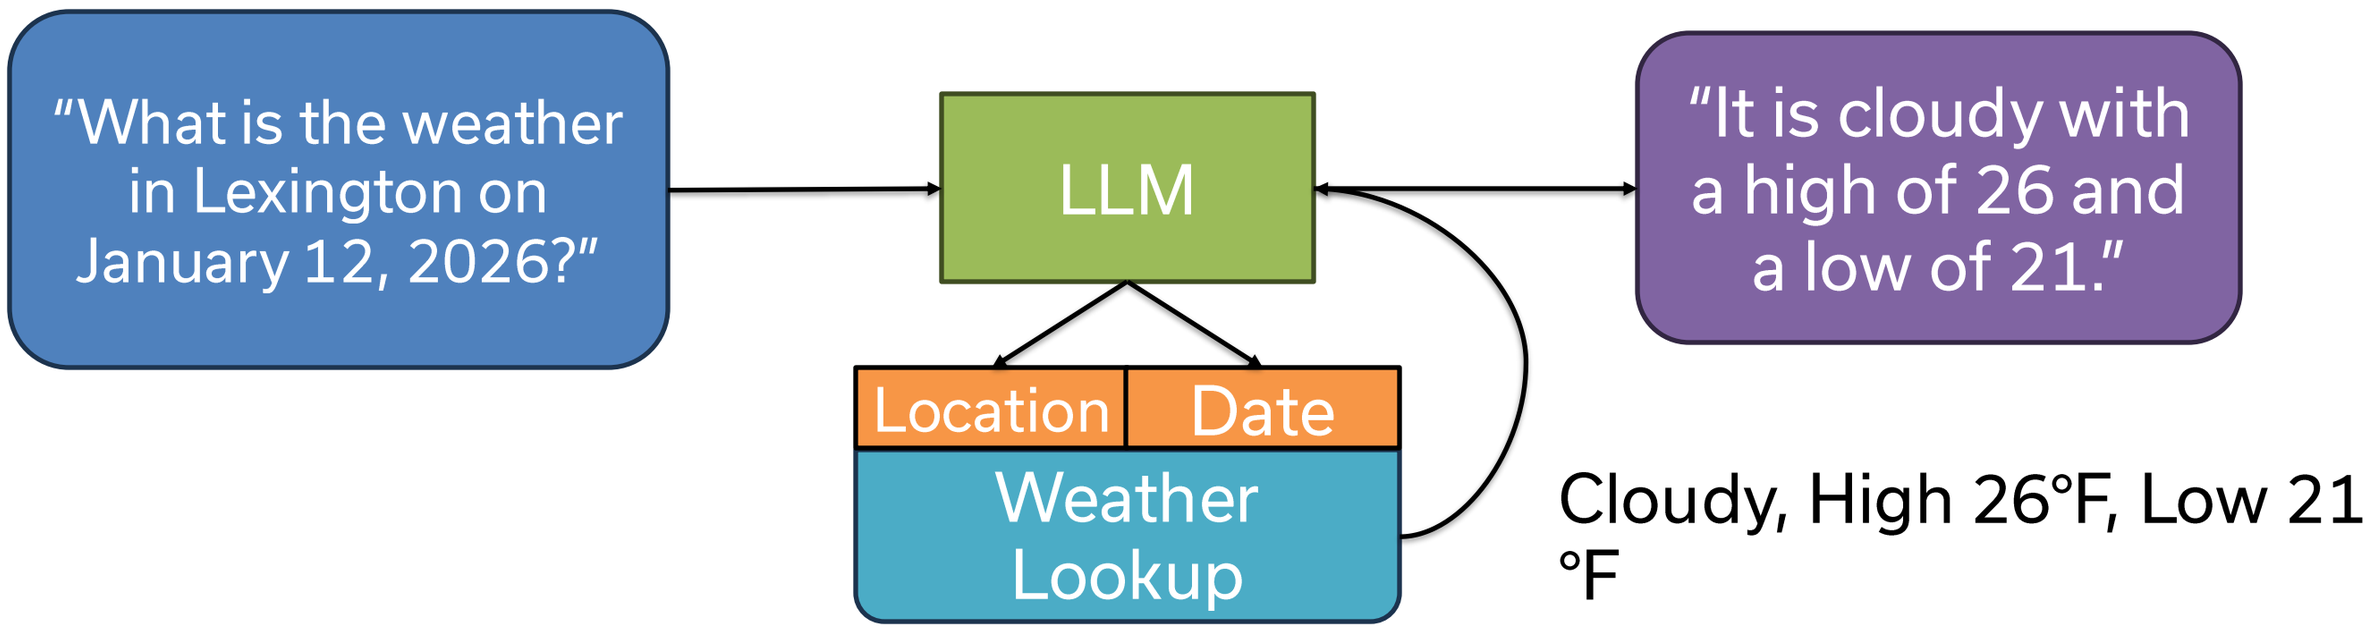
Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities



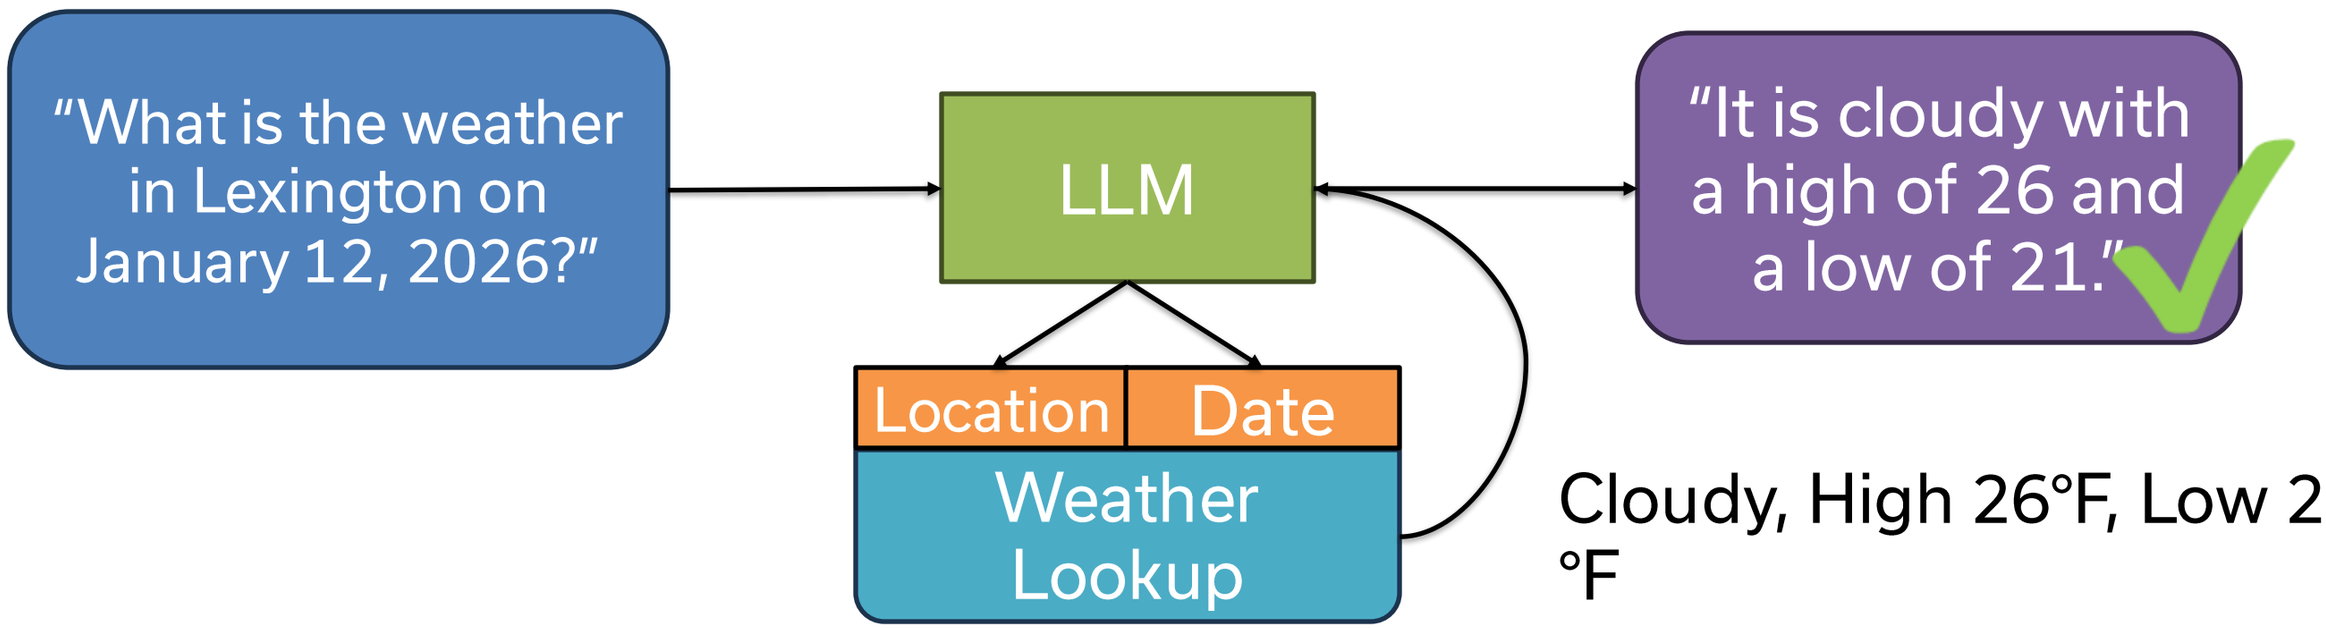
Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities



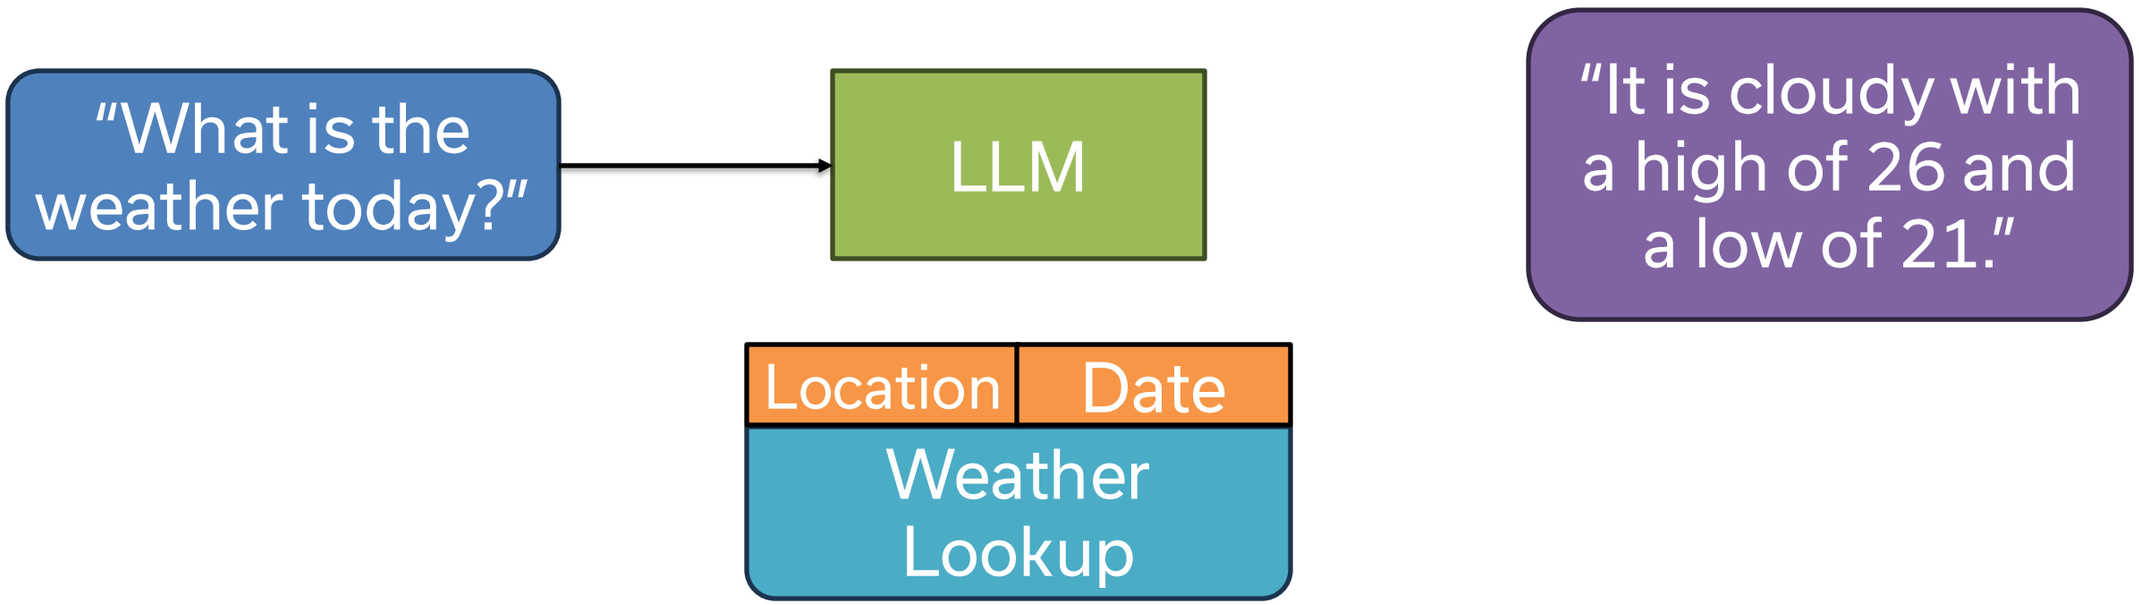
Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities



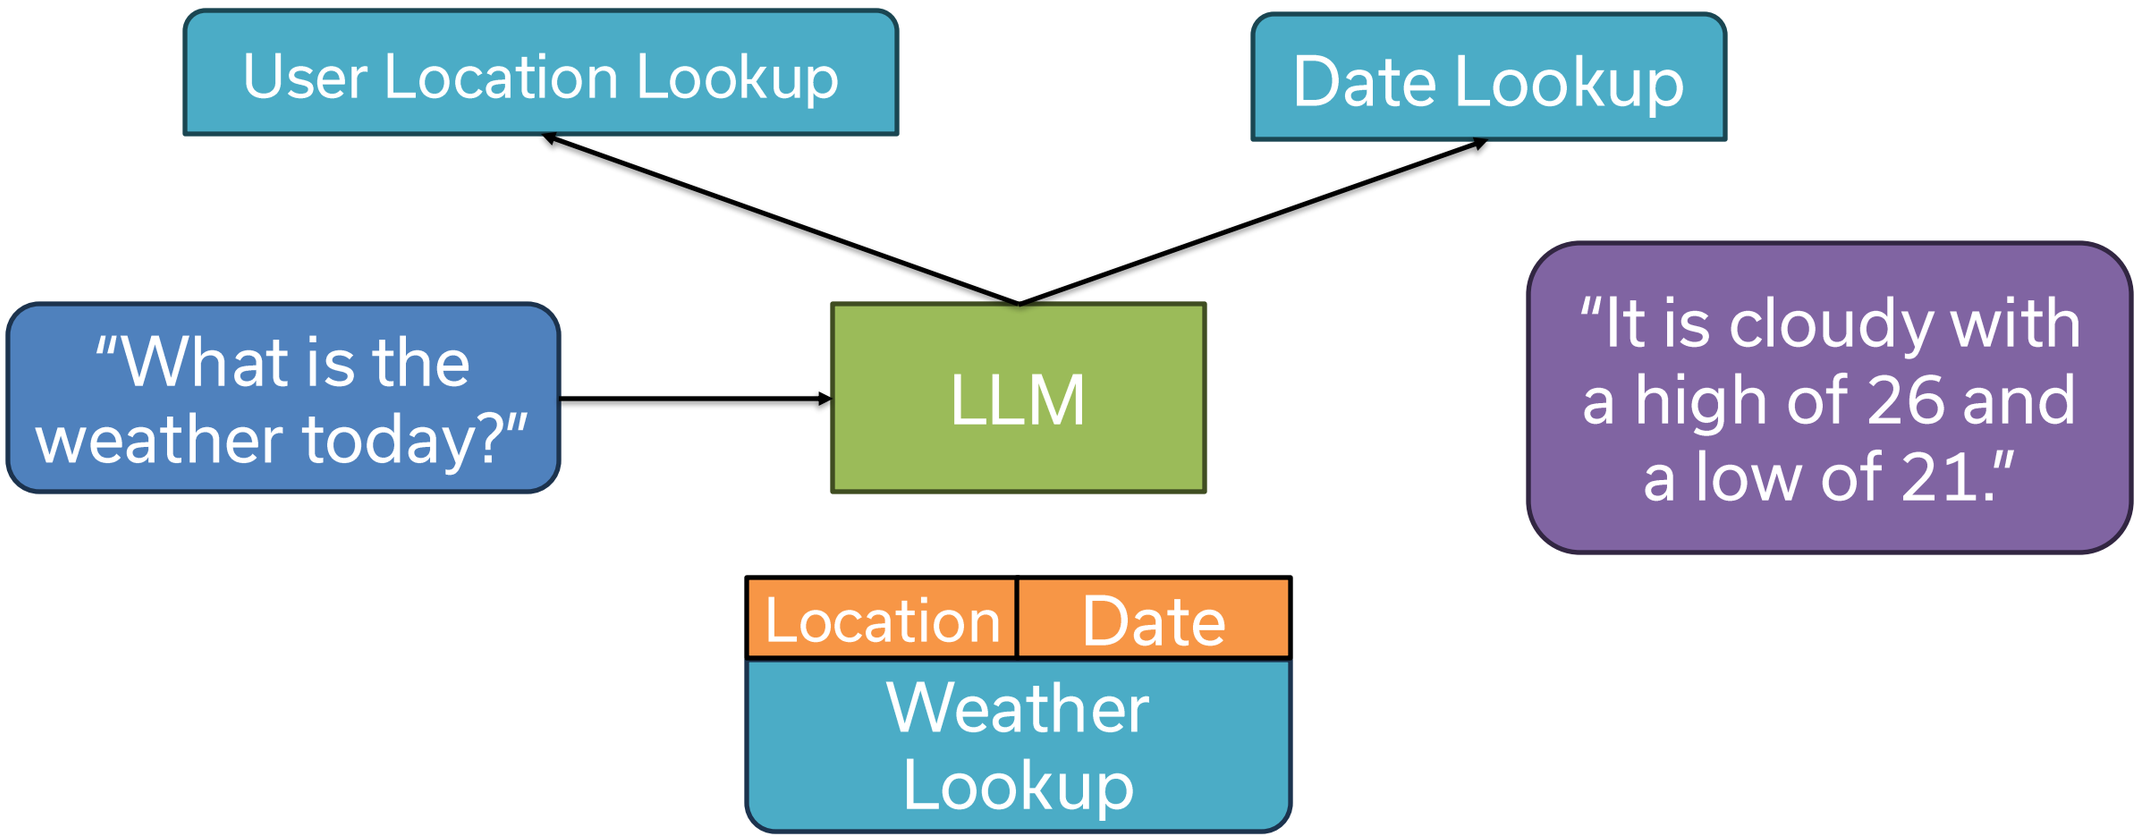
Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents**: An LLM assistant provided with specific tool capabilities



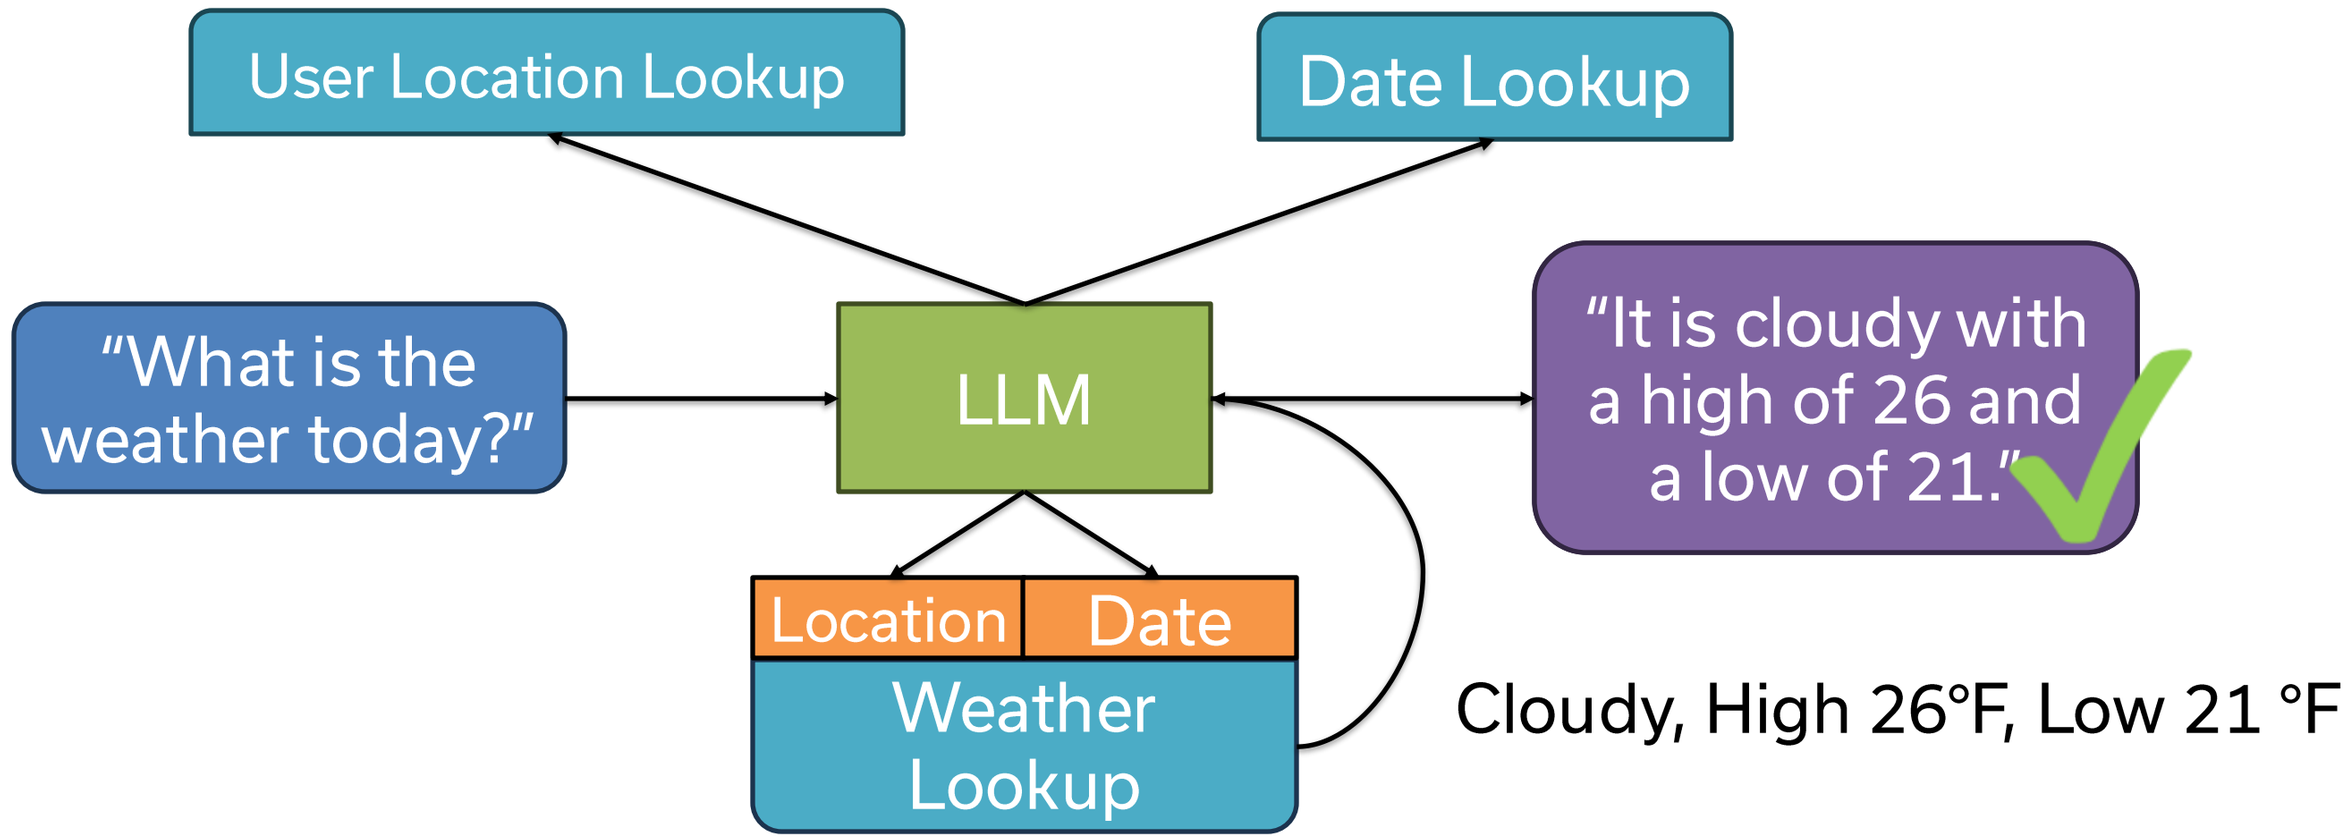
Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents:** An LLM assistant provided with specific tool capabilities



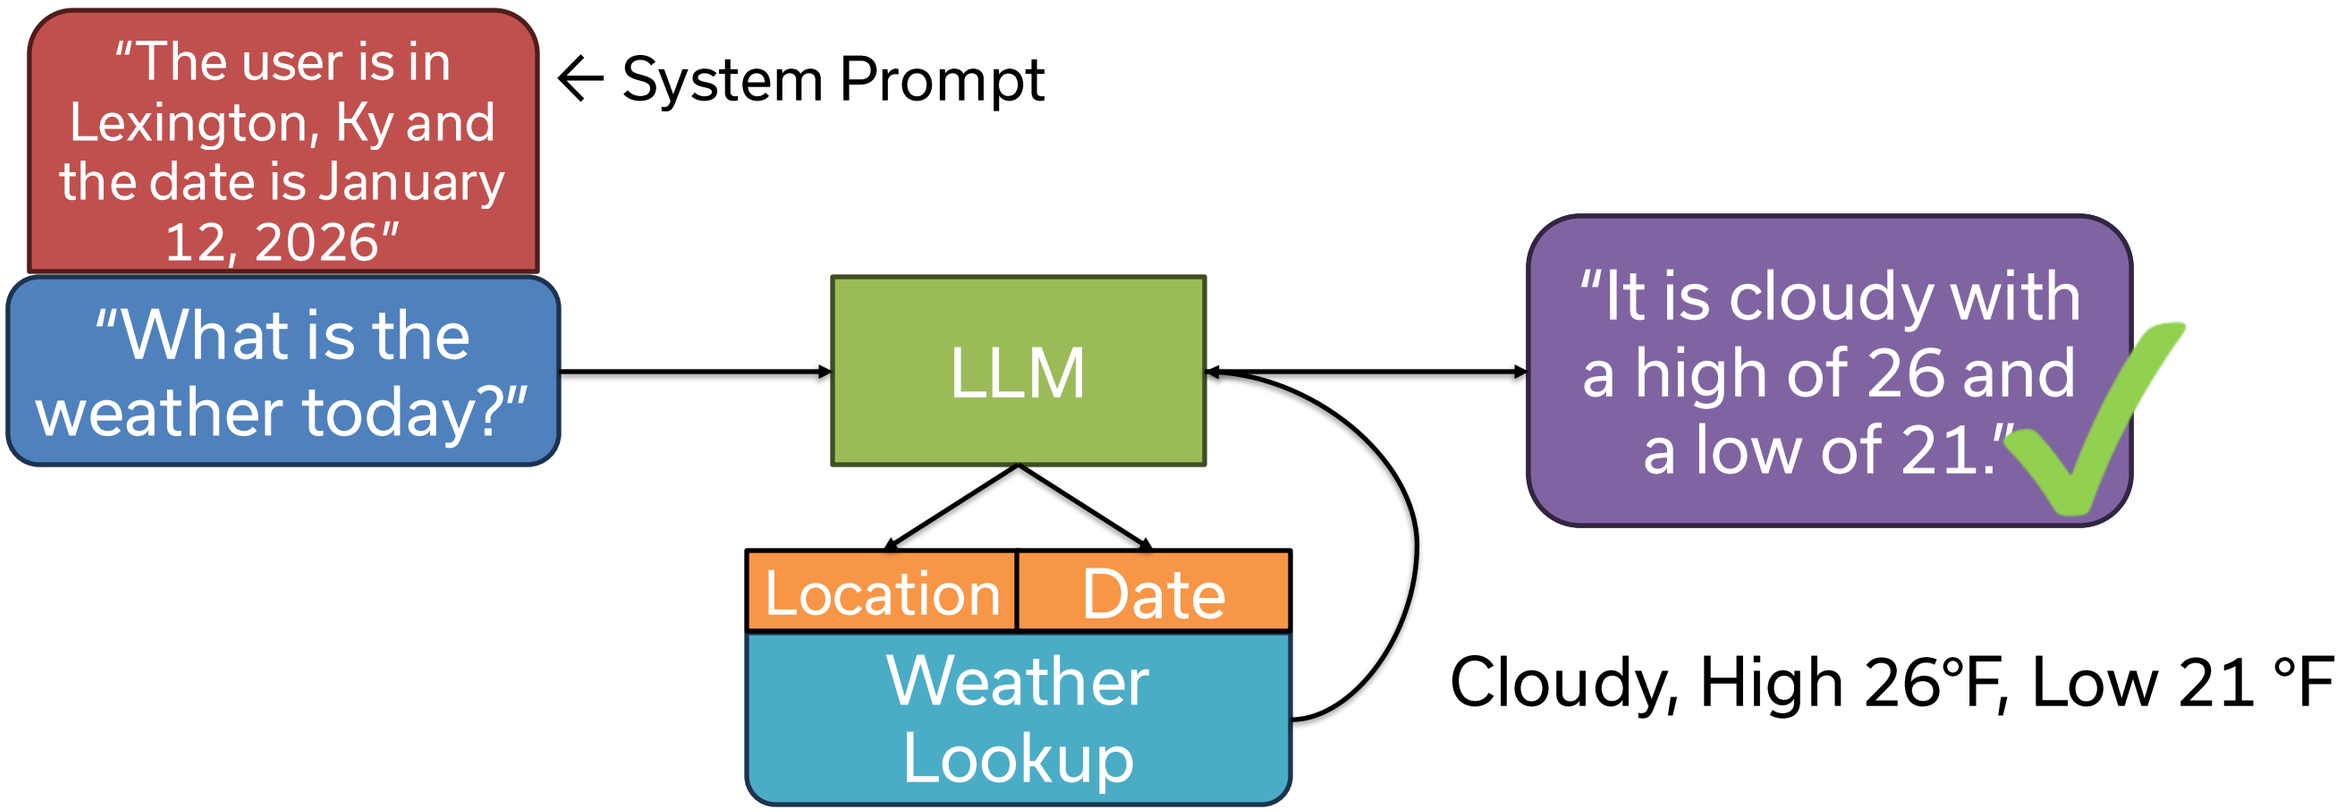
Tool Calling and Agents

- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents:** An LLM assistant provided with specific tool capabilities



Tool Calling and Agents

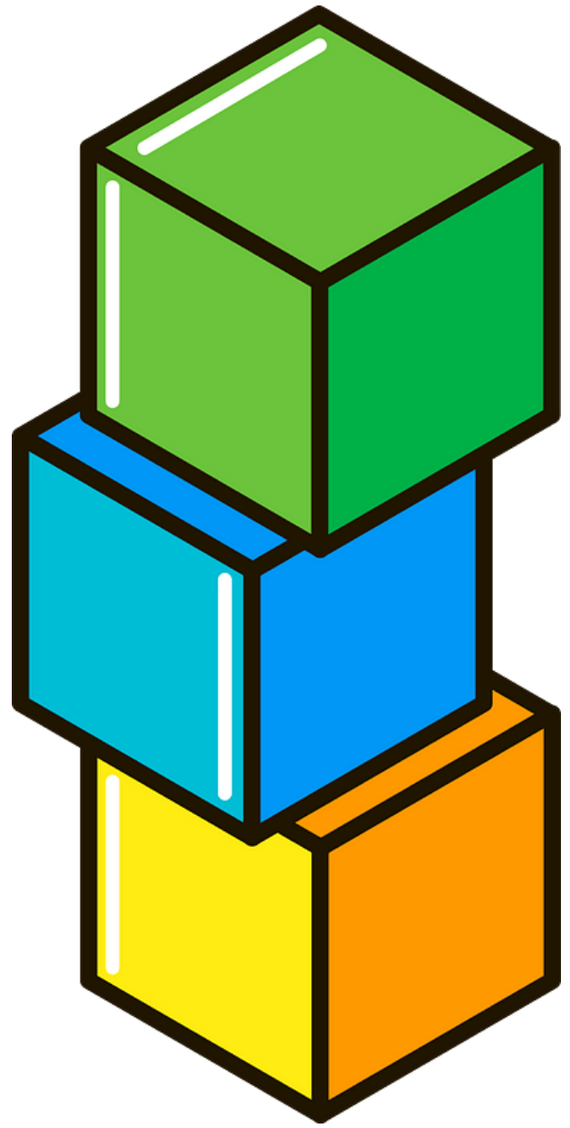
- Register external capabilities (**tools**) for LLM to call on
- LLMs are entirely text based
- From a prompt, can identify a tool that would aid in response and generate the parameters that tool needs
- **Agents:** An LLM assistant provided with specific tool capabilities



Training a Large Language Model

- Building block of language AI
- Further advancements expand on these fundamentals

- *How do you integrate **tool calling** and **RAG** to create your unique agent?*



Computer Vision

Vaiden Logan, BS

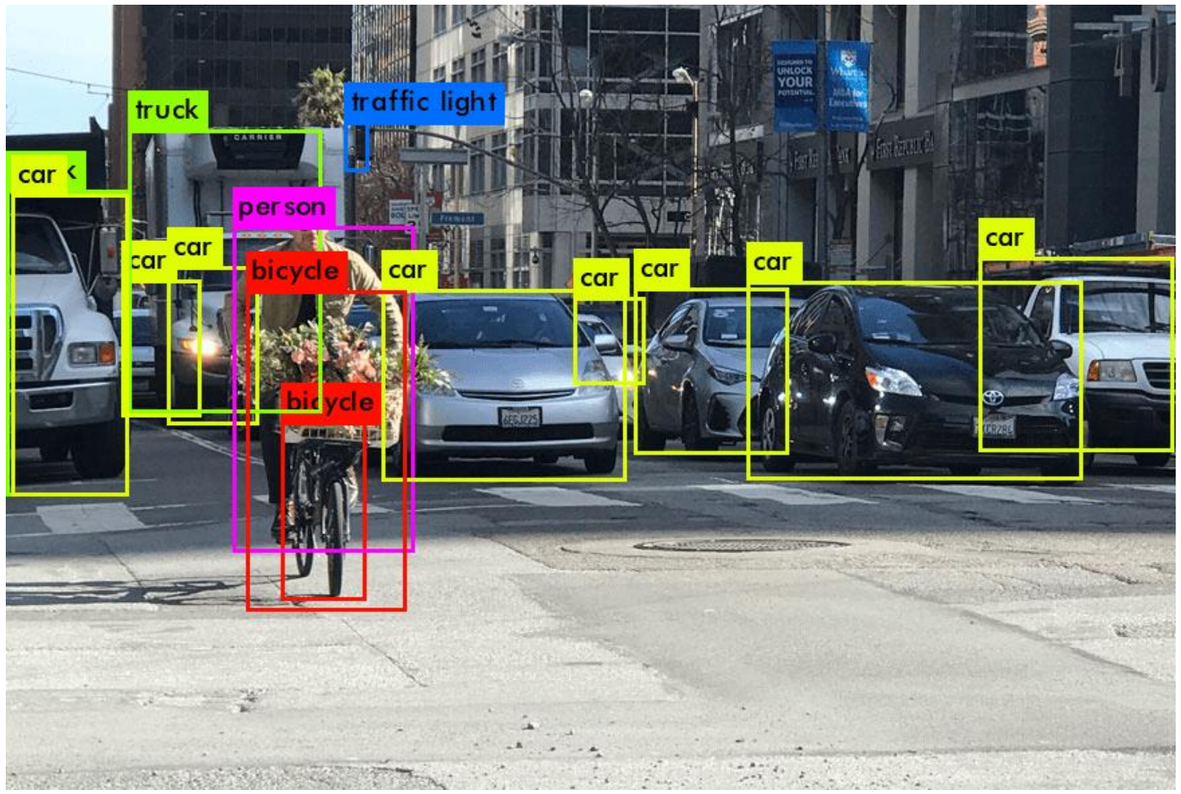
Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

Computer Vision Models

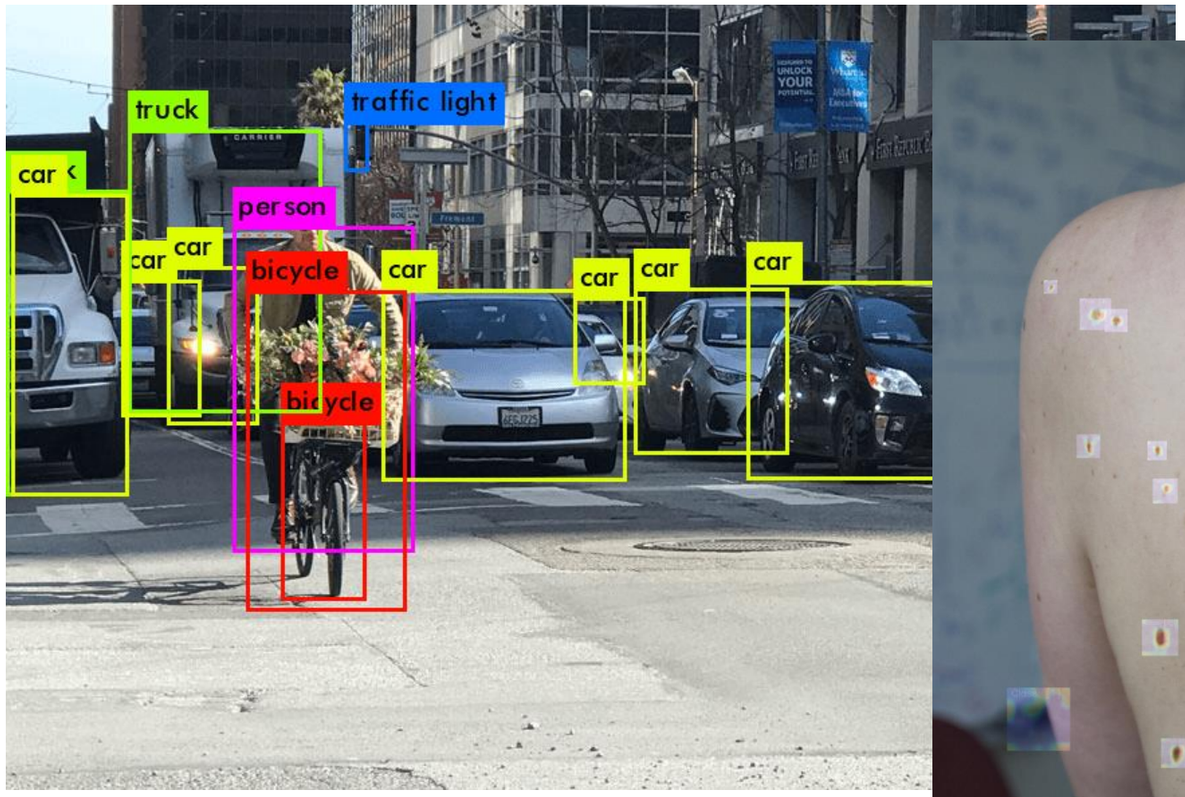
Where do we see vision models used?

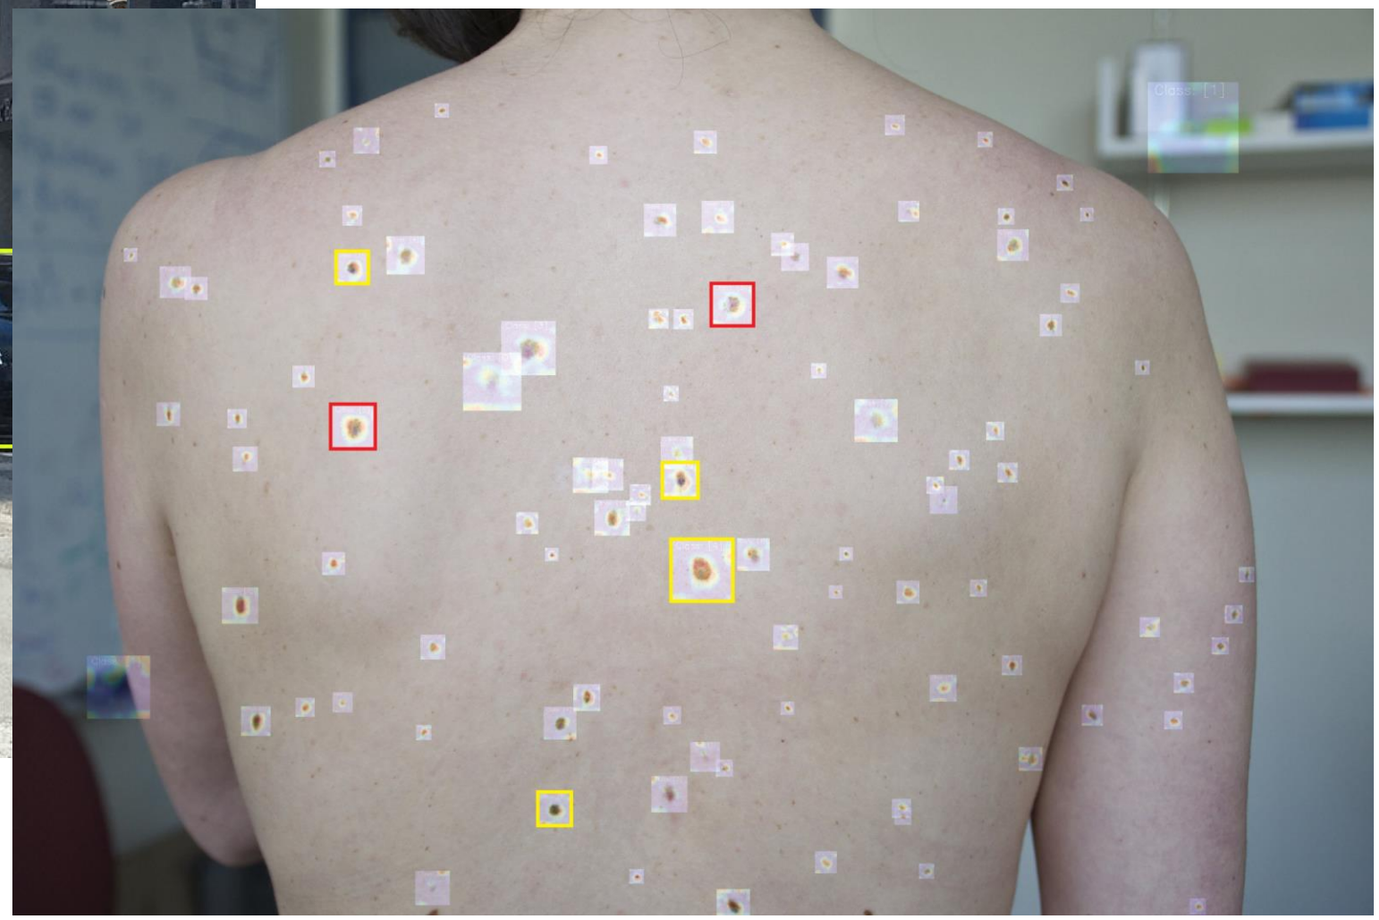
Computer Vision Models



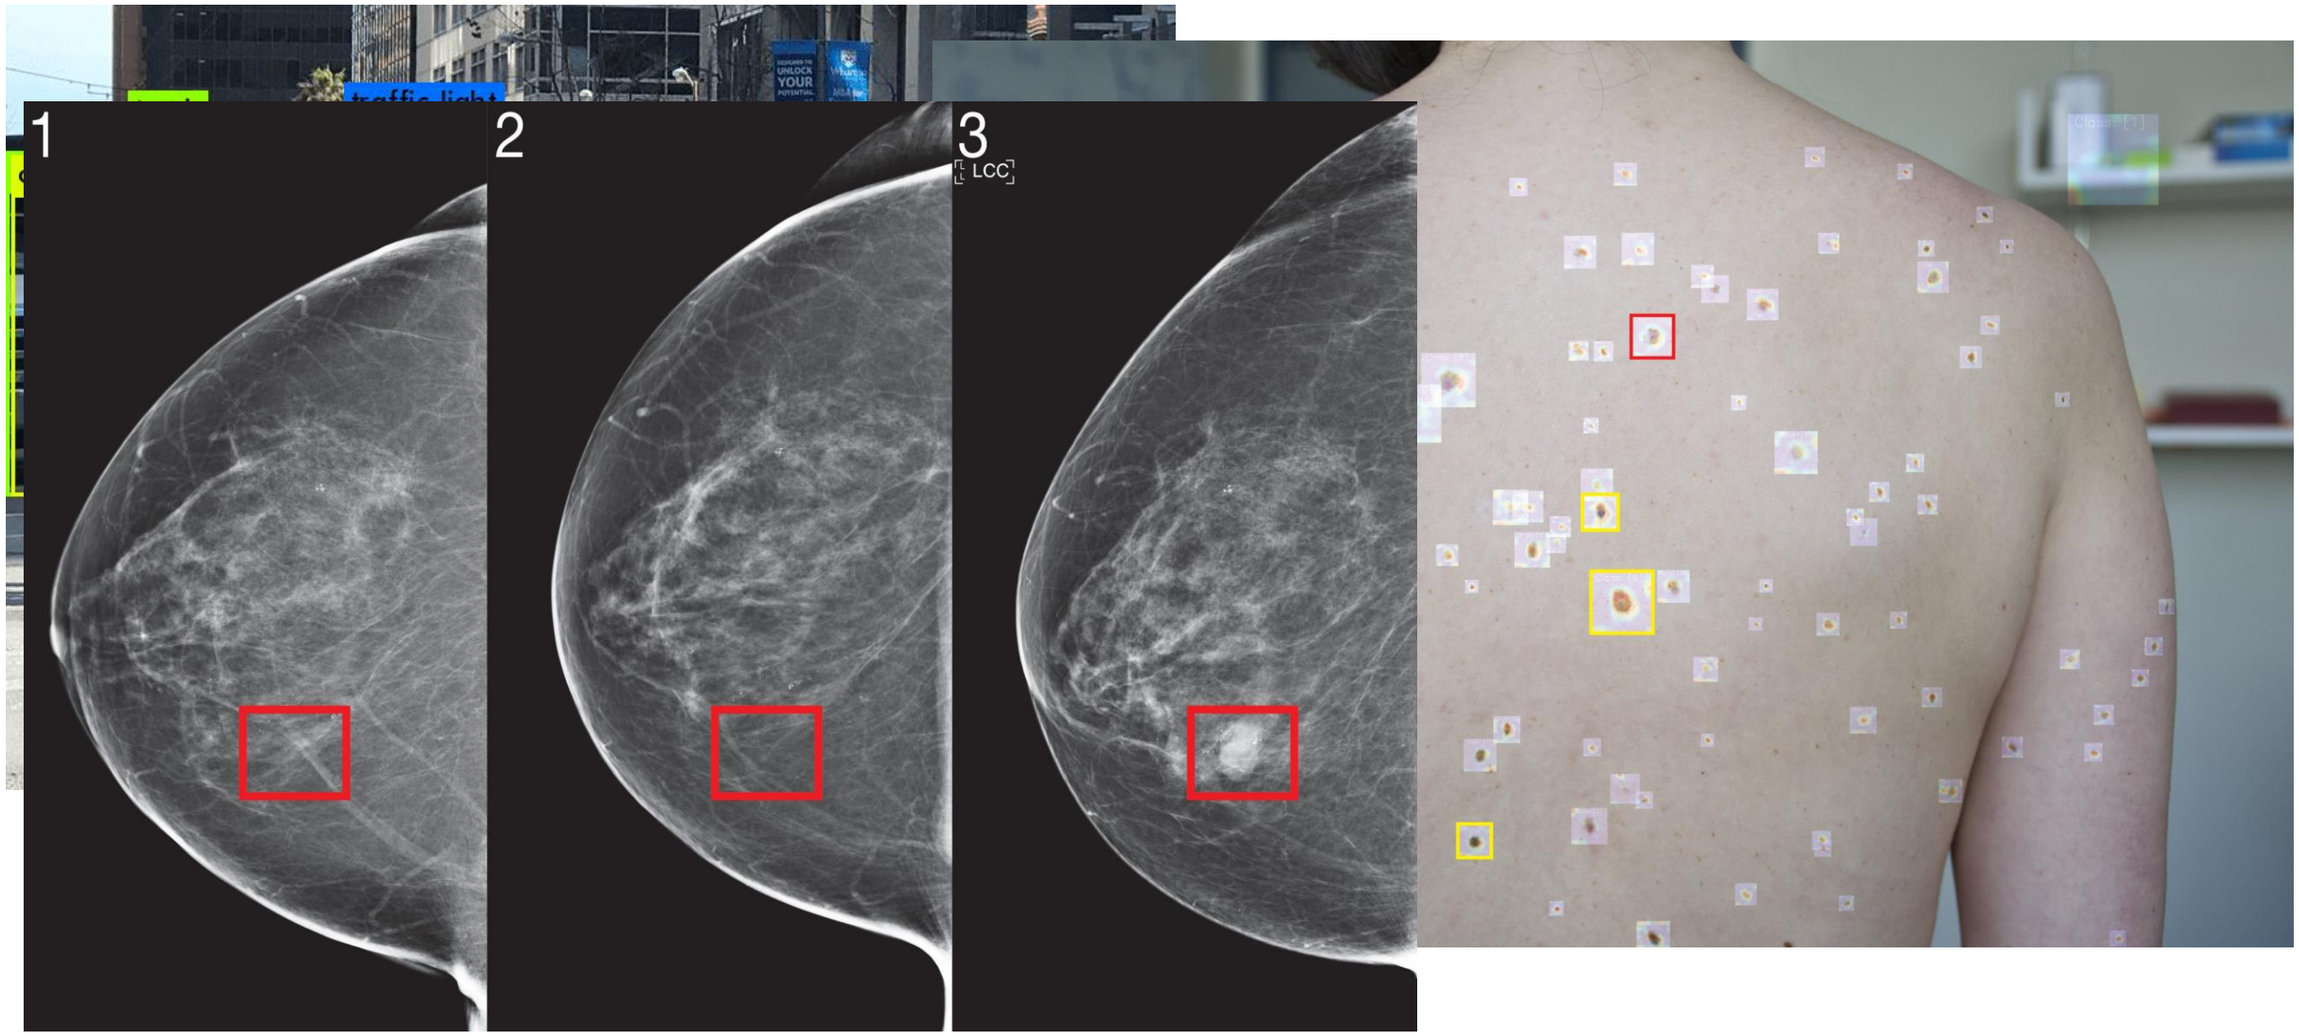
NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
CPE Higher Education Matters
University of Kentucky

Computer Vision Models

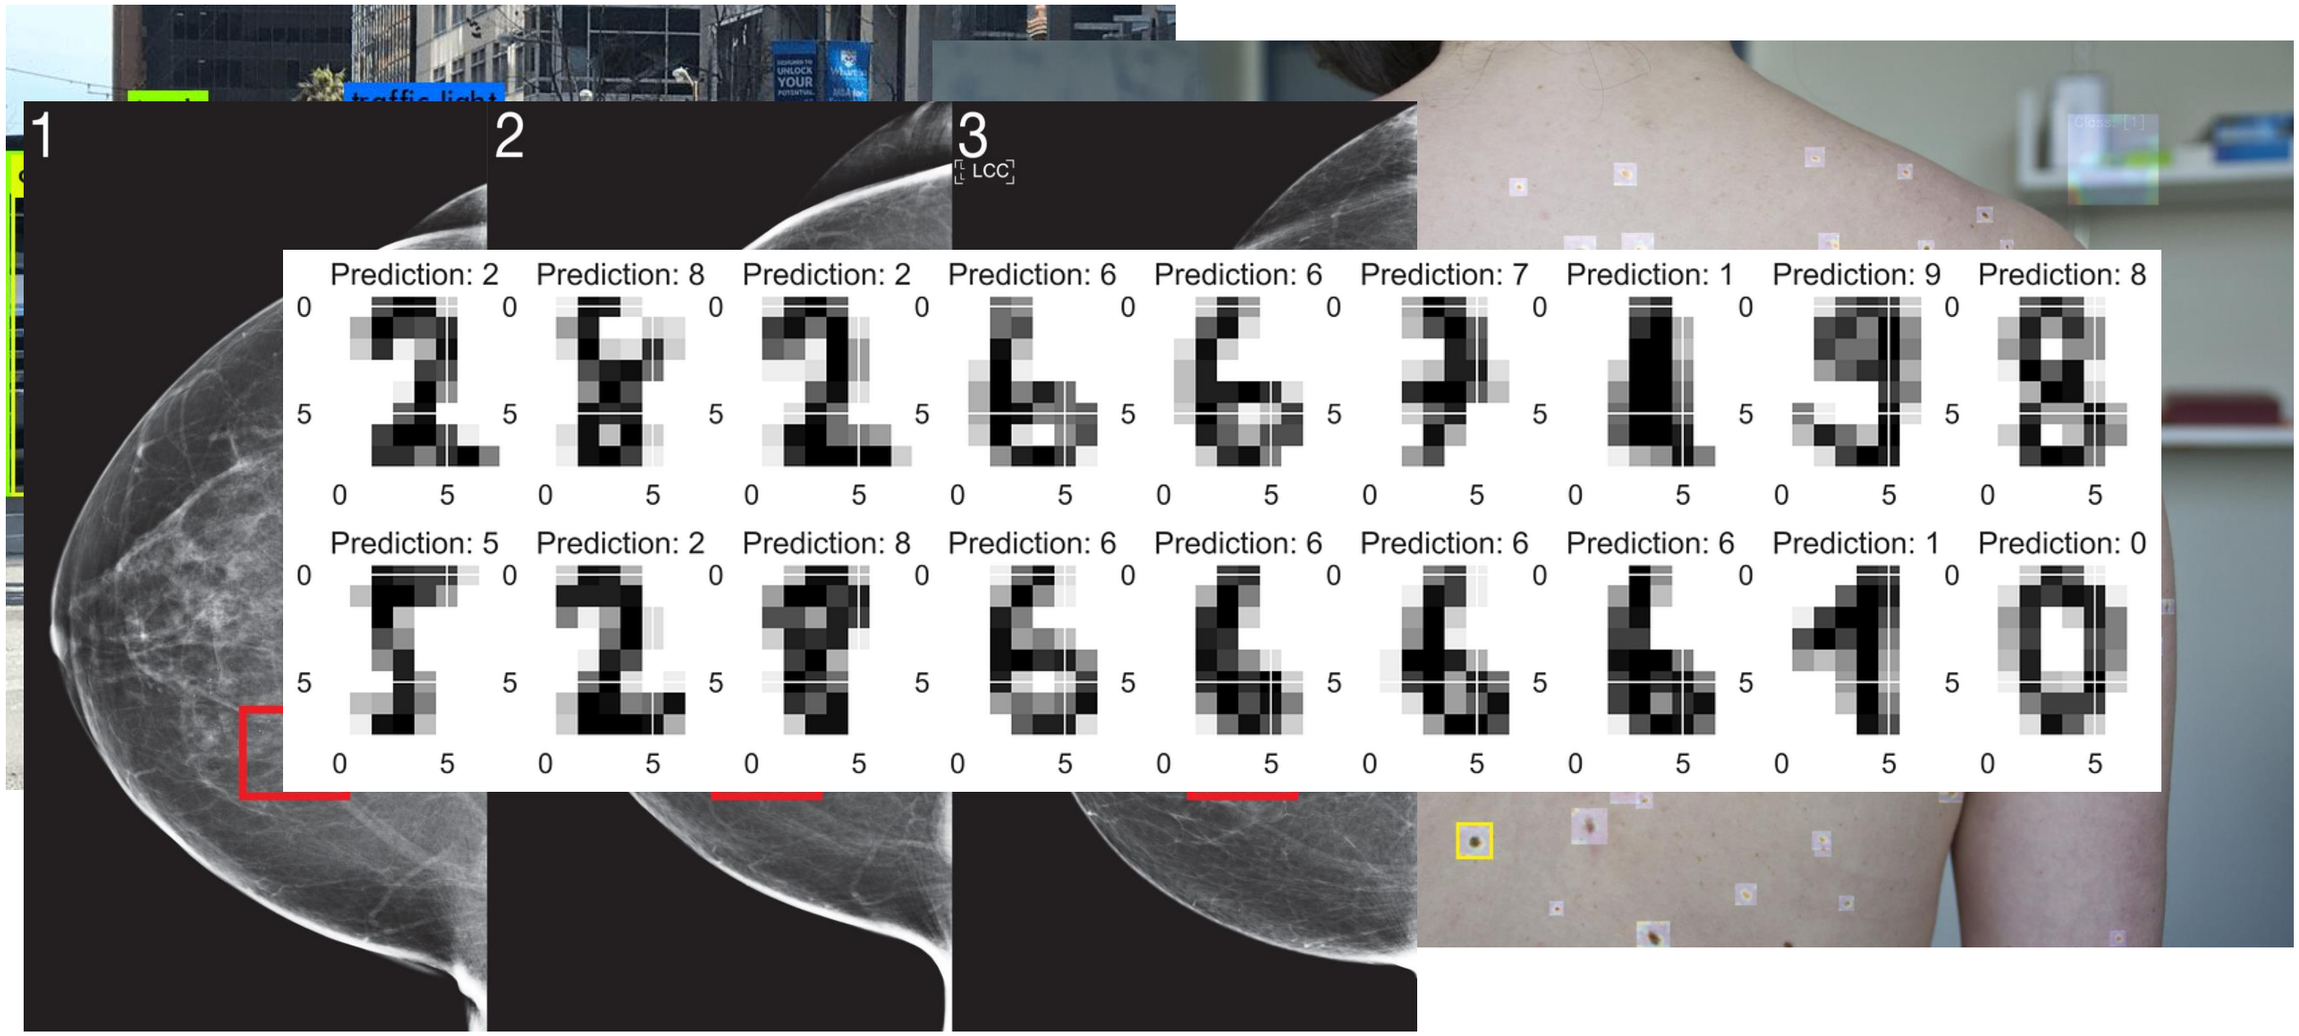




Computer Vision Models

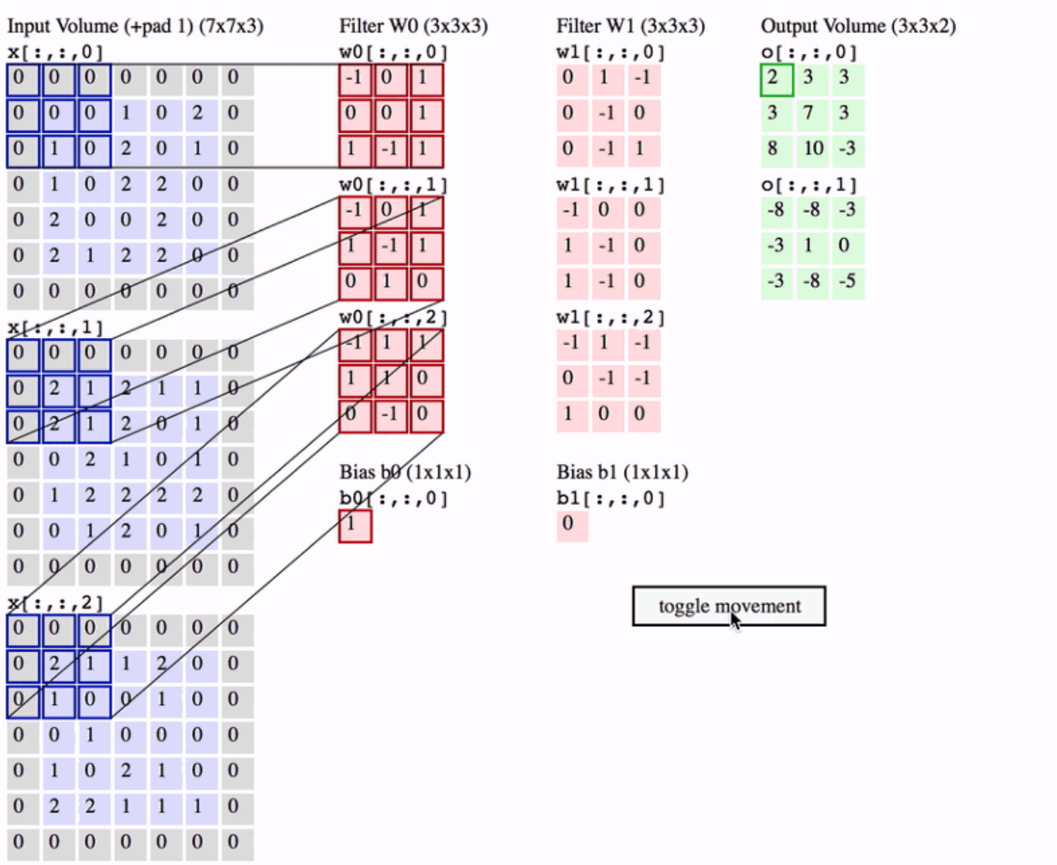


Computer Vision Models

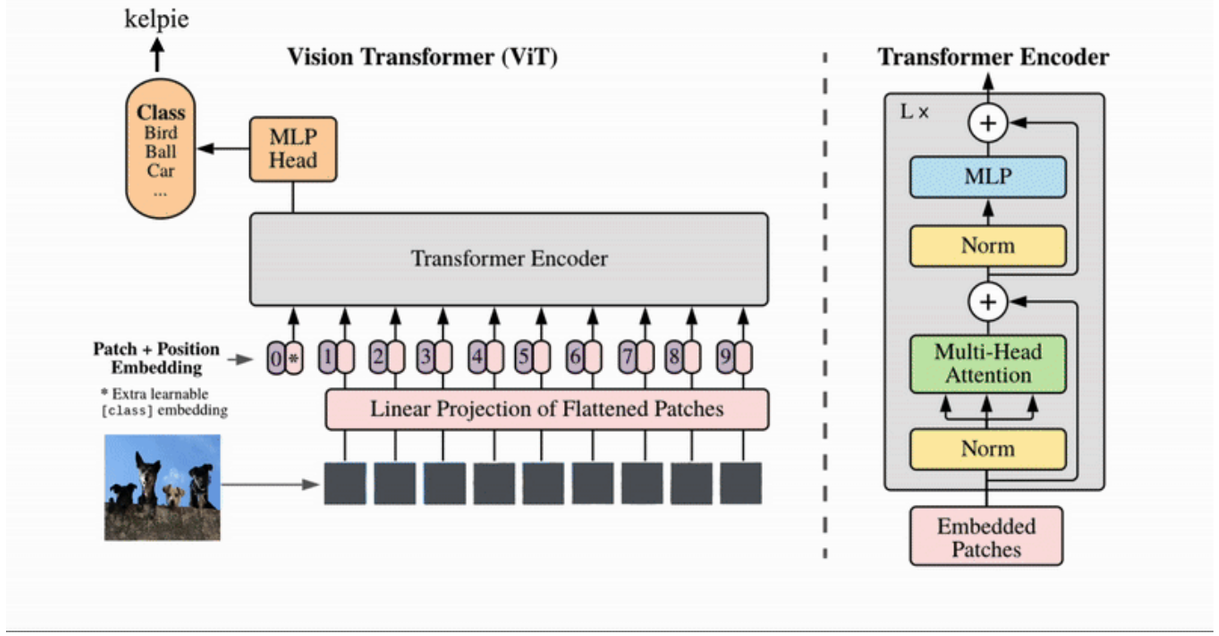


Two Main Architectures

Convolutional Neural Networks (CNNs)



Vision Transformers (ViTs)



NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
ICPE Higher Education Matters
University of Kentucky

Two Main Architectures

Convolutional Neural Networks (CNNs)

- Strengths:
  - Data Efficiency
  - Computational Efficiency
  - Excellent for Local Features
  - Strong Performance without massive-scale pre-training
- Weaknesses:
  - Limited Global Context
  - Rigid Architecture

Vision Transformers (ViTs)

- Strengths:
  - Superior Global Context
  - Scalability
  - Flexibility
  - Ideal for Multimodality
- Weaknesses:
  - Data-Hungry
  - Computationally Expensive

Convolutional Neural Networks

- How can a program determine if this is a cat, or not a cat?





Cats can look like this...



Cats can look like this... or like this!

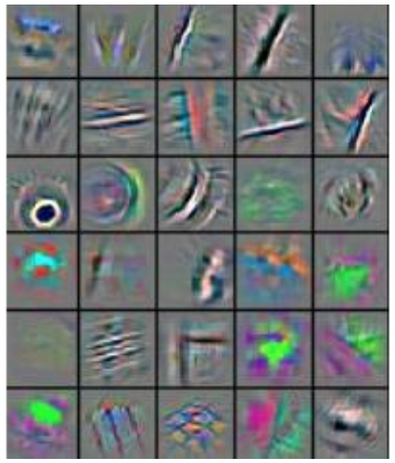


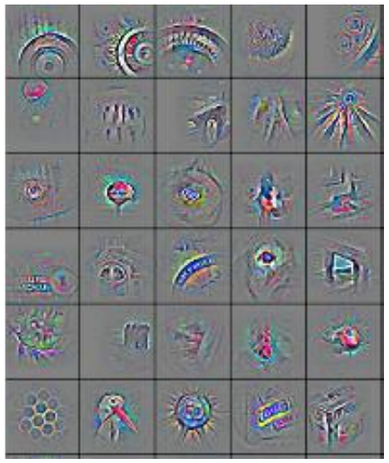


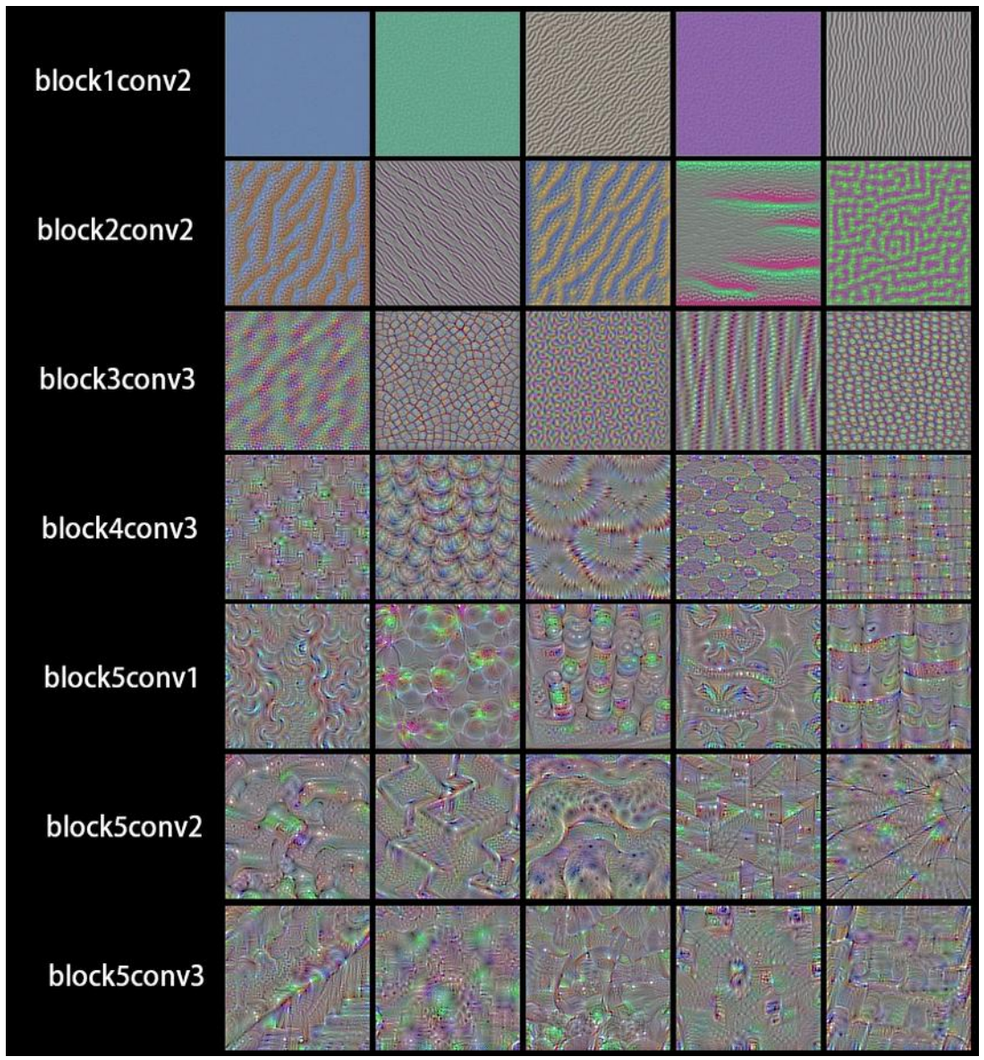
Convolutional Neural Networks (CNNs)

- Hierarchical feature extraction using filters (kernels)

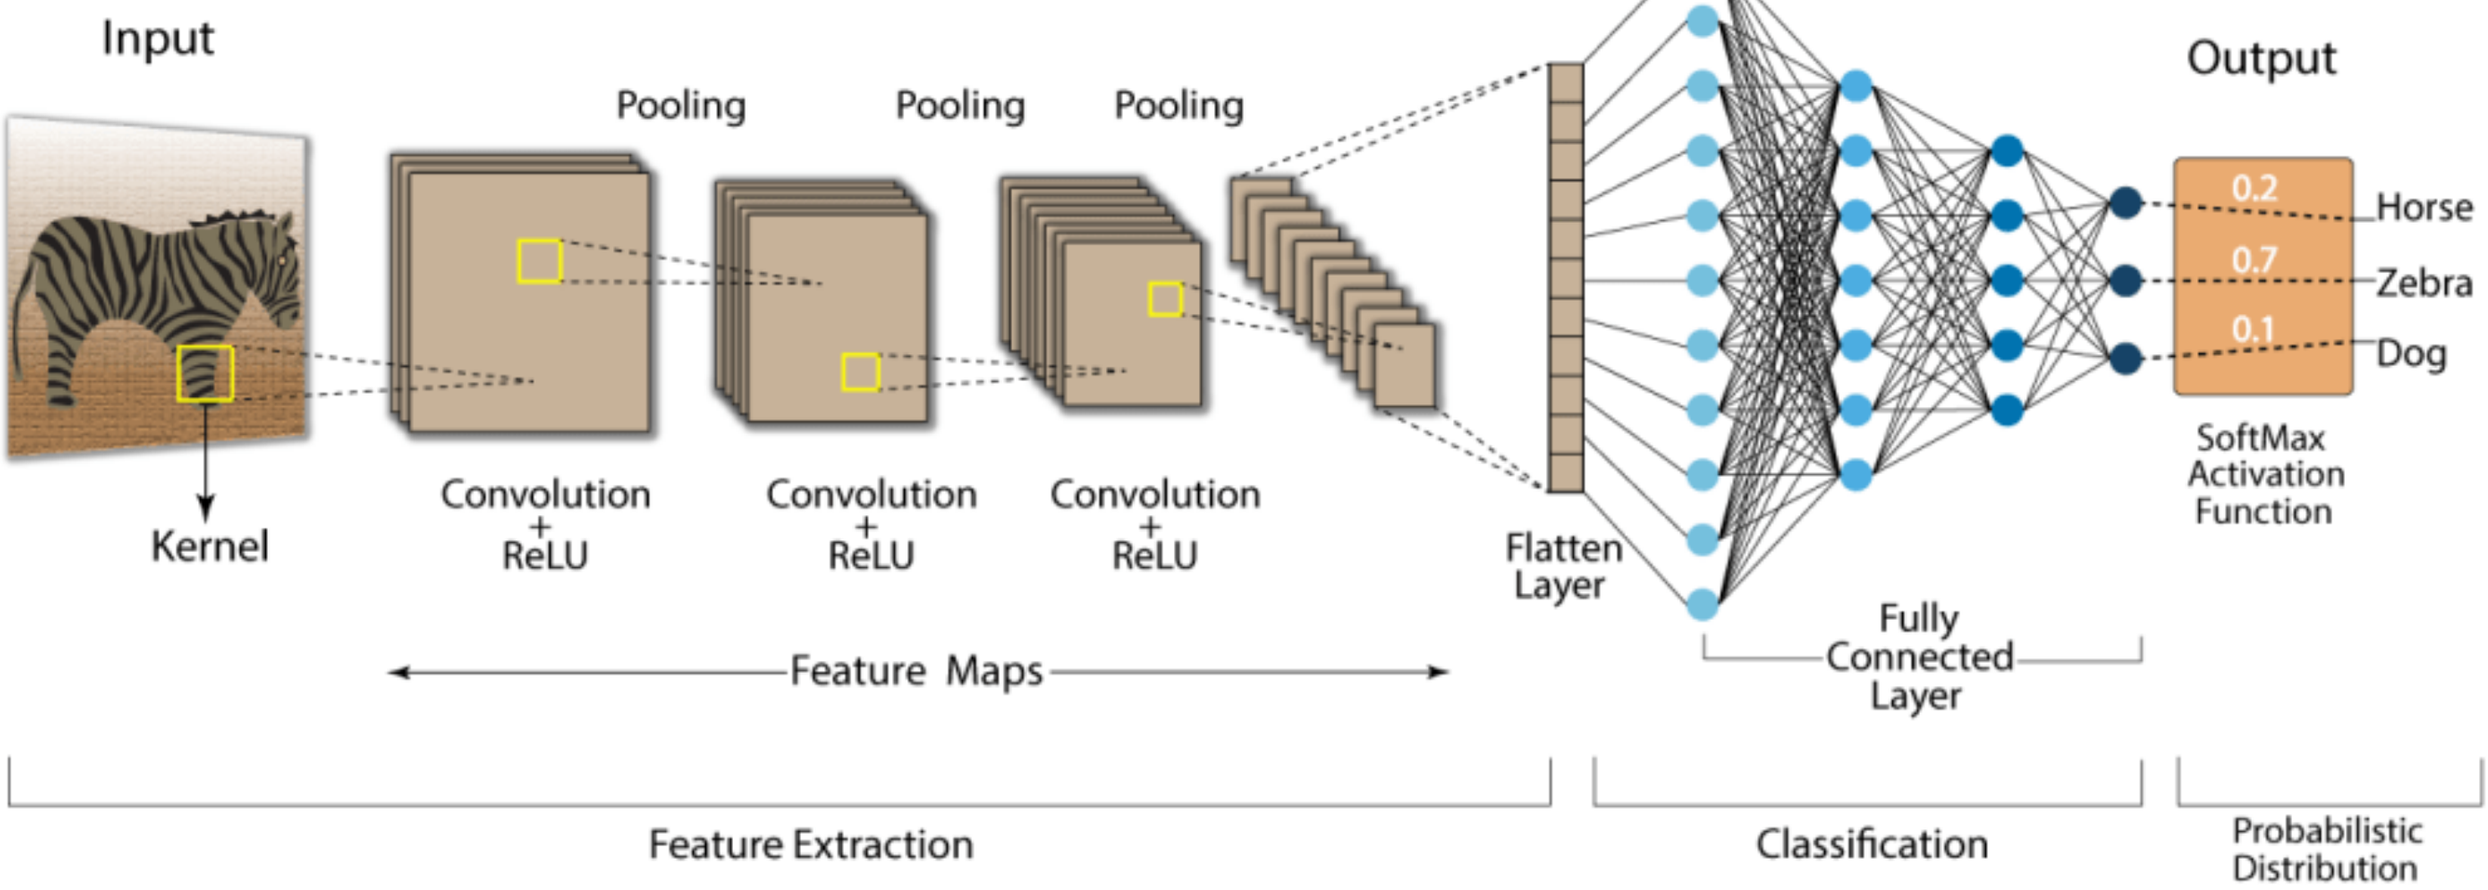








Convolution Neural Network (CNN)



NAIRR Pilot
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM
CPE Higher Education Matters
University of Kentucky

Convolutional Neural Network (CNN)



Convolutional Neural Networks (CNNs)

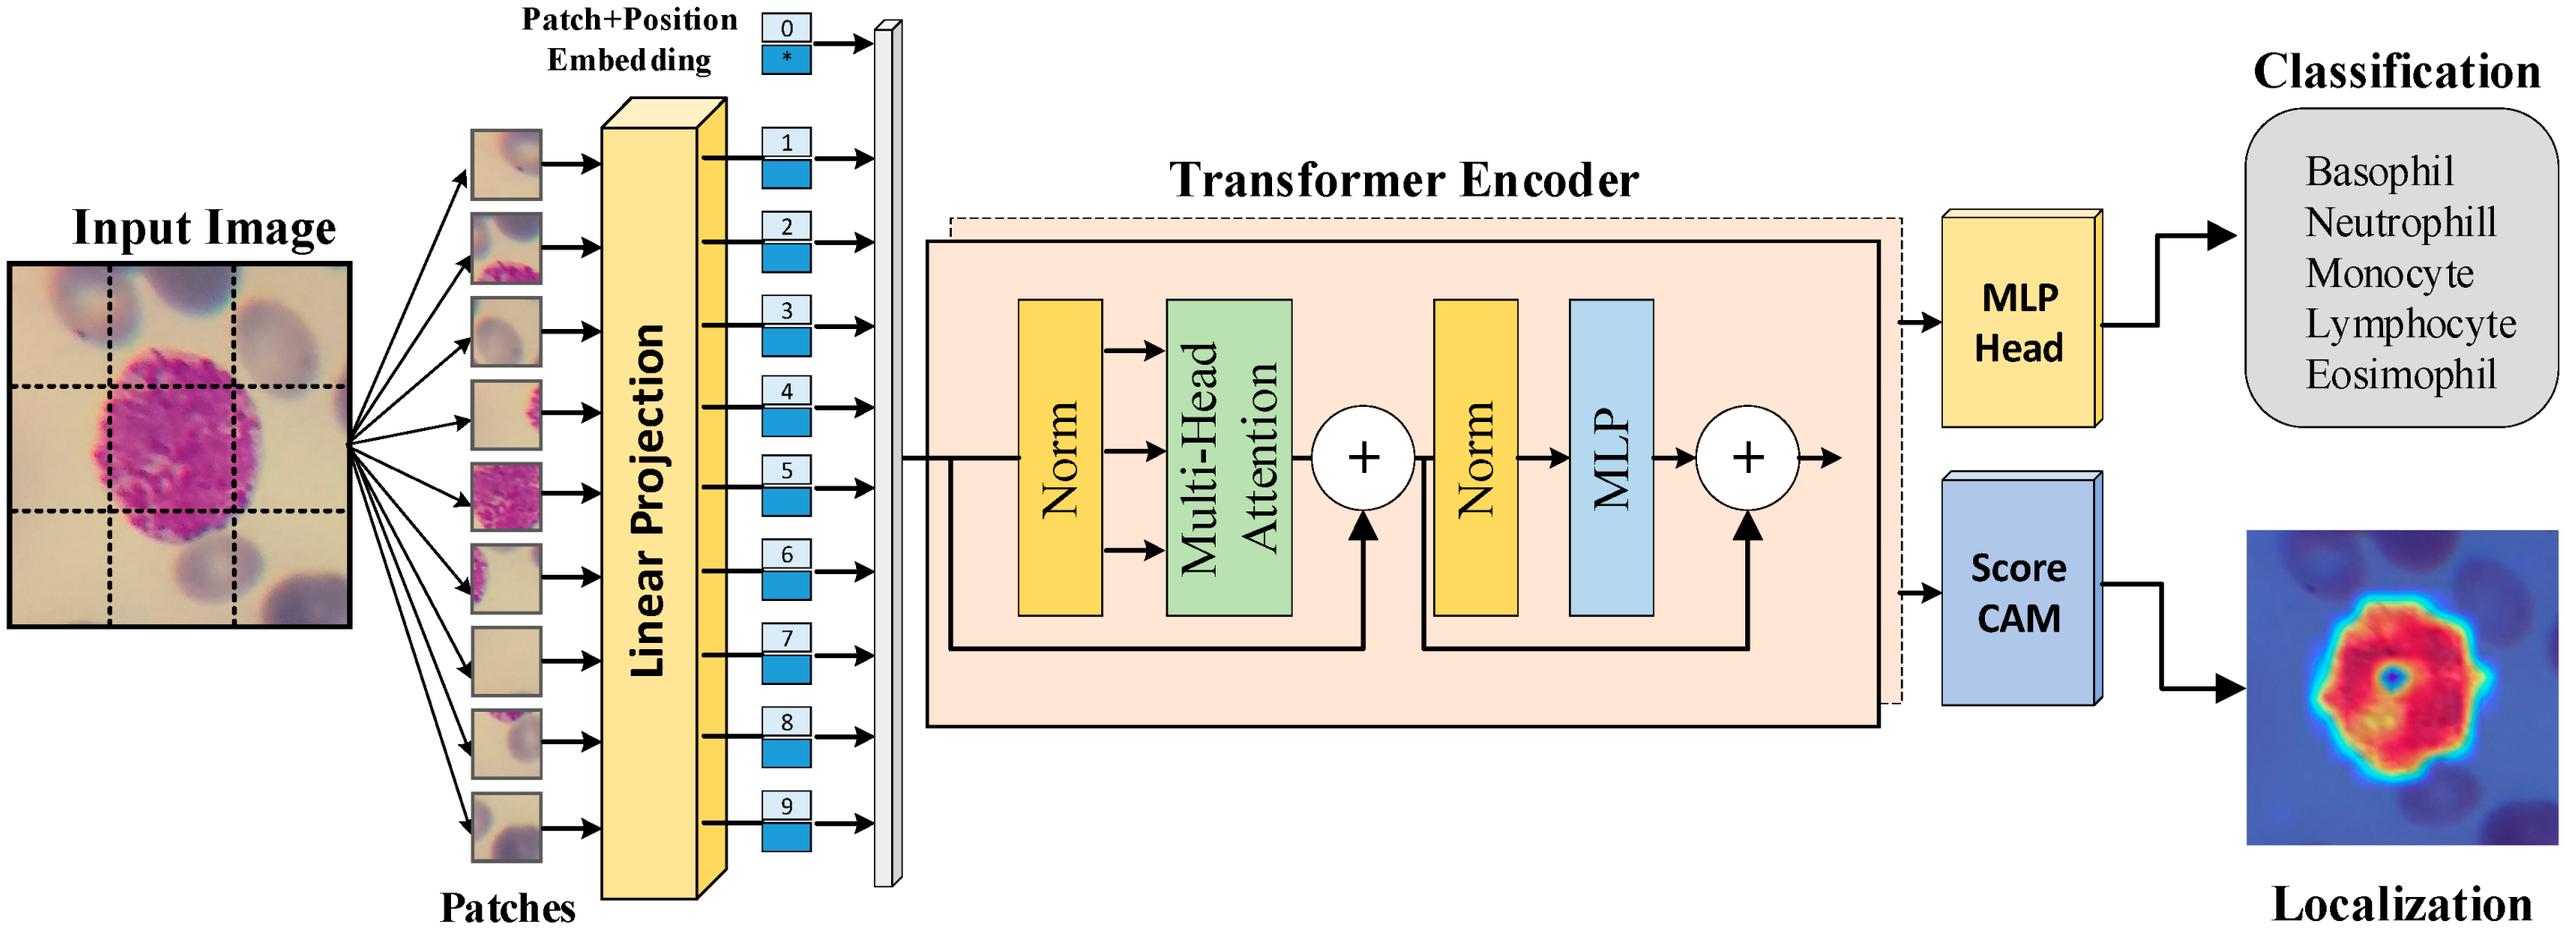
CNN Demo

NAIRR Pilot Resources

- NCSA Delta AI
  - Best For: Massive Scale Training
  - Why it helps: If a user wants to train a massive Vision Transformer (ViT) or a multimodal model (like CLIP or LLaVA) from scratch, standard clusters might be too slow. The Grace Hopper nodes here provide the massive bandwidth needed for huge image datasets.
- PSC Neocortex CS-2
  - Best For: Training from scratch
  - Why it helps: This system is architected specifically to eliminate bottlenecks in training large foundation models. It is a strong alternative for users whose models are too big for a single DGX node to handle efficiently.
- Lexset Seahaven:
  - Best For: Synthetic Data Generation
  - Why it helps: Vision models often fail because they lack diverse training data. This tool allows users to generate synthetic 3D environments and images to train models before testing them in the real world.

Vision Transformers (ViTs)

Unlike CNNs, which scan pixels with sliding windows, ViTs treat an image like a sentence of words



NAIRR Pilot

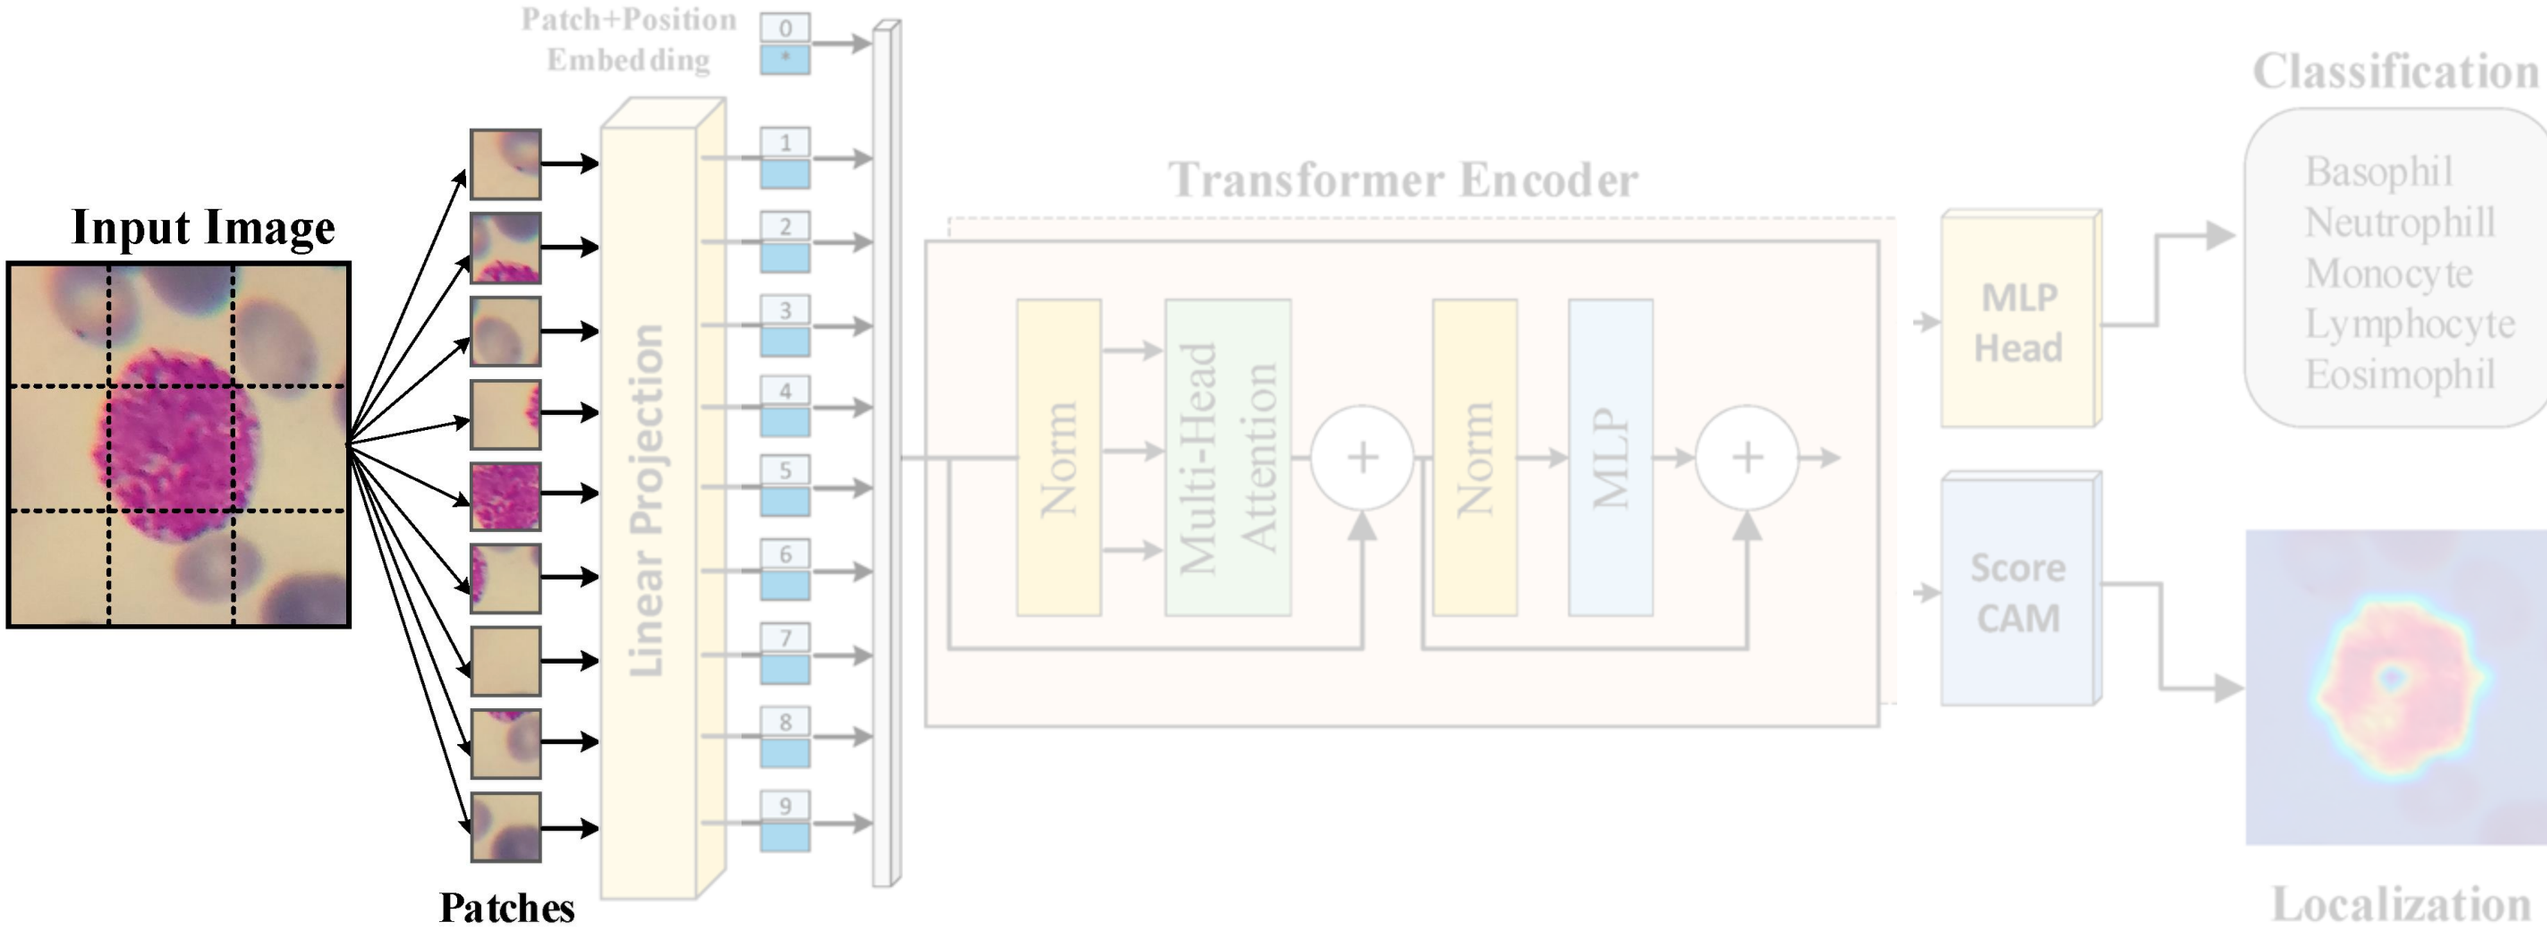
KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

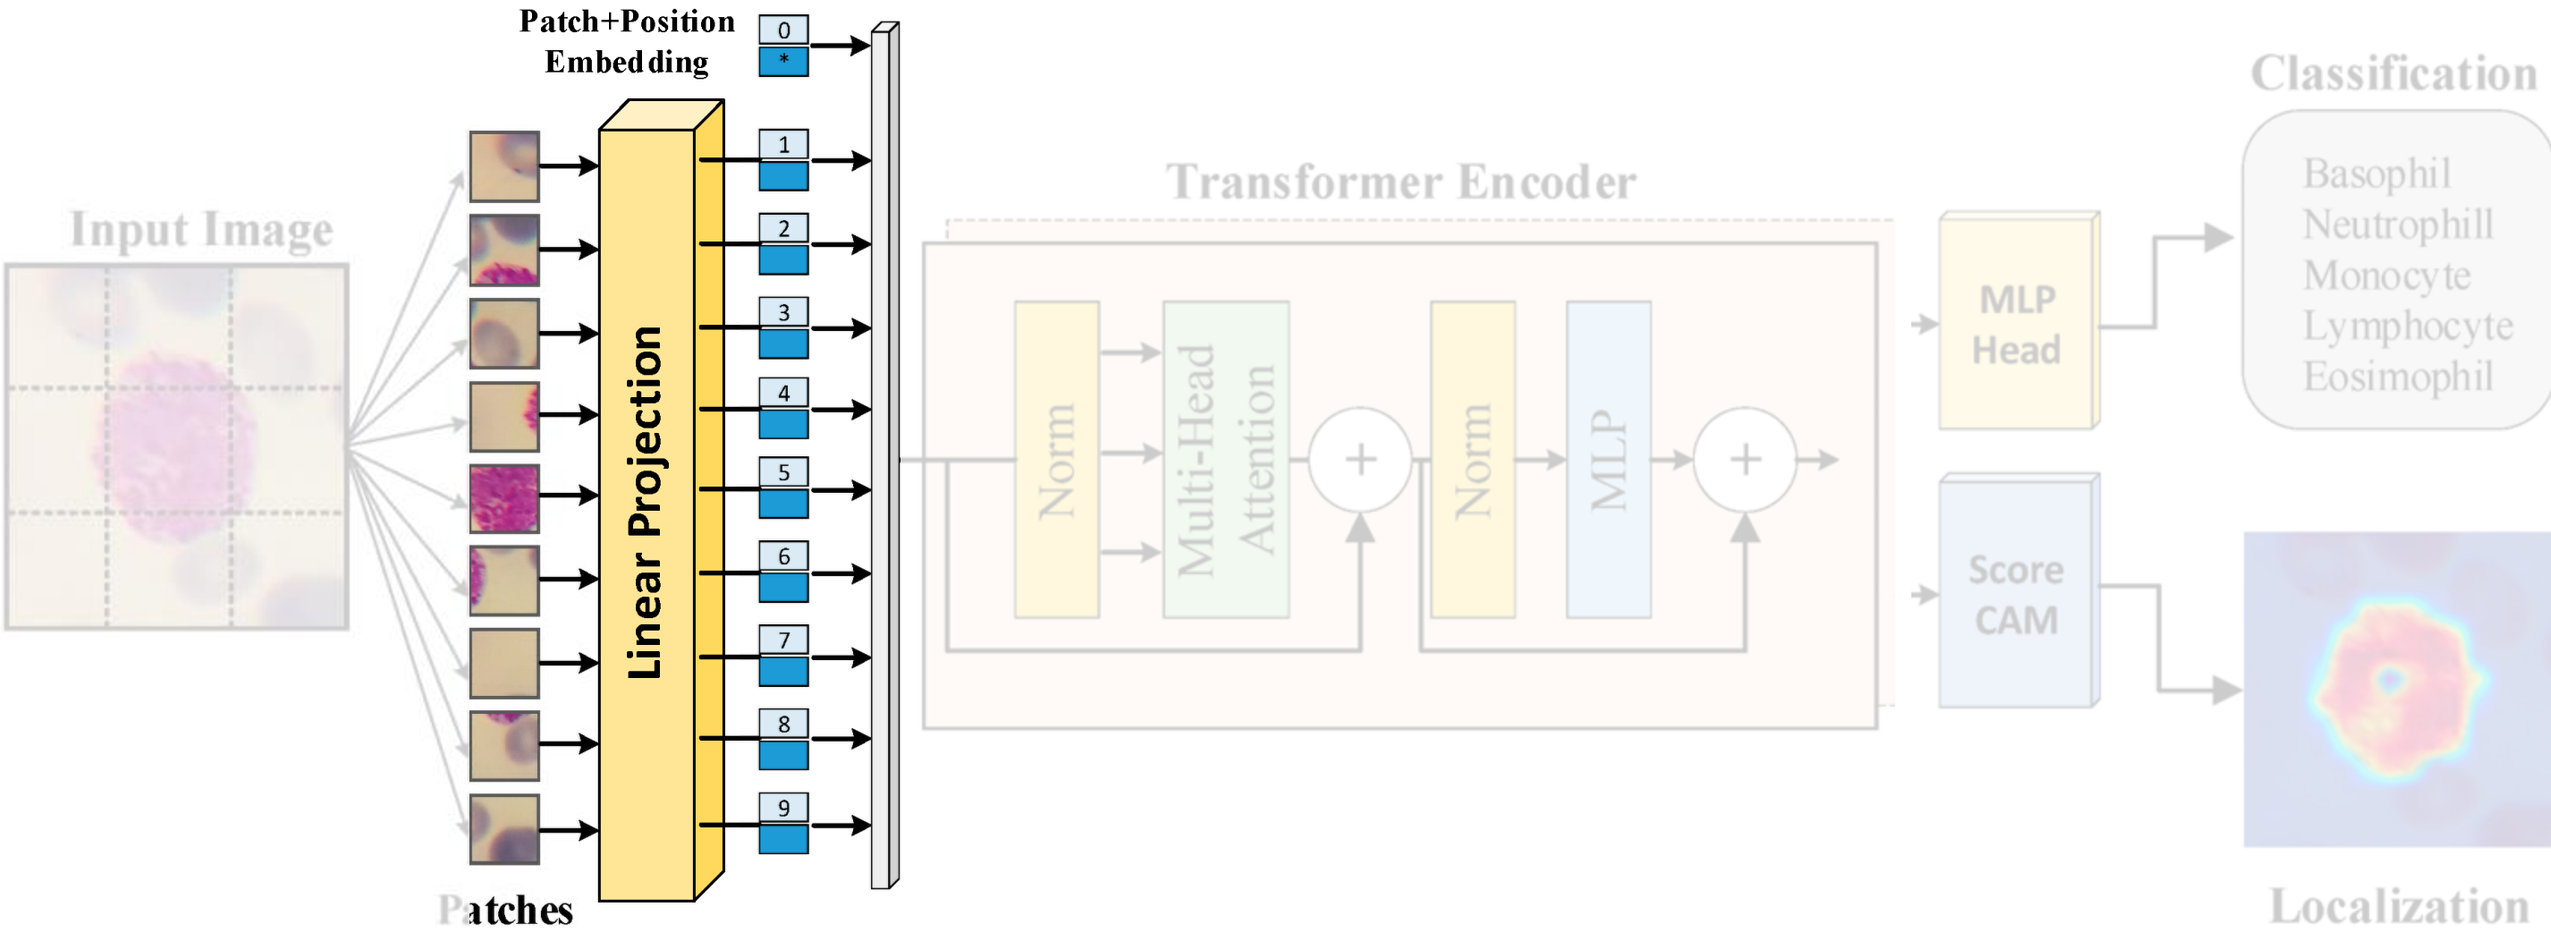
Vision Transformers (ViTs)

1. Input and Patching



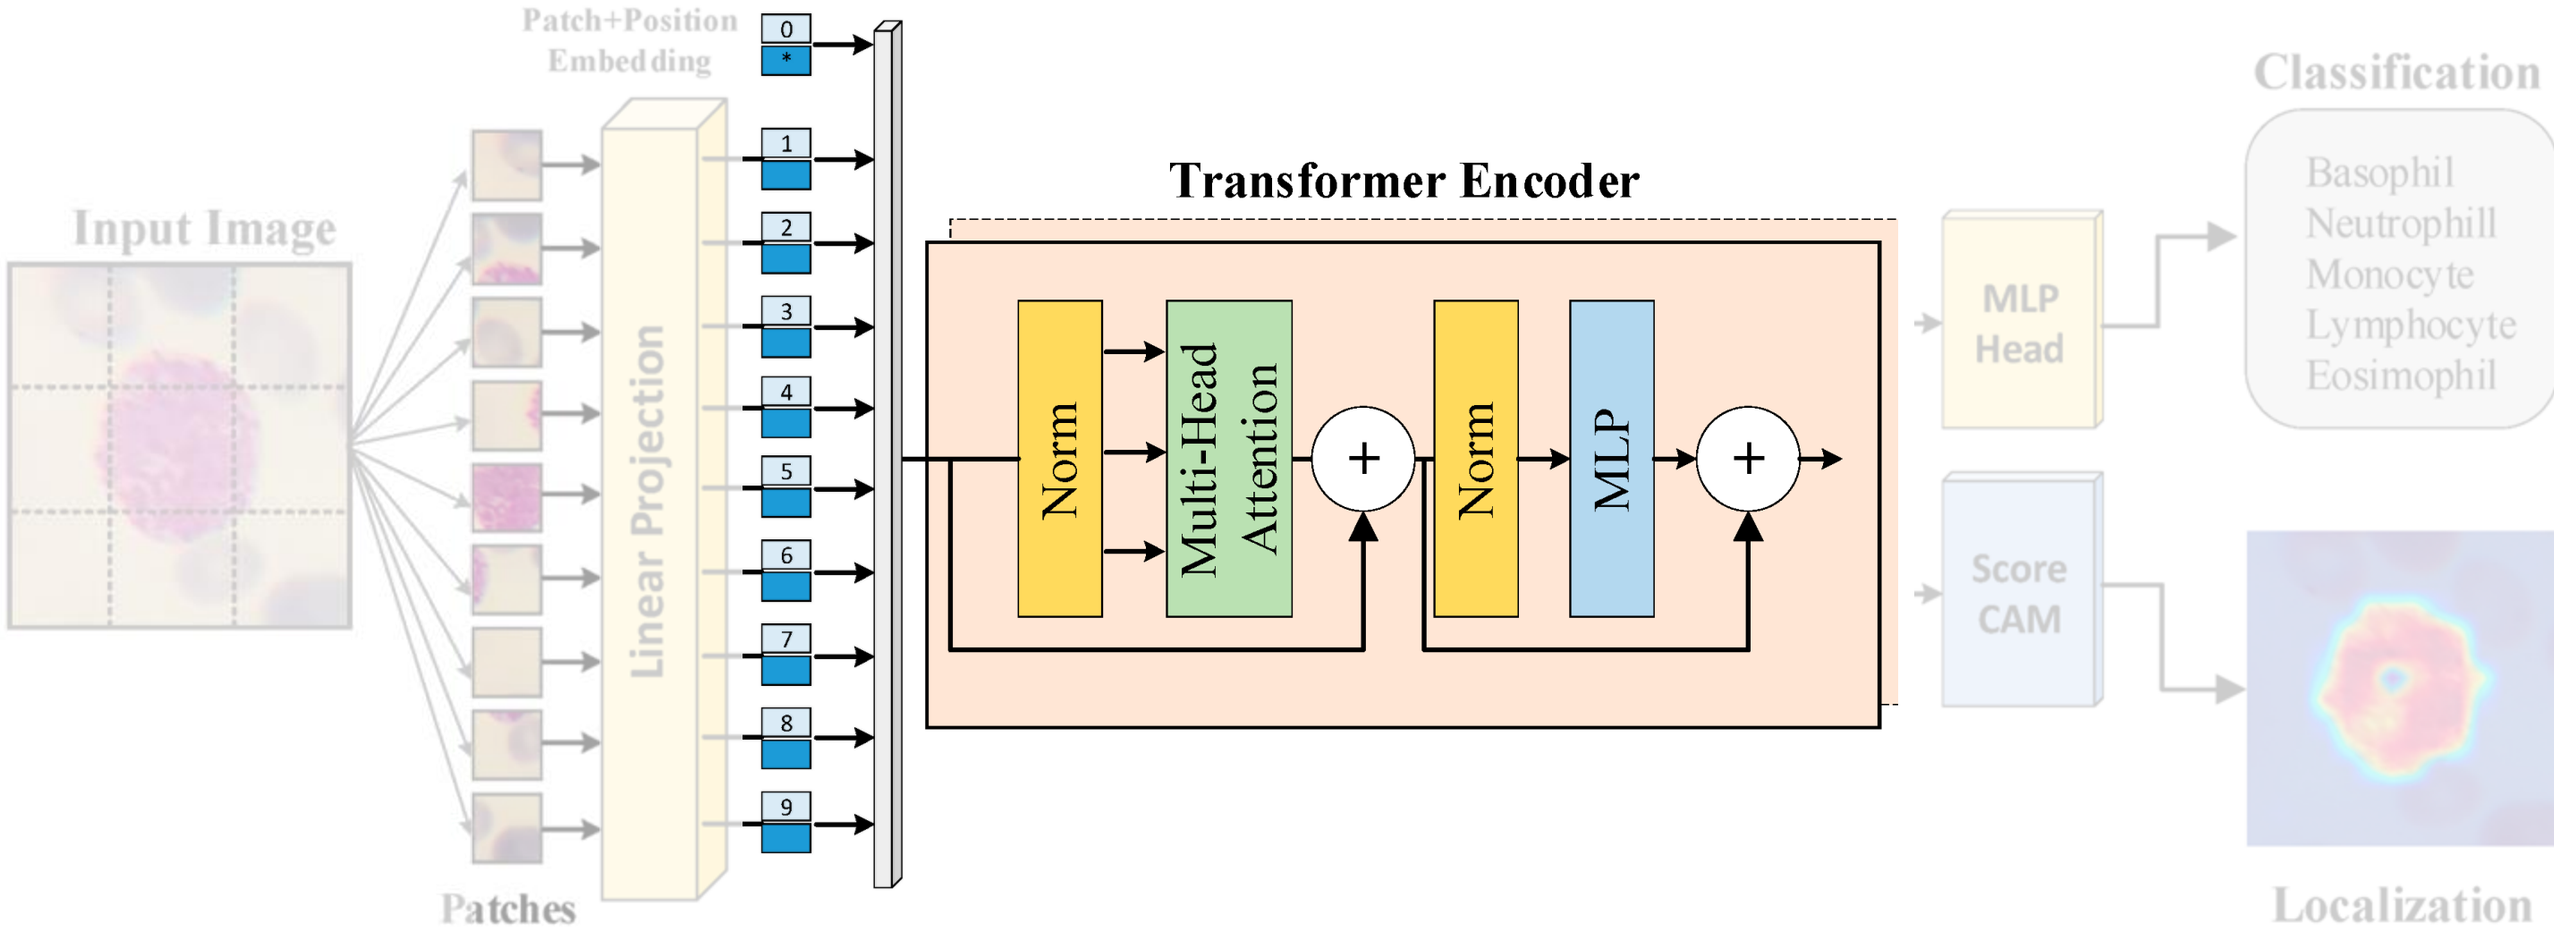
Vision Transformers (ViTs)

2. Linear Projection and Embeddings



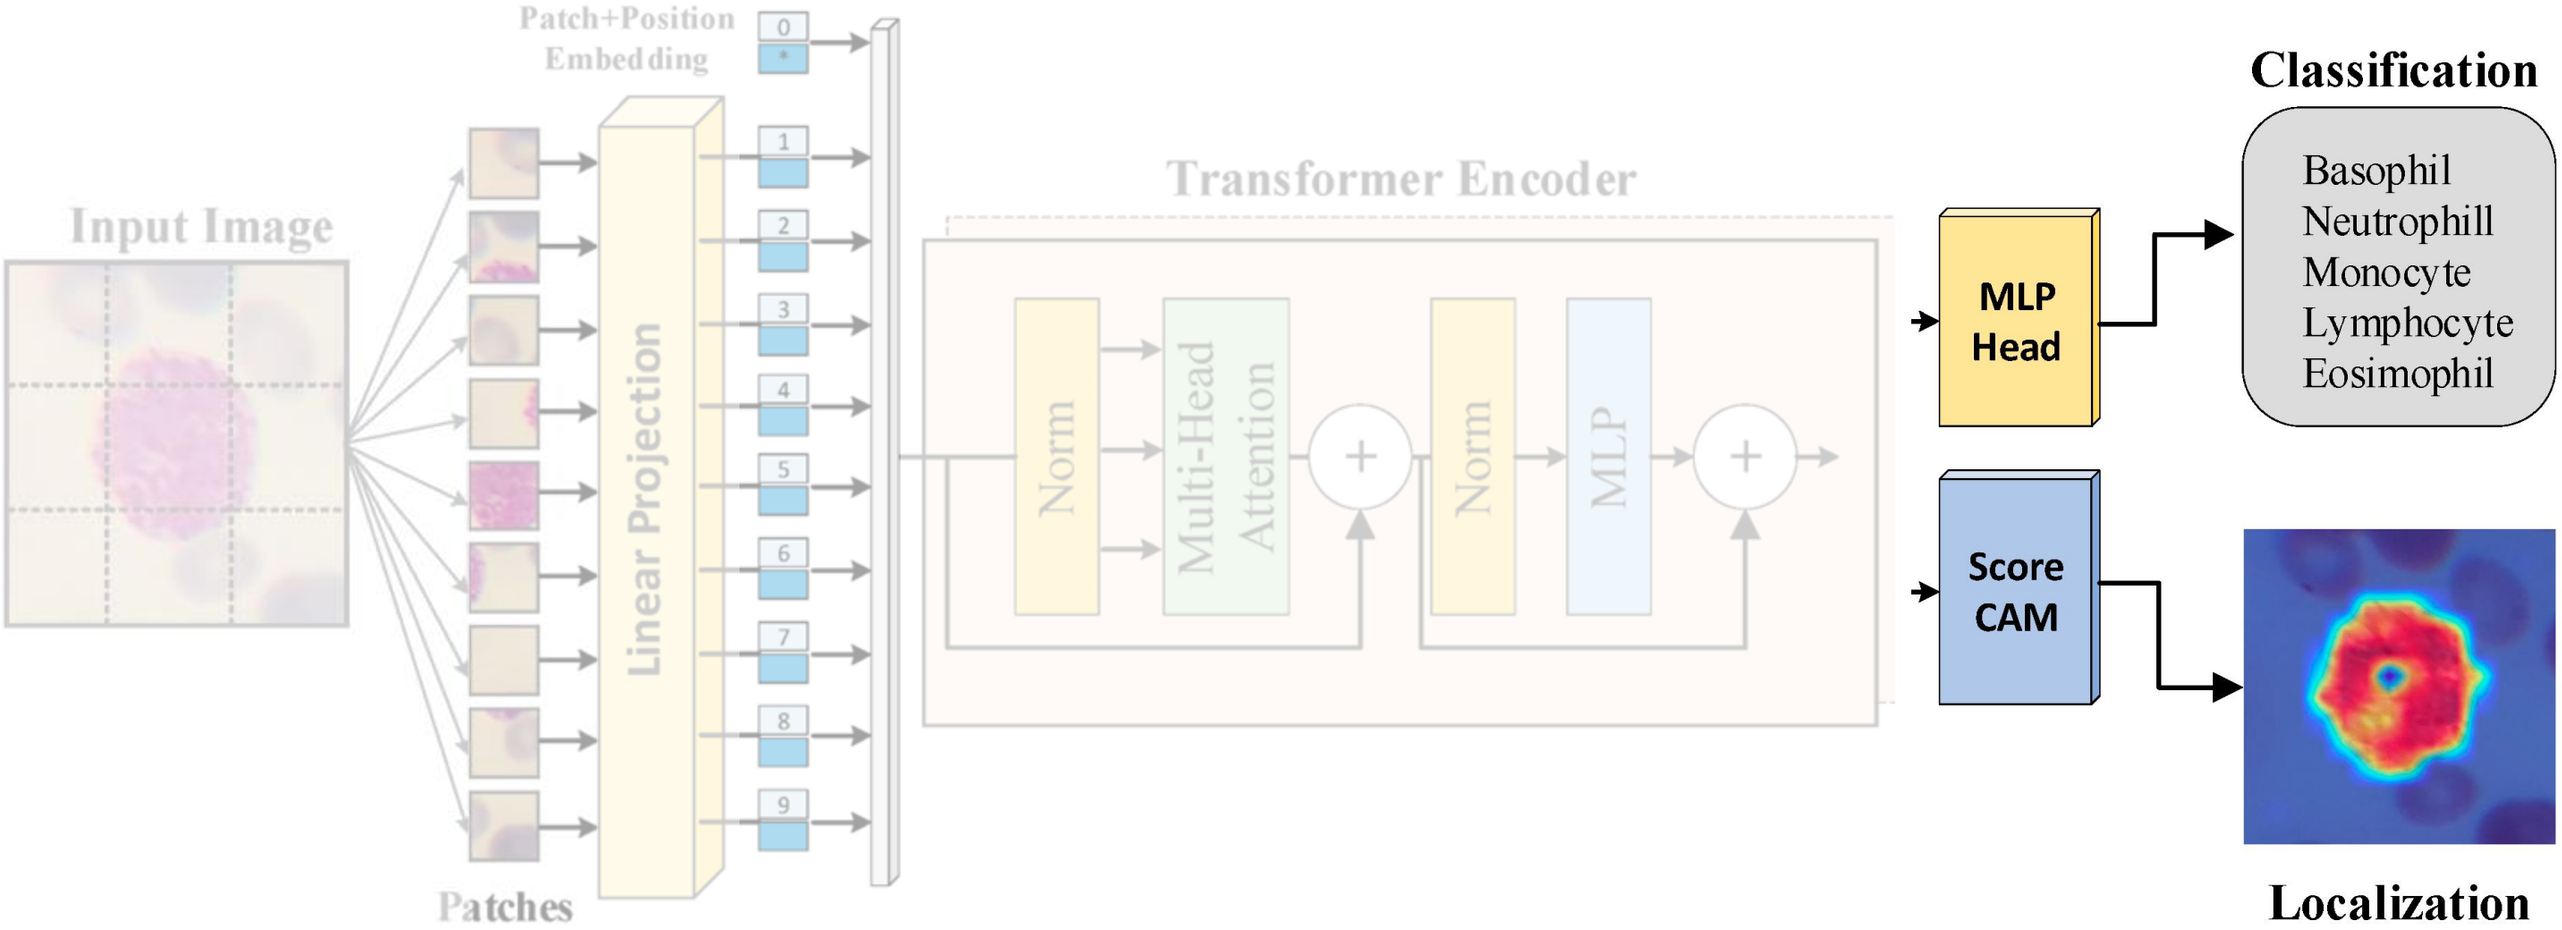
Vision Transformers (ViTs)

3. The Transformer Encoder



Vision Transformers (ViTs)

4. The Output



Vision Transformers (ViTs)



CNNs vs ViTs





Vision Transformers (ViTs)

ViT Demo

Vision Transformers (ViTs) vs CNNs

Summary:

1. Create patches from the input image
2. Create linear vectors (embeddings) from the patches
3. Perform Attention
4. Output results

NAIRR Pilot Resources

- Sage Edge Computing Platform
  - Best For: Edge Deployment and Real-Time Vision
  - Why it helps: Vision models are frequently deployed on cameras and robots, not just servers. Sage provides access to nodes with "Edge" GPUs (like NVIDIA Jetson or T4) across the US, allowing researchers to test how their models handle real-world sensor data and lower power constraints.

NAIRR Pilot Resources

- NCSA Delta AI
  - –Best For: Massive Scale Training
  - –Why it helps: If a user wants to train a massive Vision Transformer (ViT) or a multimodal model (like CLIP or LLaVA) from scratch, standard clusters might be too slow. The Grace Hopper nodes here provide the massive bandwidth needed for huge image datasets.
- PSC Neocortex CS-2
  - –Best For: Training from scratch
  - –Why it helps: This system is architected specifically to eliminate bottlenecks in training large foundation models. It is a strong alternative for users whose models are too big for a single DGX node to handle efficiently.

Coffee Break



Join us in the lobby

Ask questions & network with your peers!

Reconvene by 2:45 PM

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

AI Pipelines

Mitchell Klusty, BS

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

What are AI Pipelines and Why are They Important?

- Standardized processes for application management
- Integrating AI with your unique data and systems
- Deploying appropriate resources

- *Each project will have unique needs*
- What you need to do will be dependent on your project's goals and resources

- *Identify where AI is useful, and where it isn't*

Consider:

Where does AI fit into my workflow?

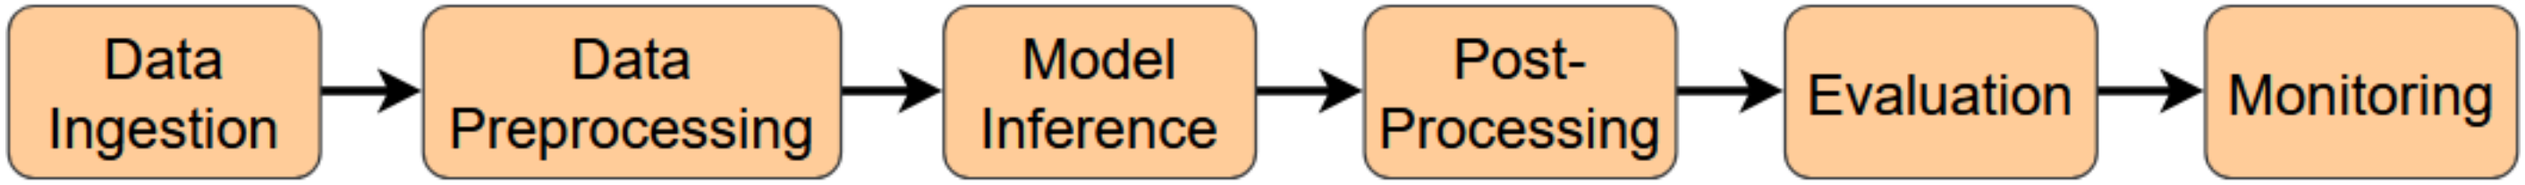
- How would introducing AI benefit individual features?

Not every problem needs an AI solution

- AI can be costly
- Is there another solution that will work just as well?

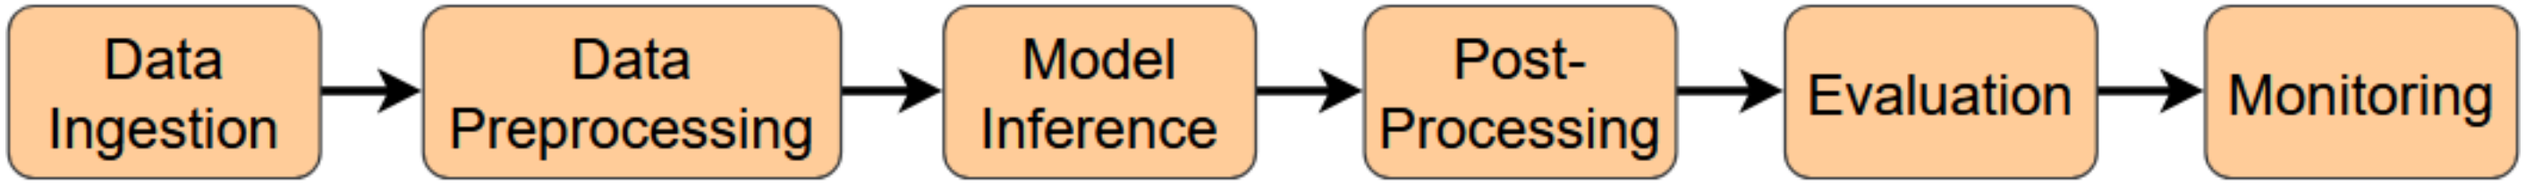
Components of an AI Workflow

- AI needs data
  - How to handling project-specific data?
- How to support model inference?
- How to evaluate the performance?
- Where to look for performance issues/improvements



Components of an AI Workflow

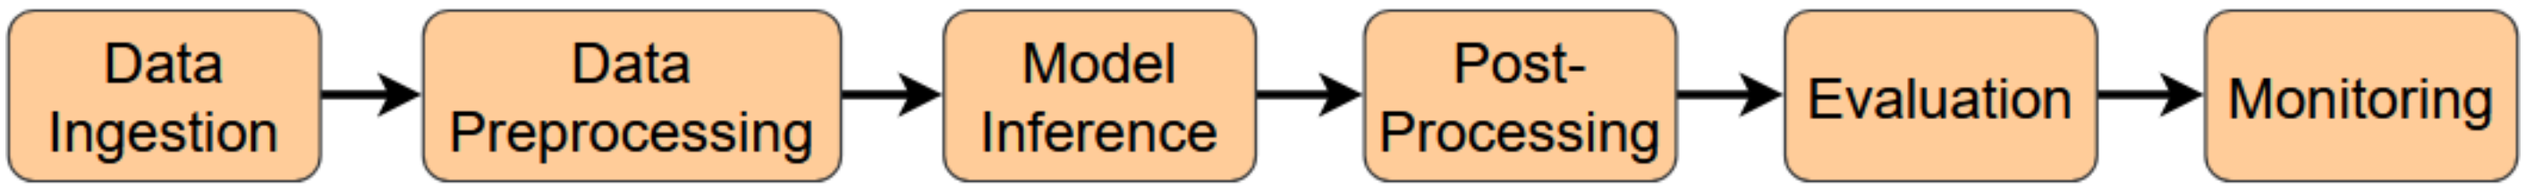
- AI needs data
  - How to handling project-specific data?
- How to support model inference?
- How to evaluate the performance?
- Where to look for performance issues/improvements



Every piece is dependent on what comes before,
and depended on by what comes after

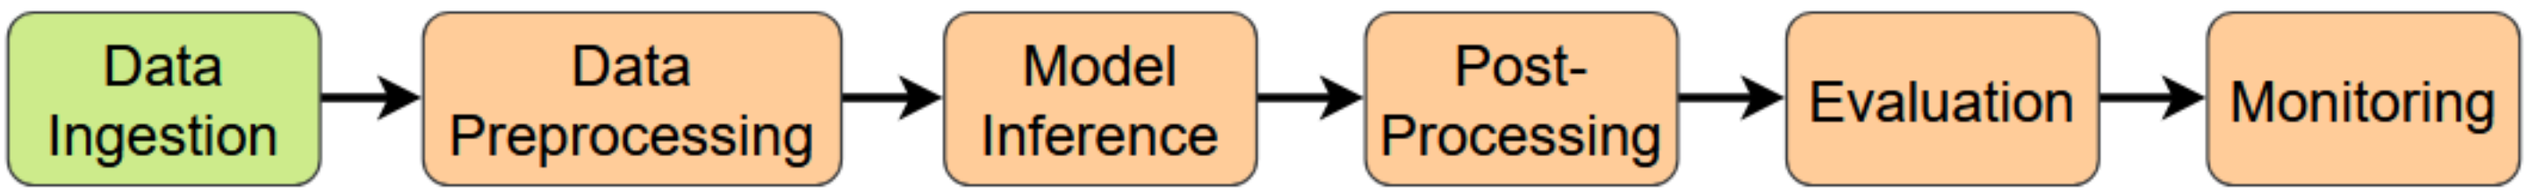
Components of an AI Workflow

- AI Agricultural Assistant
  - Sources information from University of Kentucky College of Agriculture resources



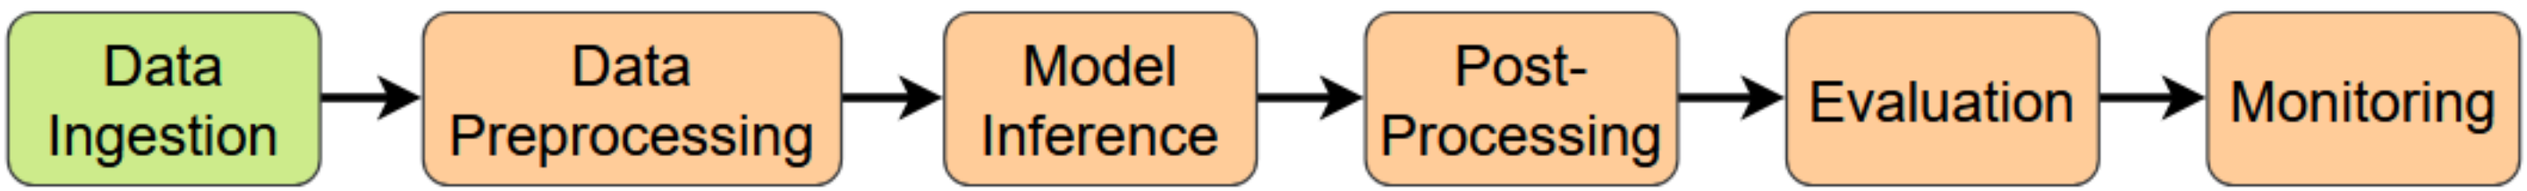
Components of an AI Workflow – Data Ingestion

- What data do you need?
- How will users supply data?
- How to collect data?
- **Structure data to be adaptive and reusable**
  - Exact goals/conditions may change



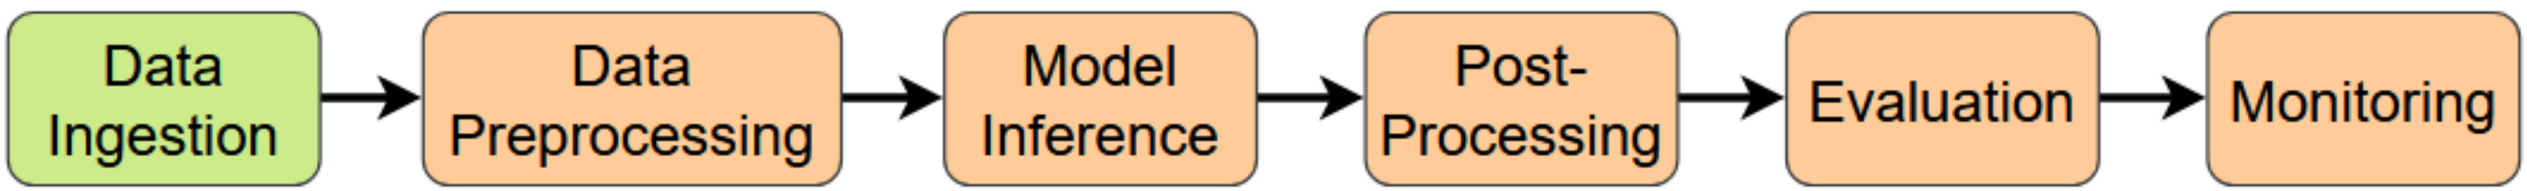
Components of an AI Workflow – Data Ingestion

- Do you need to make a public facing interface?
  - Web Hosting
  - Application Programming Interface (API)
- How much data to ingest
  - Network loads
  - Storage considerations



Components of an AI Workflow – Data Ingestion

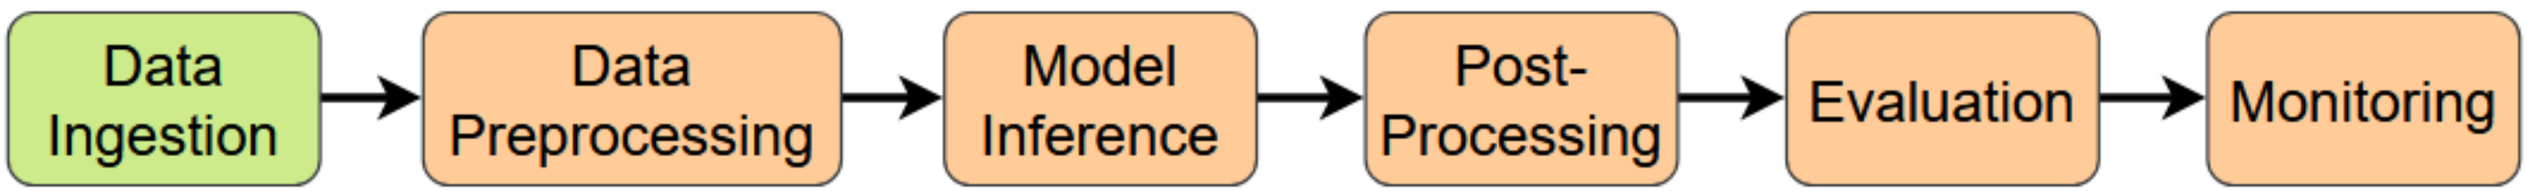
- UK Ag web sources
  - Find URLs
  - Pull content



Components of an AI Workflow – Data Ingestion

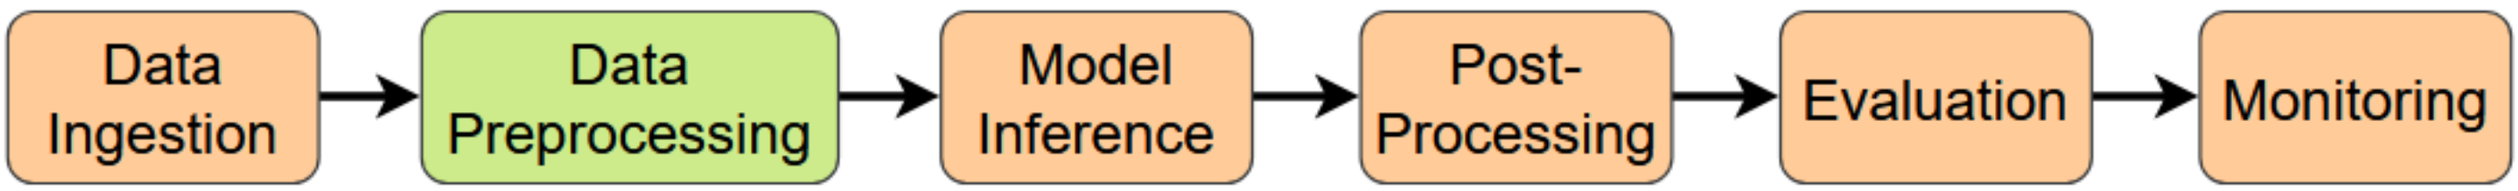
Common Oversights or Problems

- Not considering network constraints
- Not considering storage constraints



Components of an AI Workflow – Data Preprocessing

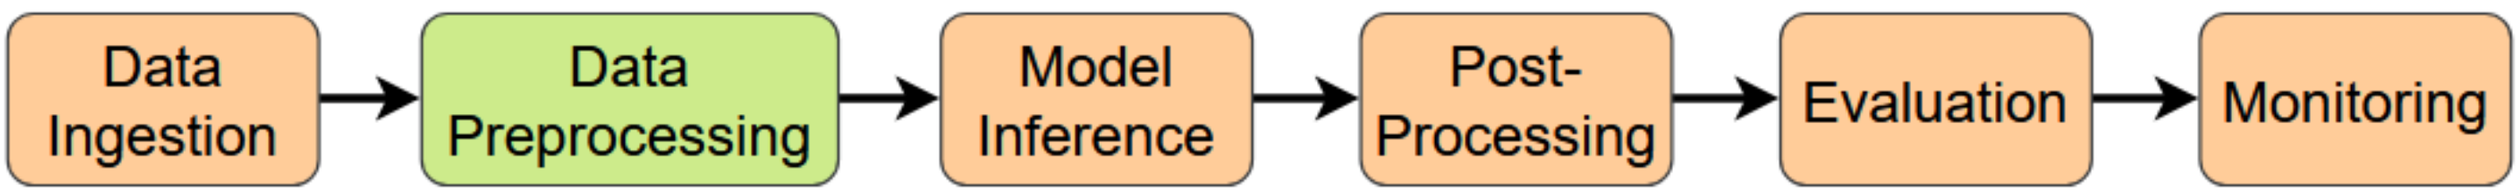
- Massive datasets exist that aren't "AI-ready"
  - Large sets of images, no labels/annotations
  - Large bodies of text
- How to prepare data for training?
- How to best perform inference on data?
  - Resources for inference can be expensive
  - How to make most efficient use of those resources when applying data?



Components of an AI Workflow – Data Preprocessing

BLUE

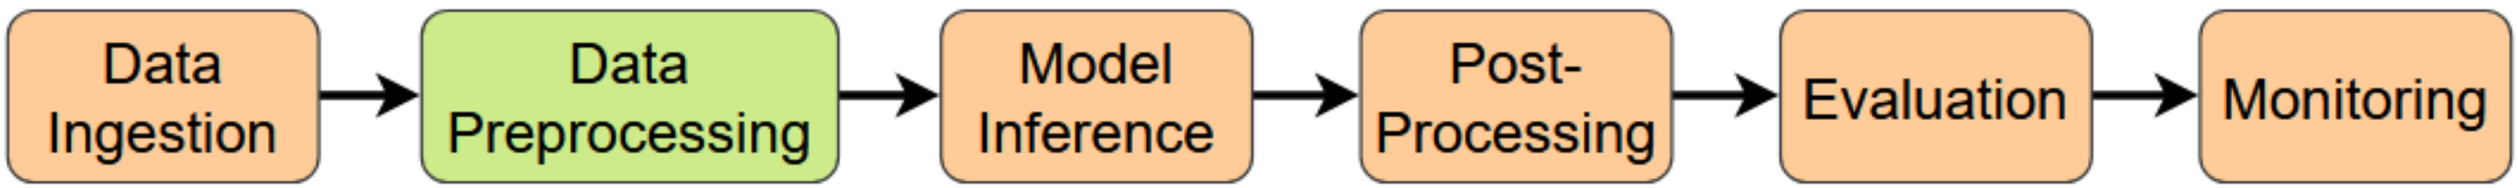
- Embed text from websites
- Store in Vector DB
  - Must be efficient to pull from



Components of an AI Workflow – Data Preprocessing

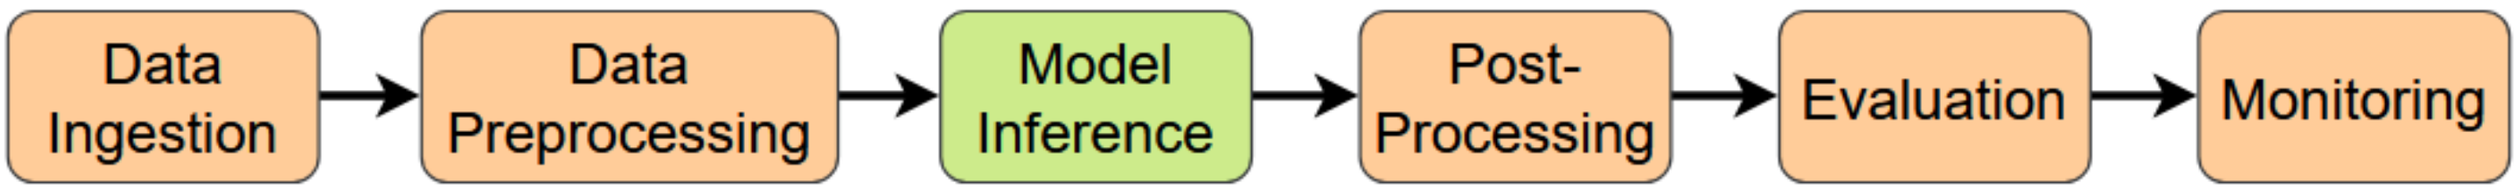
Common Oversights or Problems

- Not storing data efficiently
- Not making data easily accessible for next steps
- Garbage In, Garbage Out
  - Applies to both training and inference
  - Must ensure data is cleaned after intake



Components of an AI Workflow – Model Inference

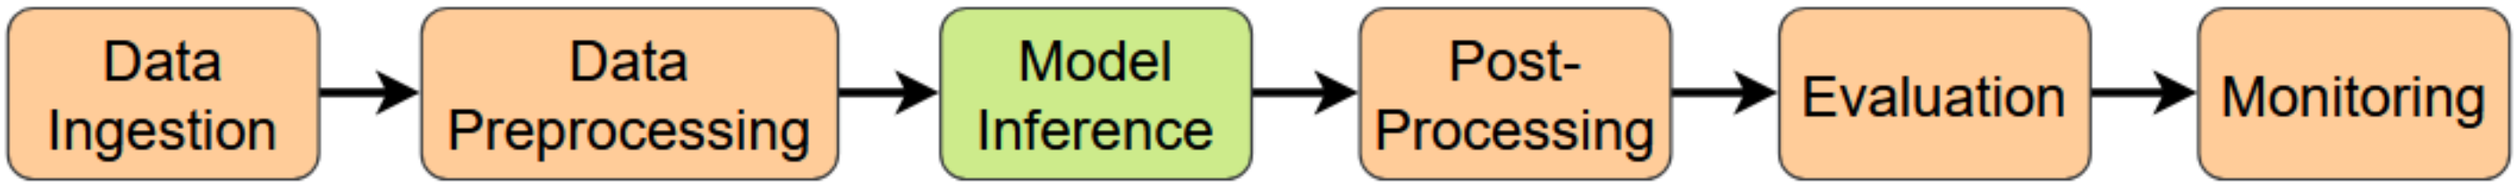
- What type of model are you running?
- What resources does that need?



Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

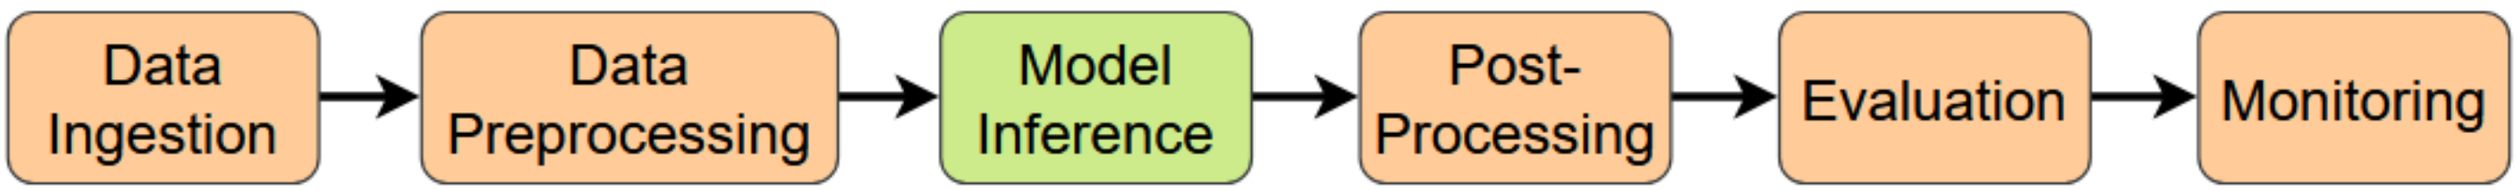
- LLMs require higher GPU memory to host
  - Dependent on size of model
    - Number of parameters
    - Precision/Quantization



Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- LLMs require higher GPU memory to host
  - Dependent on size of model
    - Number of parameters (internal weights)
      - Llama 3 8B, 70B, 405B
      - DeepSeek-V3 37B active, 671 total
      - Gemini 3 estimated 20B active, 1 trillion total
  - Precision/Quantization

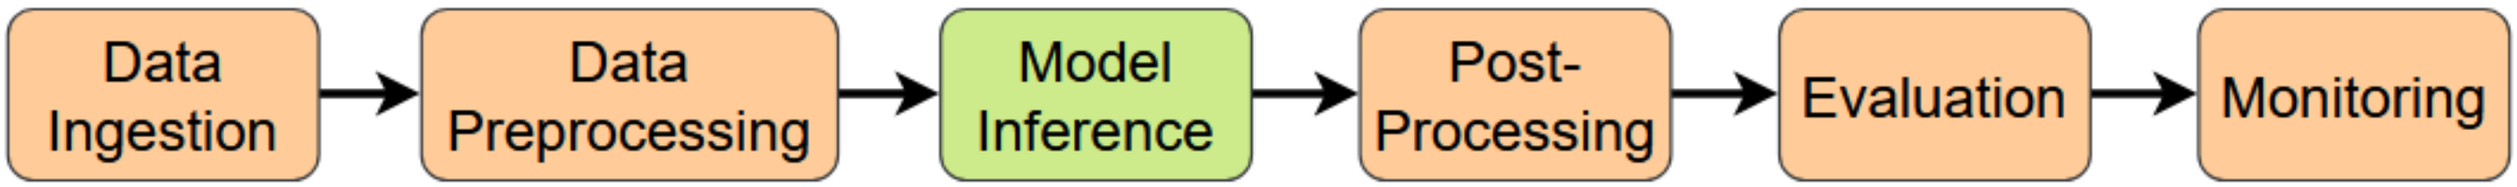


Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- LLMs require higher GPU memory to host
  - Dependent on size of model
    - Number of parameters
    - Precision/Quantization
    - How many bits in each parameter

	FP32	FP16	INT8	INT4
Bits per parameter	32	16	8	4
Bytes per parameter	4	2	1	0.5

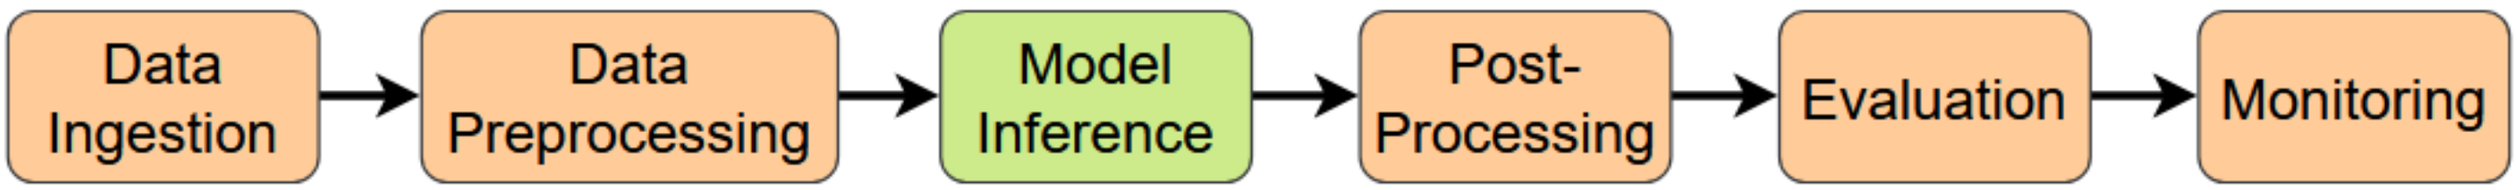


Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- LLMs require higher GPU memory to host
  - Dependent on size of model
    - Number of parameters
    - Precision/Quantization

$$Required\ VRAM\ (GB) \approx \frac{parameters * bytes\ per\ parameter}{1\ billion}$$
$$\approx \#\ billion\ parameters * bytes\ per\ parameter$$

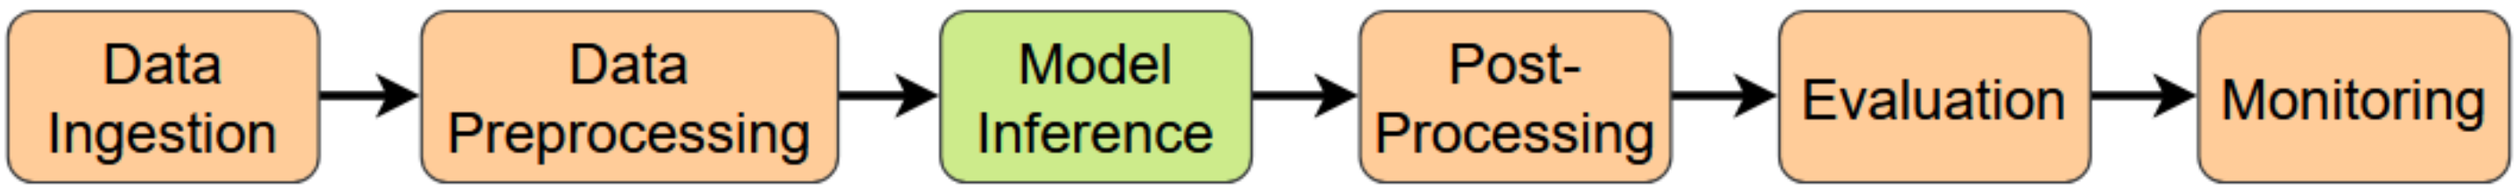


Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- LLMs require higher GPU memory to host
  - Dependent on size of model
    - Number of parameters
    - Precision/Quantization

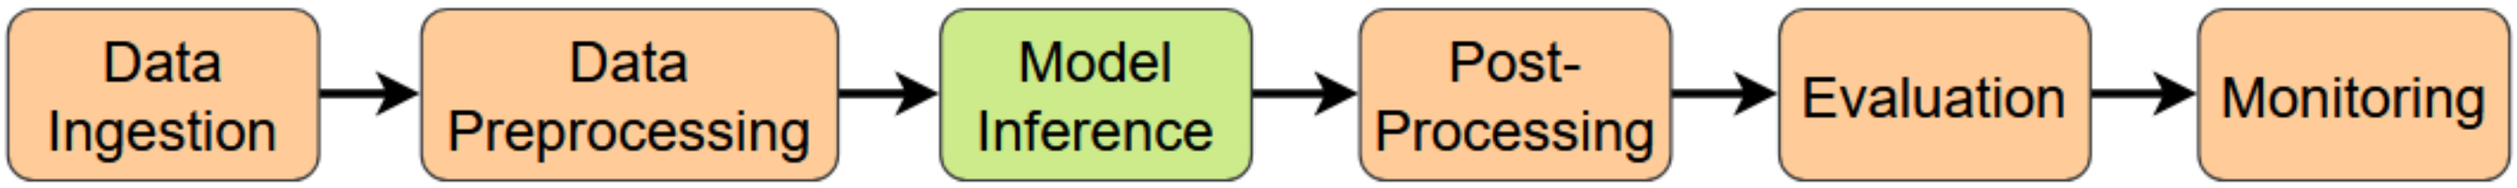
$$Required\ VRAM\ (GB) \approx \frac{parameters * bytes\ per\ parameter}{1\ billion}$$
$$\approx \#\ billion\ parameters * bytes\ per\ parameter$$
$$+\ overhead$$



Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- Hosting LLMs
  - 8B model at FP16: ~20GB VRAM
  - 400B model at FP16: ~850GB VRAM
  - 1T model at FP16: ~2TB VRAM



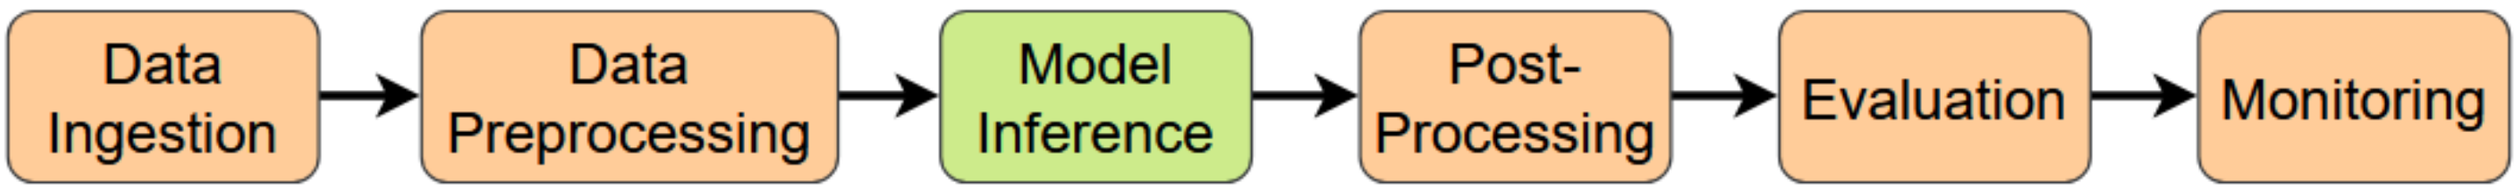
Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- Hosting LLMs
  - 8B model at FP16: ~20GB VRAM
  - 400B model at FP16: ~850GB VRAM
  - 1T model at FP16: ~2TB VRAM

GPU Type	VRAM (GB)	Cost ($)	Cloud Rental
A10	24	~2k	$0.50-2/hr
A100	40/80	10-17k	$1-4/hr
H100	80	27-30k	$2-7/hr
H200	141	31-32k	$4-10/hr

- NAIRR Resources can help get started



Components of an AI Workflow – Model Inference

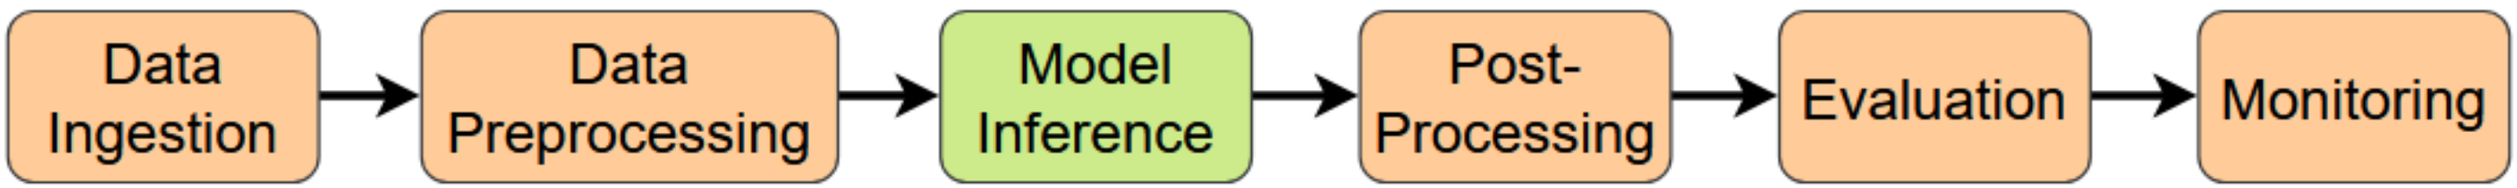
- What type of model are you running?
- What resources does that need?

- Hosting LLMs
  - 8B model at FP16: ~20GB VRAM
  - 400B model at FP16: ~850GB VRAM
  - 1T model at FP16: ~2TB VRAM

GPU Type	VRAM (GB)	Cost ($)	Cloud Rental
A10	24	~2k	$0.50-2/hr
A100	40/80	10-17k	$1-4/hr
H100	80	27-30k	$2-7/hr
H200	141	31-32k	$4-10/hr

- NAIRR Resources can help get started

1 GPU ≈ $1500/mo



Components of an AI Workflow – Model Inference

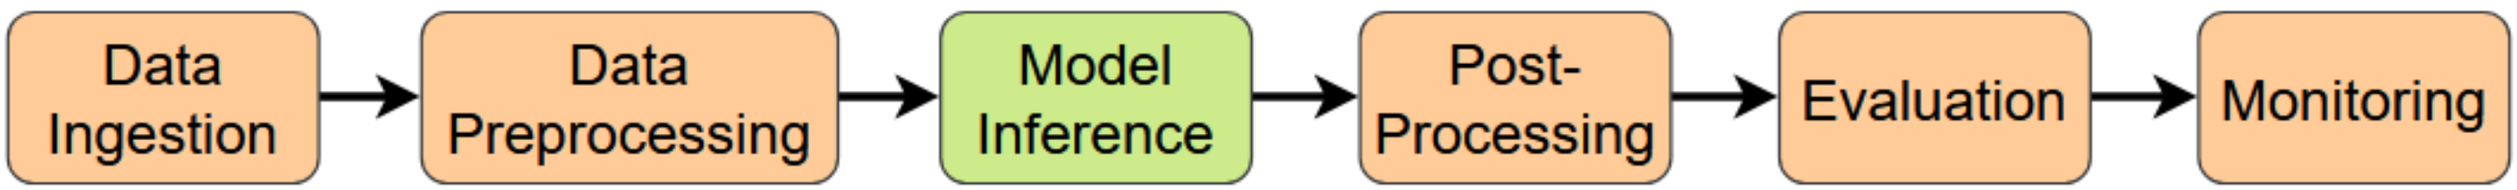
- What type of model are you running?
- What resources does that need?

- Hosting LLMs
  - 8B model at FP16: ~20GB VRAM
  - 400B model at FP16: ~850GB VRAM
  - 1T model at FP16: ~2TB VRAM

GPU Type	VRAM (GB)	Cost ($)	Cloud Rental
A10	24	~2k	$0.50-2/hr
A100	40/80	10-17k	$1-4/hr
H100	80	27-30k	$2-7/hr
H200	141	31-32k	$4-10/hr

8 GPU ≈ $24k/mo

- NAIRR Resources can help get started



Components of an AI Workflow – Model Inference

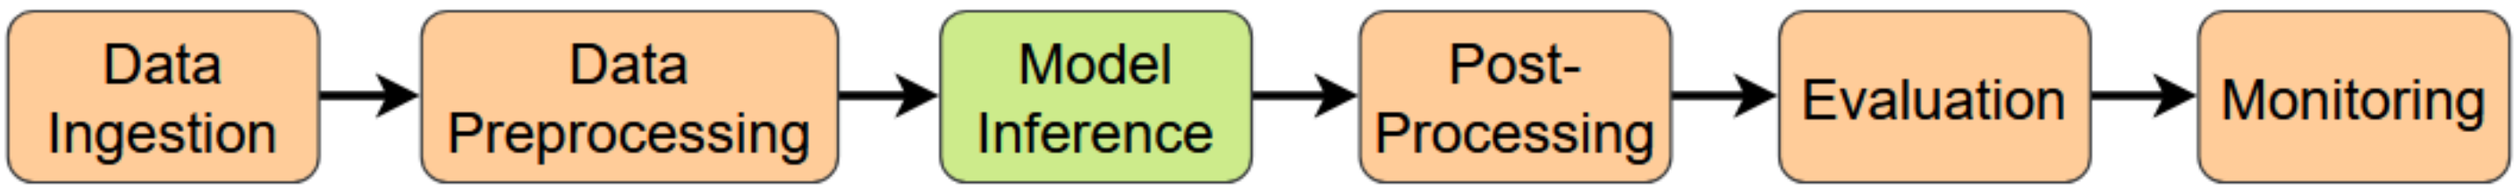
- What type of model are you running?
- What resources does that need?

- Hosting LLMs
  - 8B model at FP16: ~20GB VRAM
  - 400B model at FP16: ~850GB VRAM
  - 1T model at FP16: ~2TB VRAM

GPU Type	VRAM (GB)	Cost ($)	Cloud Rental
A10	24	~2k	$0.50-2/hr
A100	40/80	10-17k	$1-4/hr
H100	80	27-30k	$2-7/hr
H200	141	31-32k	$4-10/hr

- NAIRR Resources can help get started

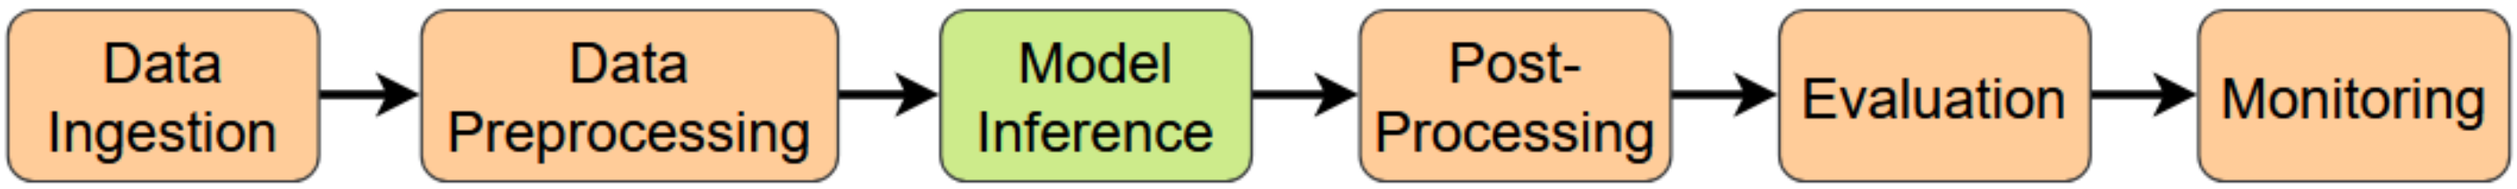
8 GPU ≈ $60k/mo



Components of an AI Workflow – Model Inference

- Model-as-a-Service (MaaS)
  - OpenAI API
  - Claude API
  - Gemini API

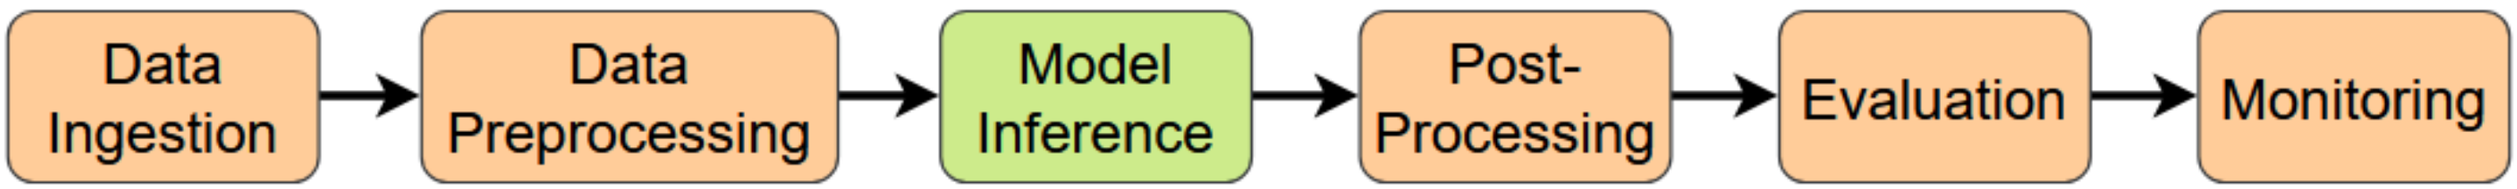
- Cost based on token usage
- Estimate token usage per month for your project
  - During evaluation step
  - NAIRR resources available



Components of an AI Workflow – Model Inference

BLUE

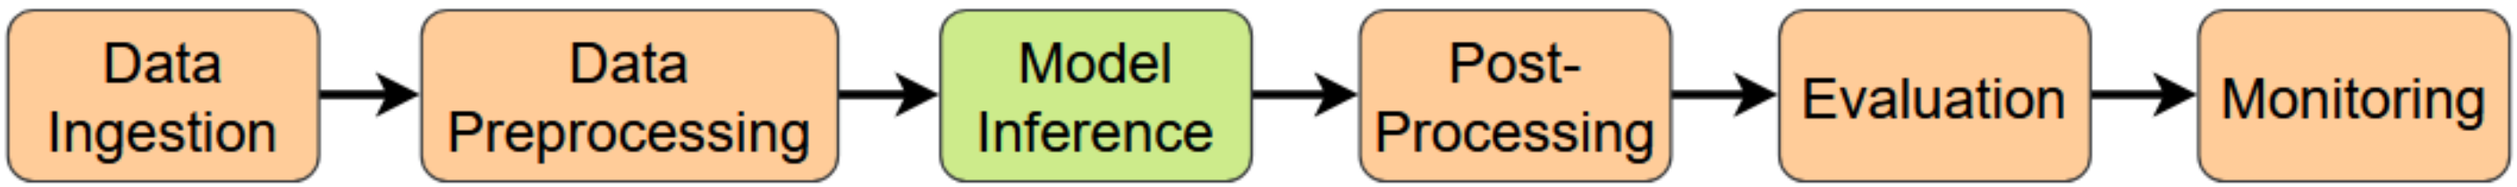
- How many tokens do we expect per month?
- What is our budget?
- Which setup will support that?



Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

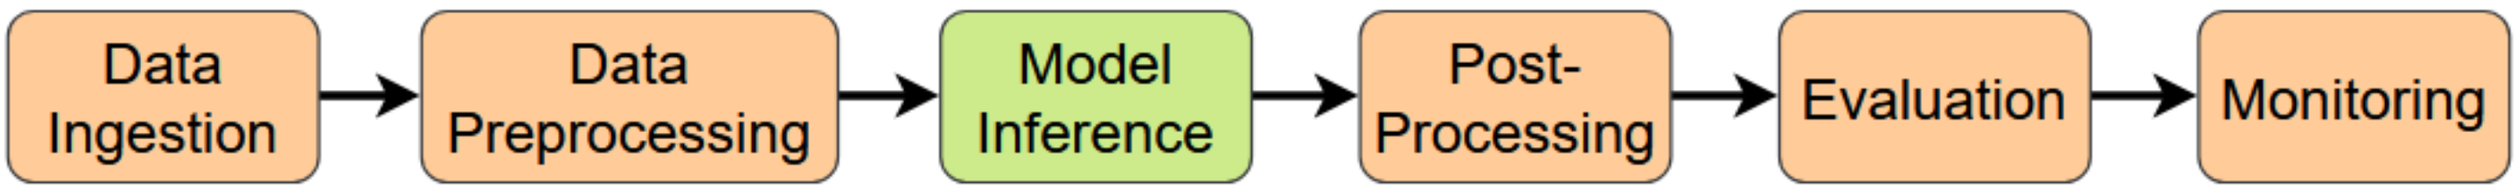
- Vision and traditional ML models more lightweight
  - Significantly less active memory requirements for inference
  - Many can run on only CPU



Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

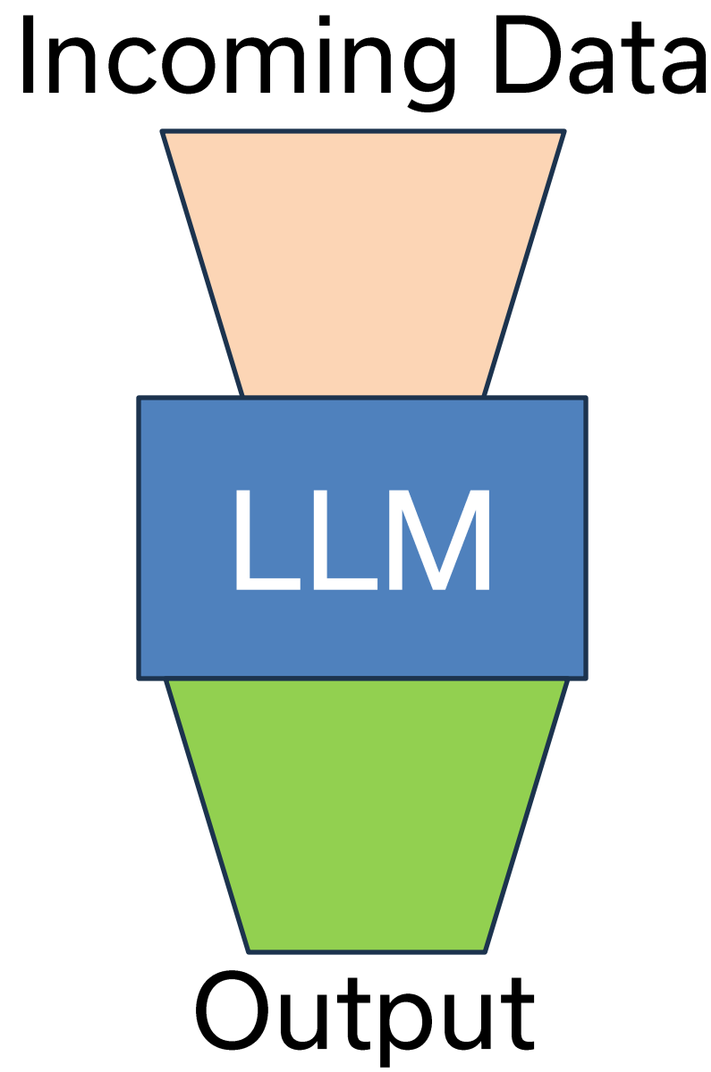
- How many users do you expect?
- Run multiple models in parallel?

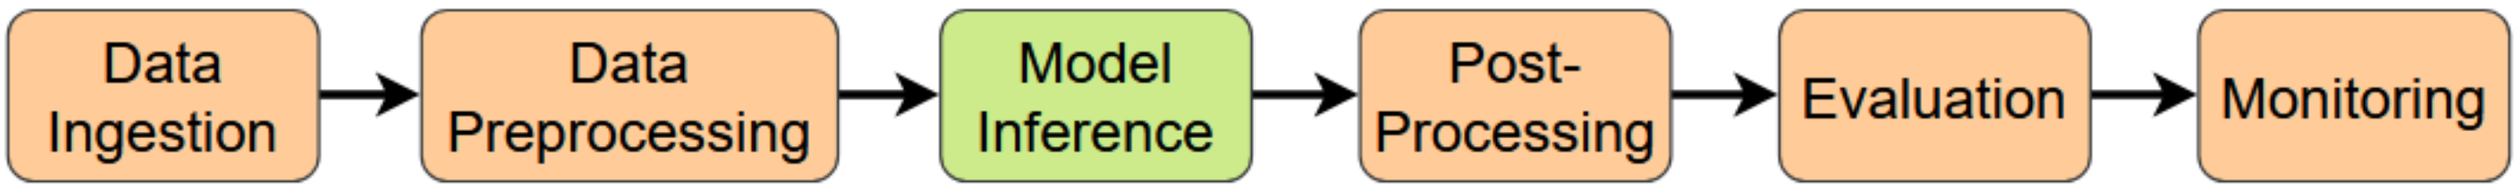


Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- How many users do you expect?
- Run multiple models in parallel?

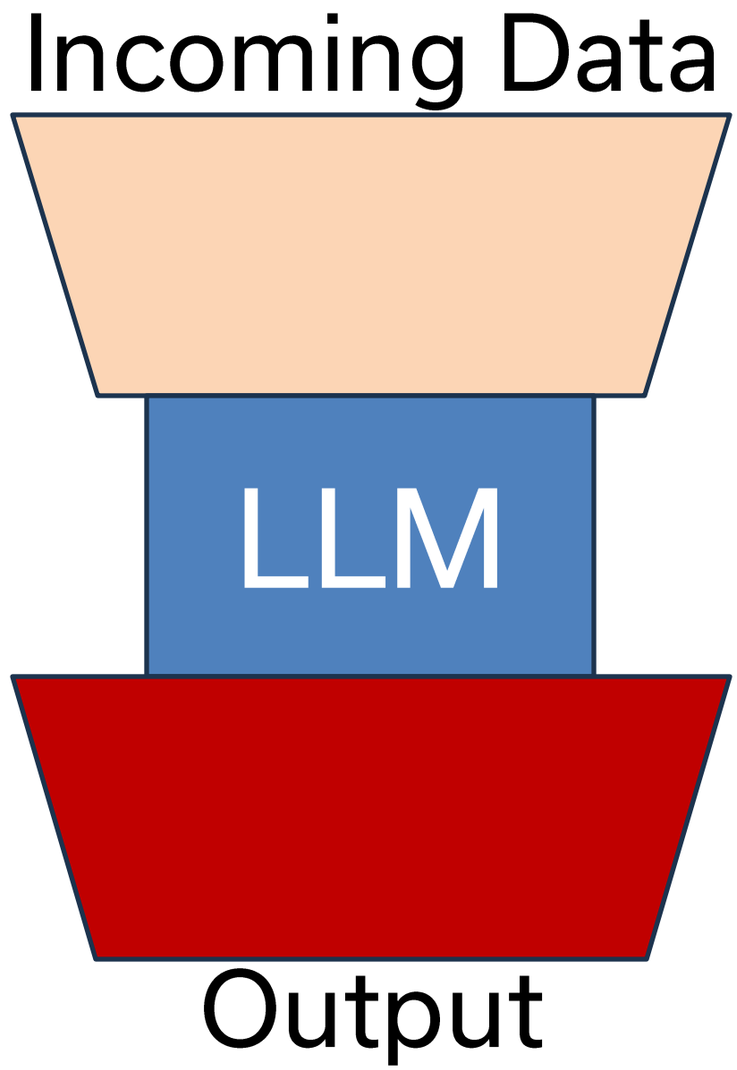


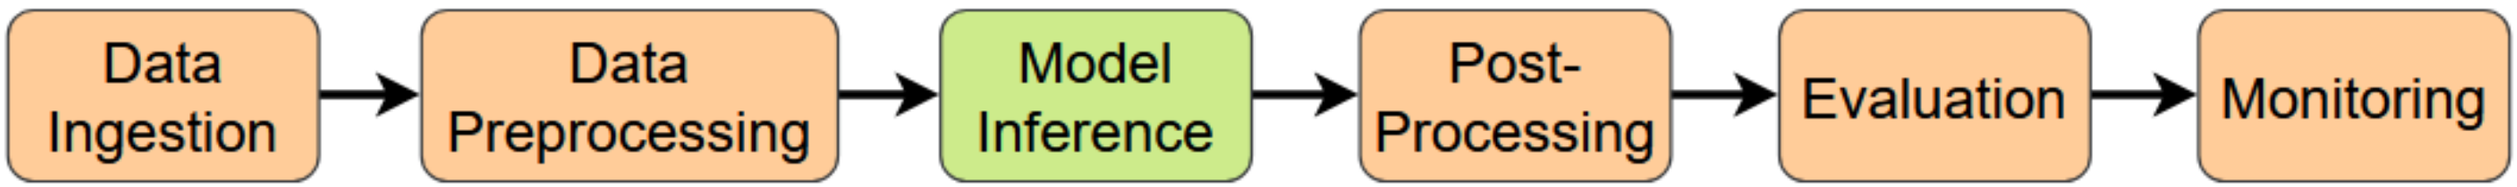


Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- How many users do you expect?
- Run multiple models in parallel?

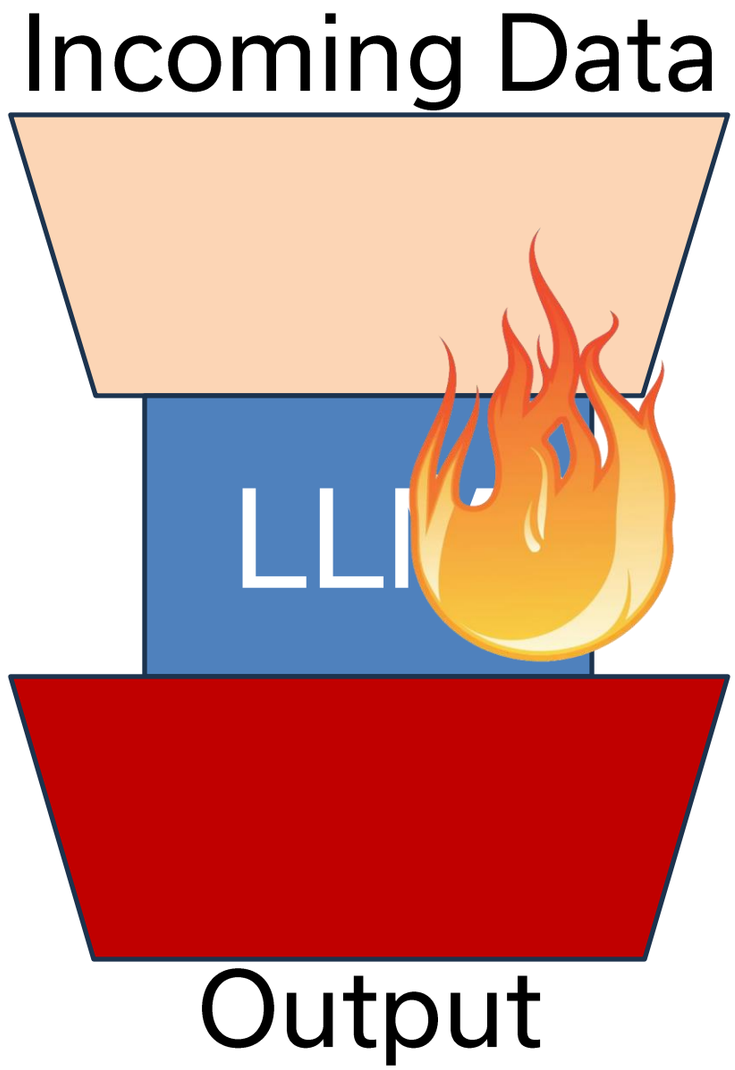


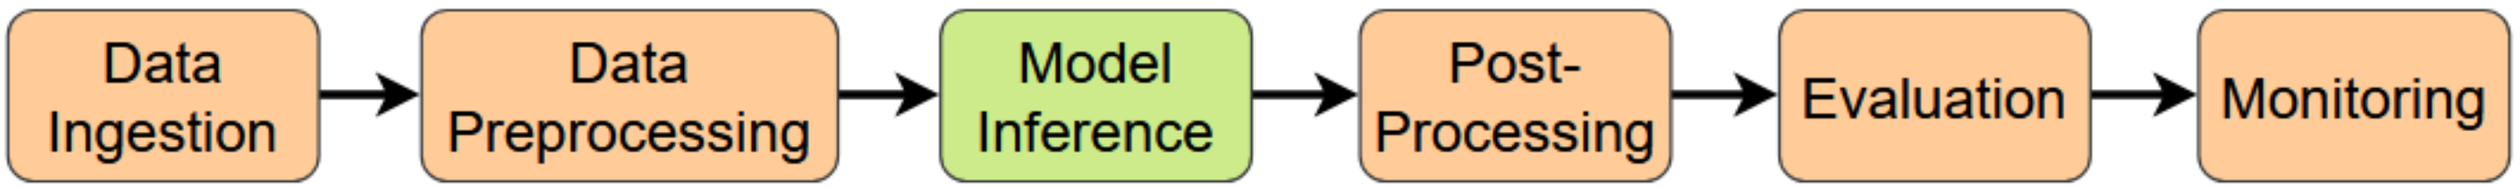


Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- How many users do you expect?
- Run multiple models in parallel?

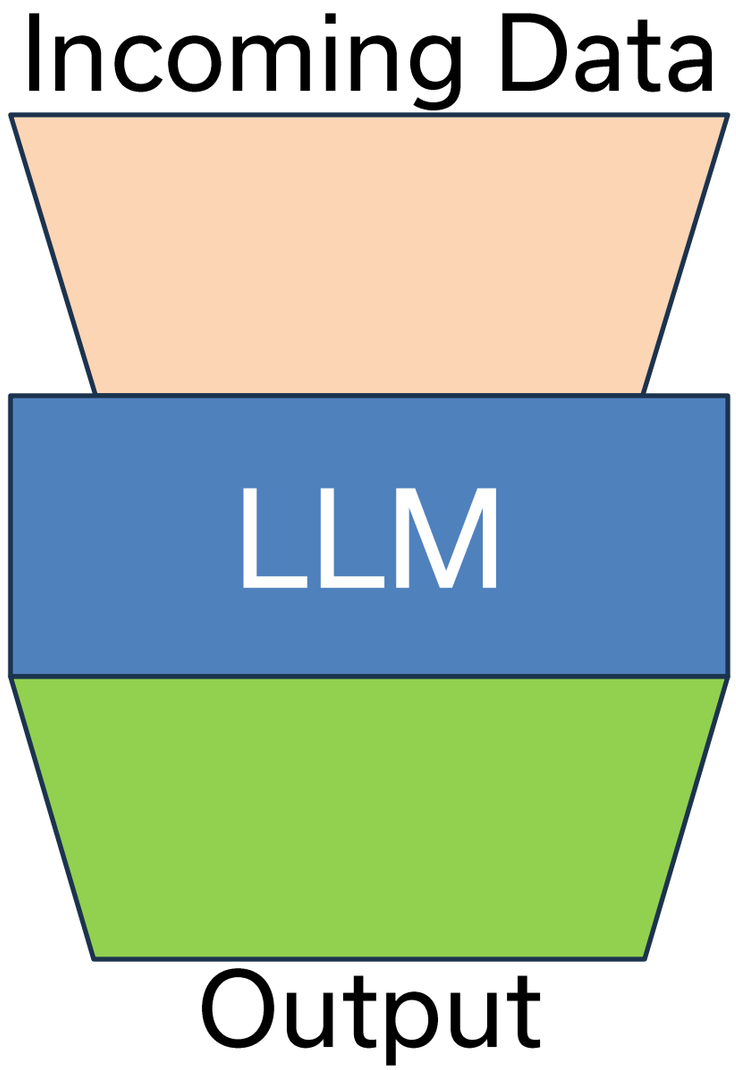


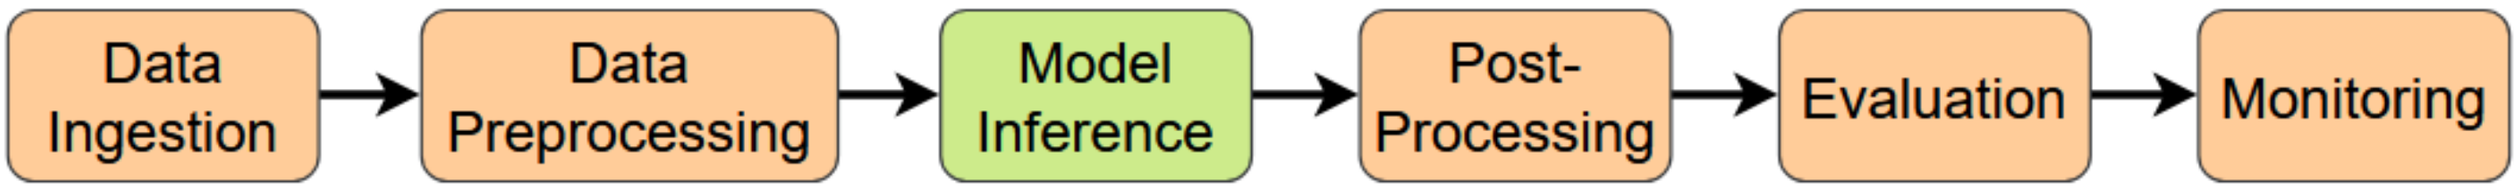


Components of an AI Workflow – Model Inference

- What type of model are you running?
- What resources does that need?

- How many users do you expect?
- Run multiple models in parallel?

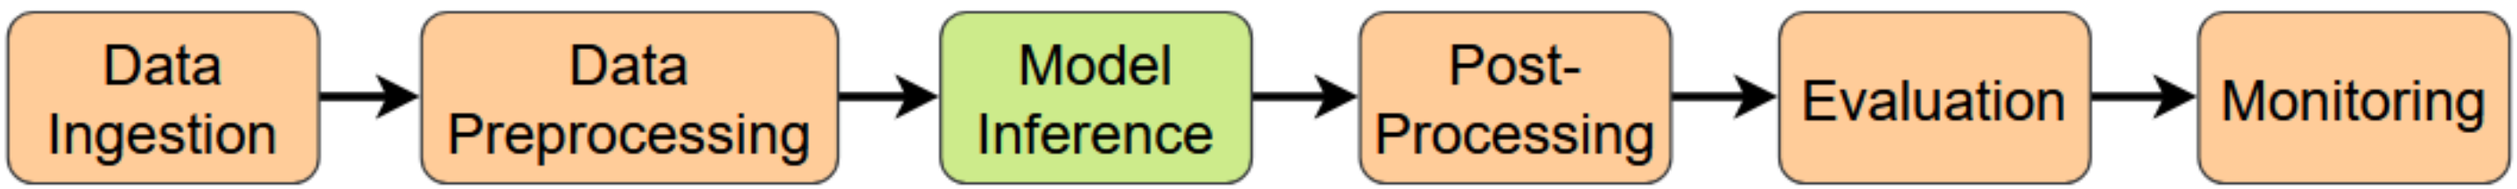




Components of an AI Workflow – Model Inference

BLUE

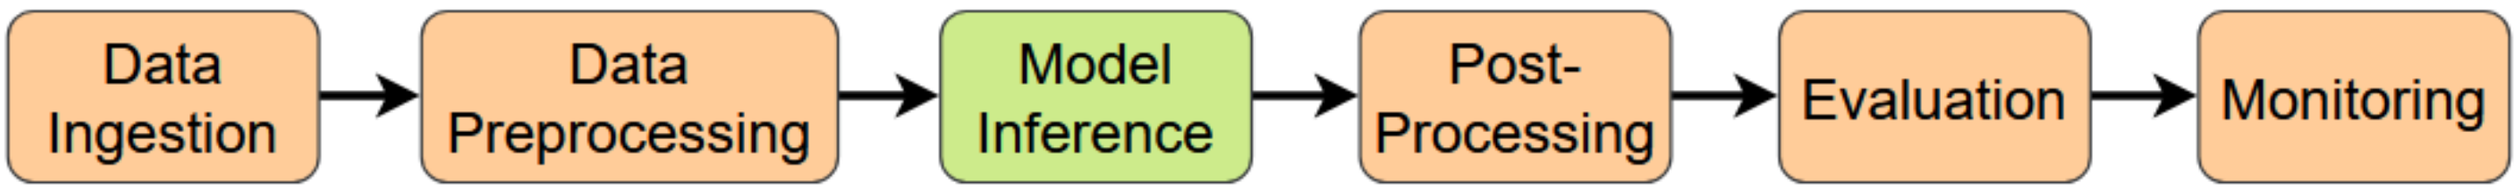
- Establish scaling methods for higher user load
- Gracefully handle crashes in parts of the system
- Design inference to be fast
- Lower latency allows more users



Components of an AI Workflow – Model Inference

Common Oversights or Problems

- Underestimating user load
- Not properly testing model
- Over-relying on the model
  - Especially important and difficult for LLMs
  - Inputs can be very broad, need measures for when model is wrong
    - Use RAG and other techniques

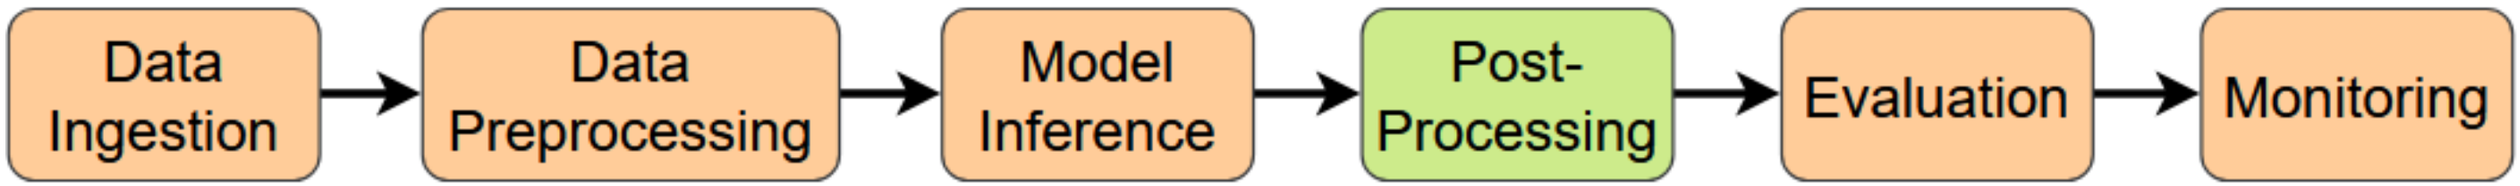


Components of an AI Workflow – Post-Processing

- How to interpret the output?
- How to present the output?

- If user-facing:
  - How to make output human-friendly?
  - How to report confidence?

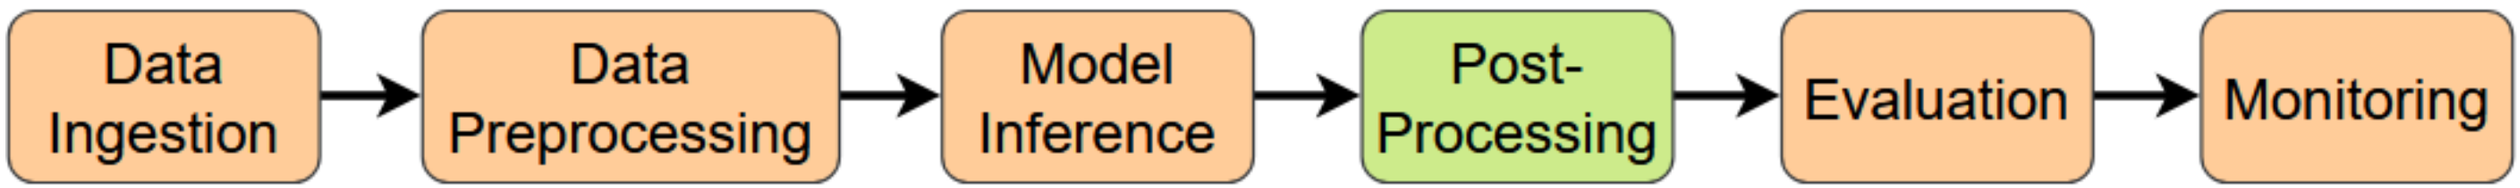
- How to limit network load with the output?



Components of an AI Workflow – Post-Processing

Common Oversights or Problems

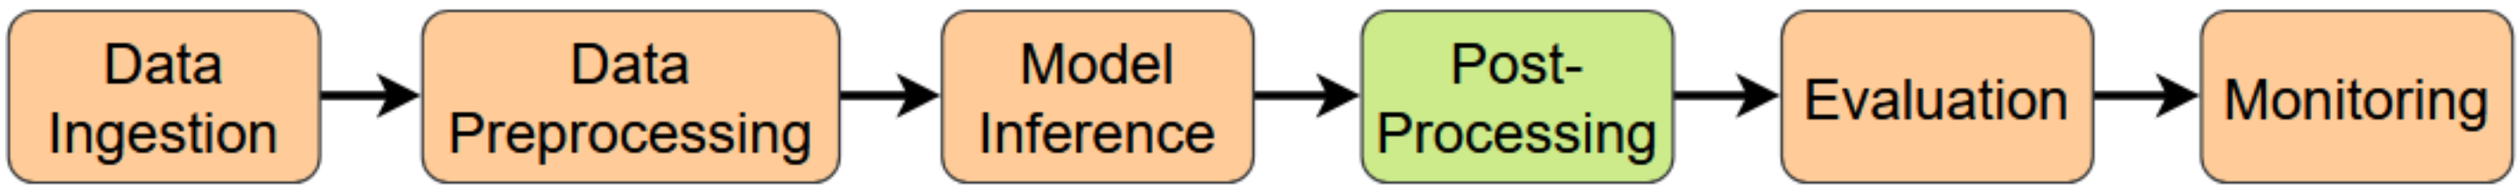
- Reporting raw model output to users is often not best-practice
- Not logging and storing useful data on model outputs



Components of an AI Workflow – Post-Processing

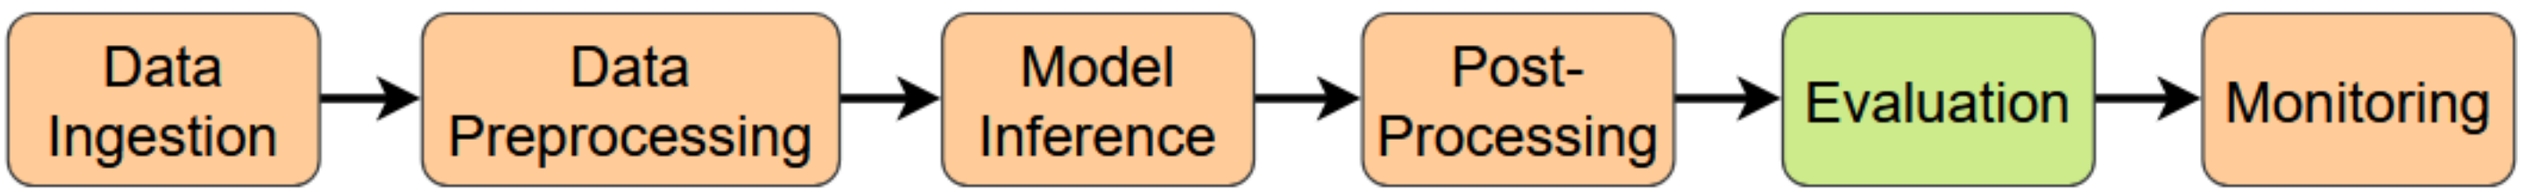
BLUE

- After RAG retrieval, provide information intelligently to LLM context
  - Allows model to respond better



Components of an AI Workflow – Evaluation

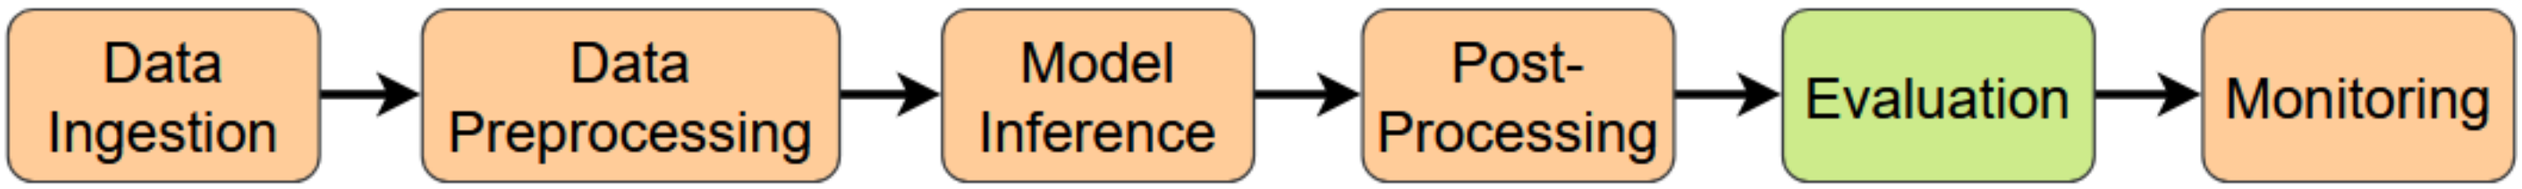
- Measuring correctness of model output
- Receiving human feedback
  - RL/HF
  - Make model better based on human review
- Handling incorrectness
- Identifying sources of incorrectness



Components of an AI Workflow – Evaluation

BLUE

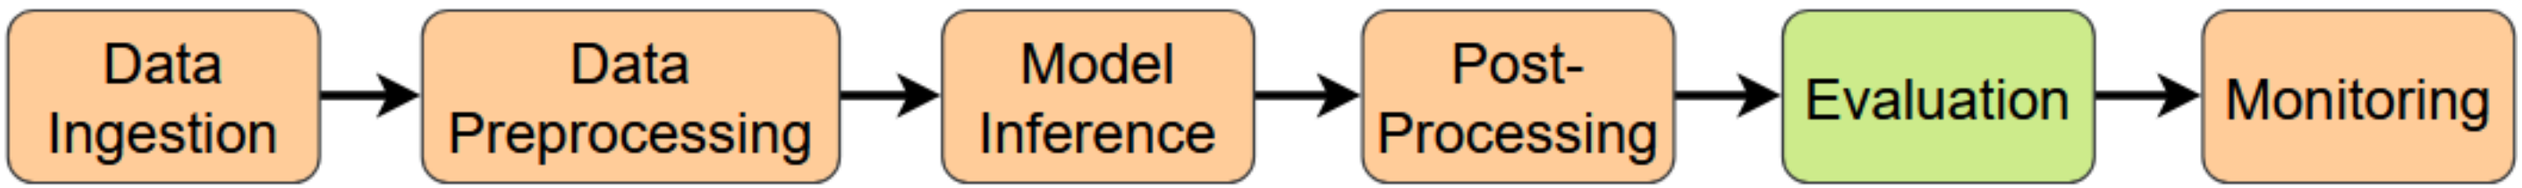
- Allow user to rate response
  - Thumbs up / Thumbs down
- Use feedback to identify weaknesses in inference



Components of an AI Workflow – Evaluation

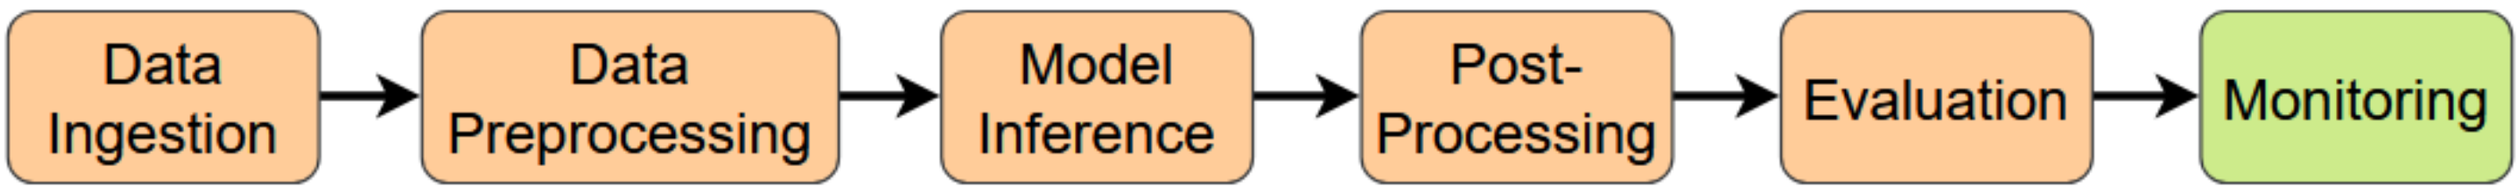
Common Oversights or Problems

- Continuous evaluation often overlooked
- Miss out on opportunities to improve model



Components of an AI Workflow – Monitoring

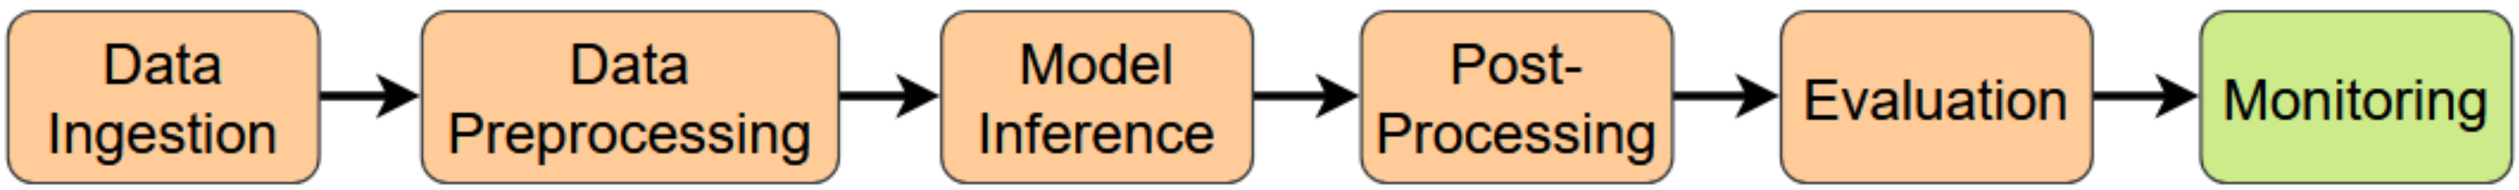
- Ensuring uptime
- Identifying any pipeline errors
- Identifying chokepoints



Components of an AI Workflow – Monitoring

BLUE

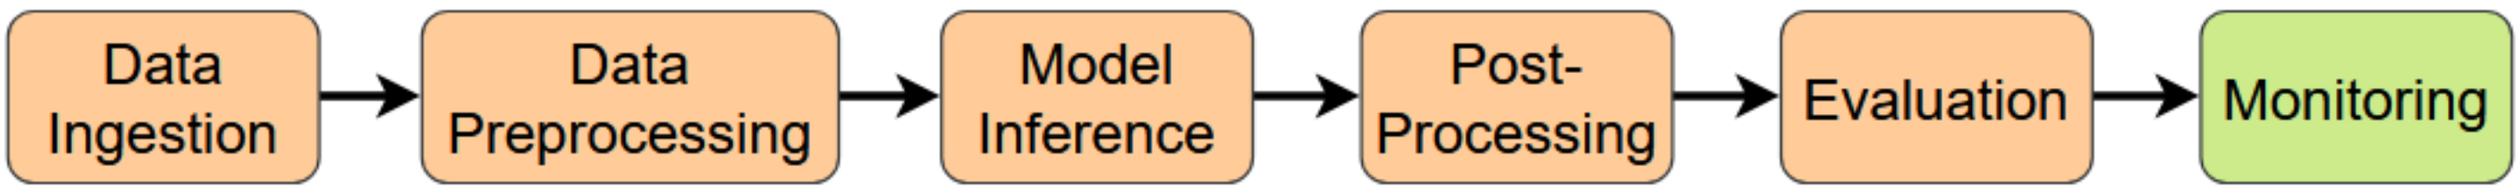
- Be notified of any system crashes
- Watch network traffic
- Identify security threats
- Etc.



Components of an AI Workflow – Monitoring

Common Oversights or Problems

- Ensure balance between each phase of the pipeline
  - Track what is useful
  - Don’t over-monitor
    - Can stall pieces
    - Can drastically increase storage requirements

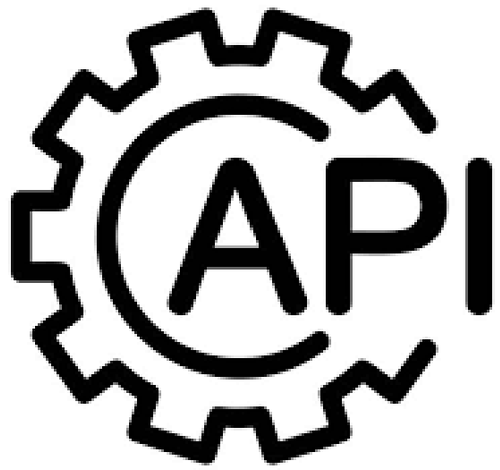


Deployment Strategies

Local Deployments vs. Cloud Deployments vs API Services







Deployment Strategies – Local Deployments

Host models on your hardware

Requirements:

- Server with GPU(s)
- Enough VRAM for your model
- Inference server software
  - vLLM
  - SGLang
- Management of servers





Deployment Strategies – Cloud Deployments

Host models on Cloud Hardware

- AWS
- Google Cloud Platform
- Microsoft Azure
- NIH Cloud Lab
- And many more

Available as NAIRR resources!

Deployment Strategies – Cloud Deployments

https://submit-nairr.xras.org/resources

Resource Catalog

Filters

AI Capabilities

- [ ] AI tools and support
- [x] Model inference services
- [ ] Model training services (GPU)
- [ ] Model training services (non-GPU)
- [ ] Research Collaboration

Resource Type

- [ ] CPU Compute
- [x] Cloud
- [x] Commercial Cloud
- [ ] GPU Compute
- [ ] Innovative / Novel Compute
- [ ] Sensors / Instruments
- [ ] Service / Other

Resources

- Amazon Web Services
- CloudBank Classroom
- CloudBank Research
- Google Cloud Platform
- Microsoft Azure
- NIH Cloud Lab
- SambaCloud (API inference services, Llama, DeepSeek, Qwen, etc.)

Deployment Strategies – API Services

Access models through existing APIs

- OpenAI API (ChatGPT)
- Anthropic API (Claude)
- Google API (Gemini)

Available as NAIRR resources!

Deployment Strategies – API Services

https://submit-nairr.xras.org/resources

Resource Catalog

Filters

AI Capabilities

- [ ] AI tools and support
- [x] Model inference services
- [ ] Model training services (GPU)
- [ ] Model training services (non-GPU)
- [ ] Research Collaboration

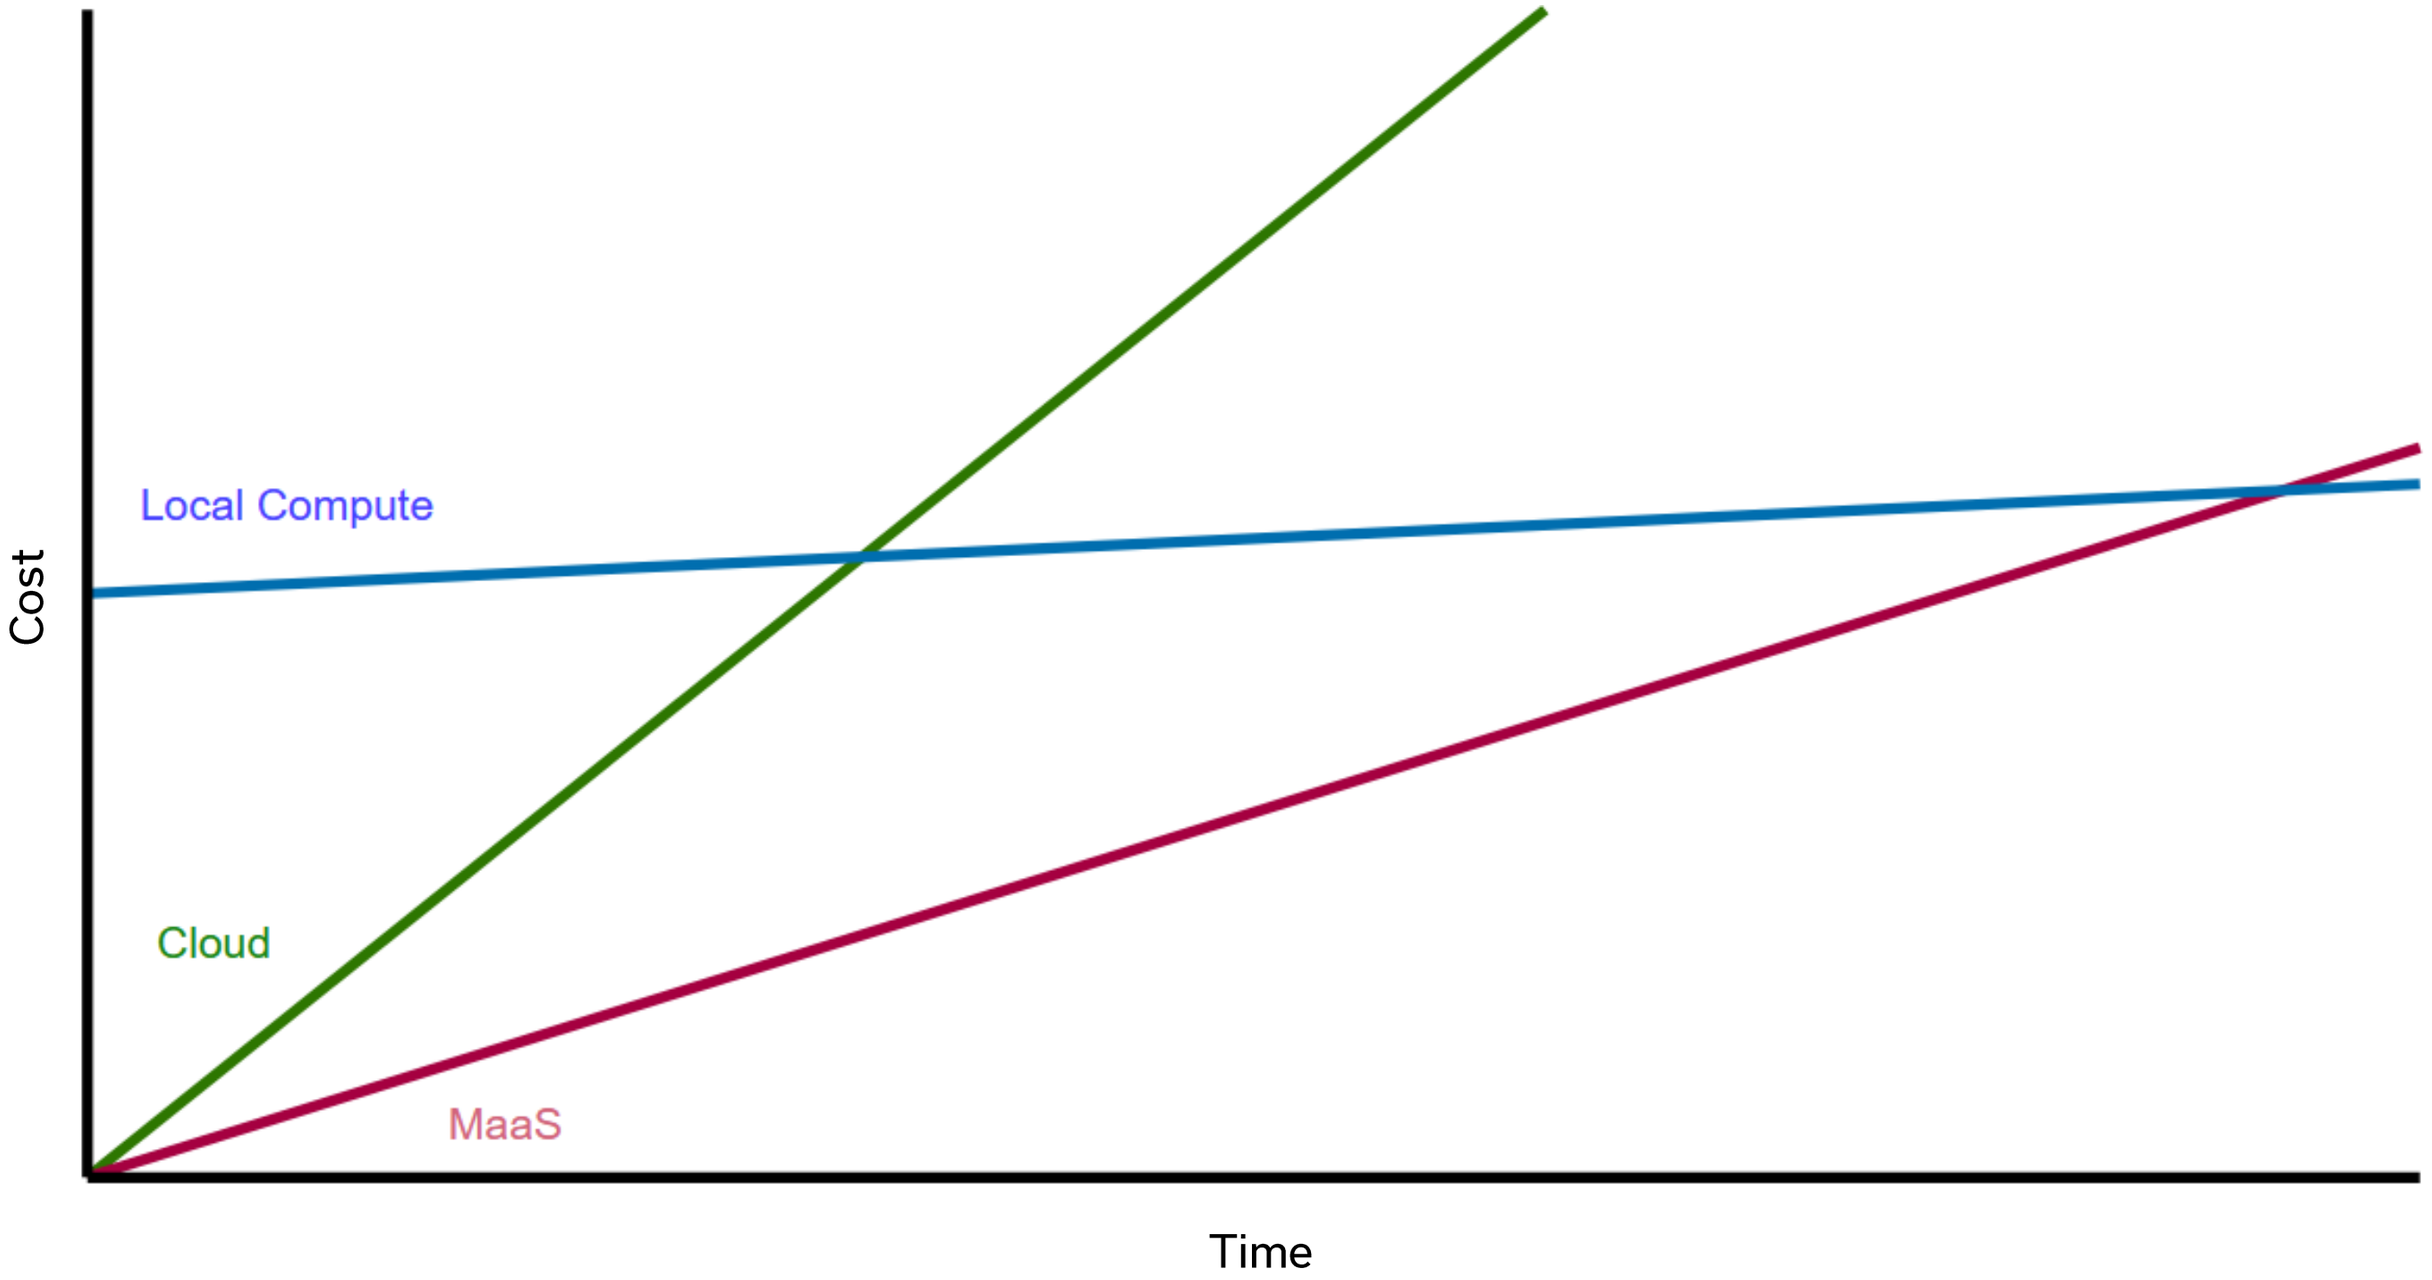
Resource Type

- [ ] CPU Compute
- [ ] Cloud
- [ ] Commercial Cloud
- [ ] GPU Compute
- [ ] Innovative / Novel Compute
- [ ] Sensors / Instruments
- [x] Service / Other

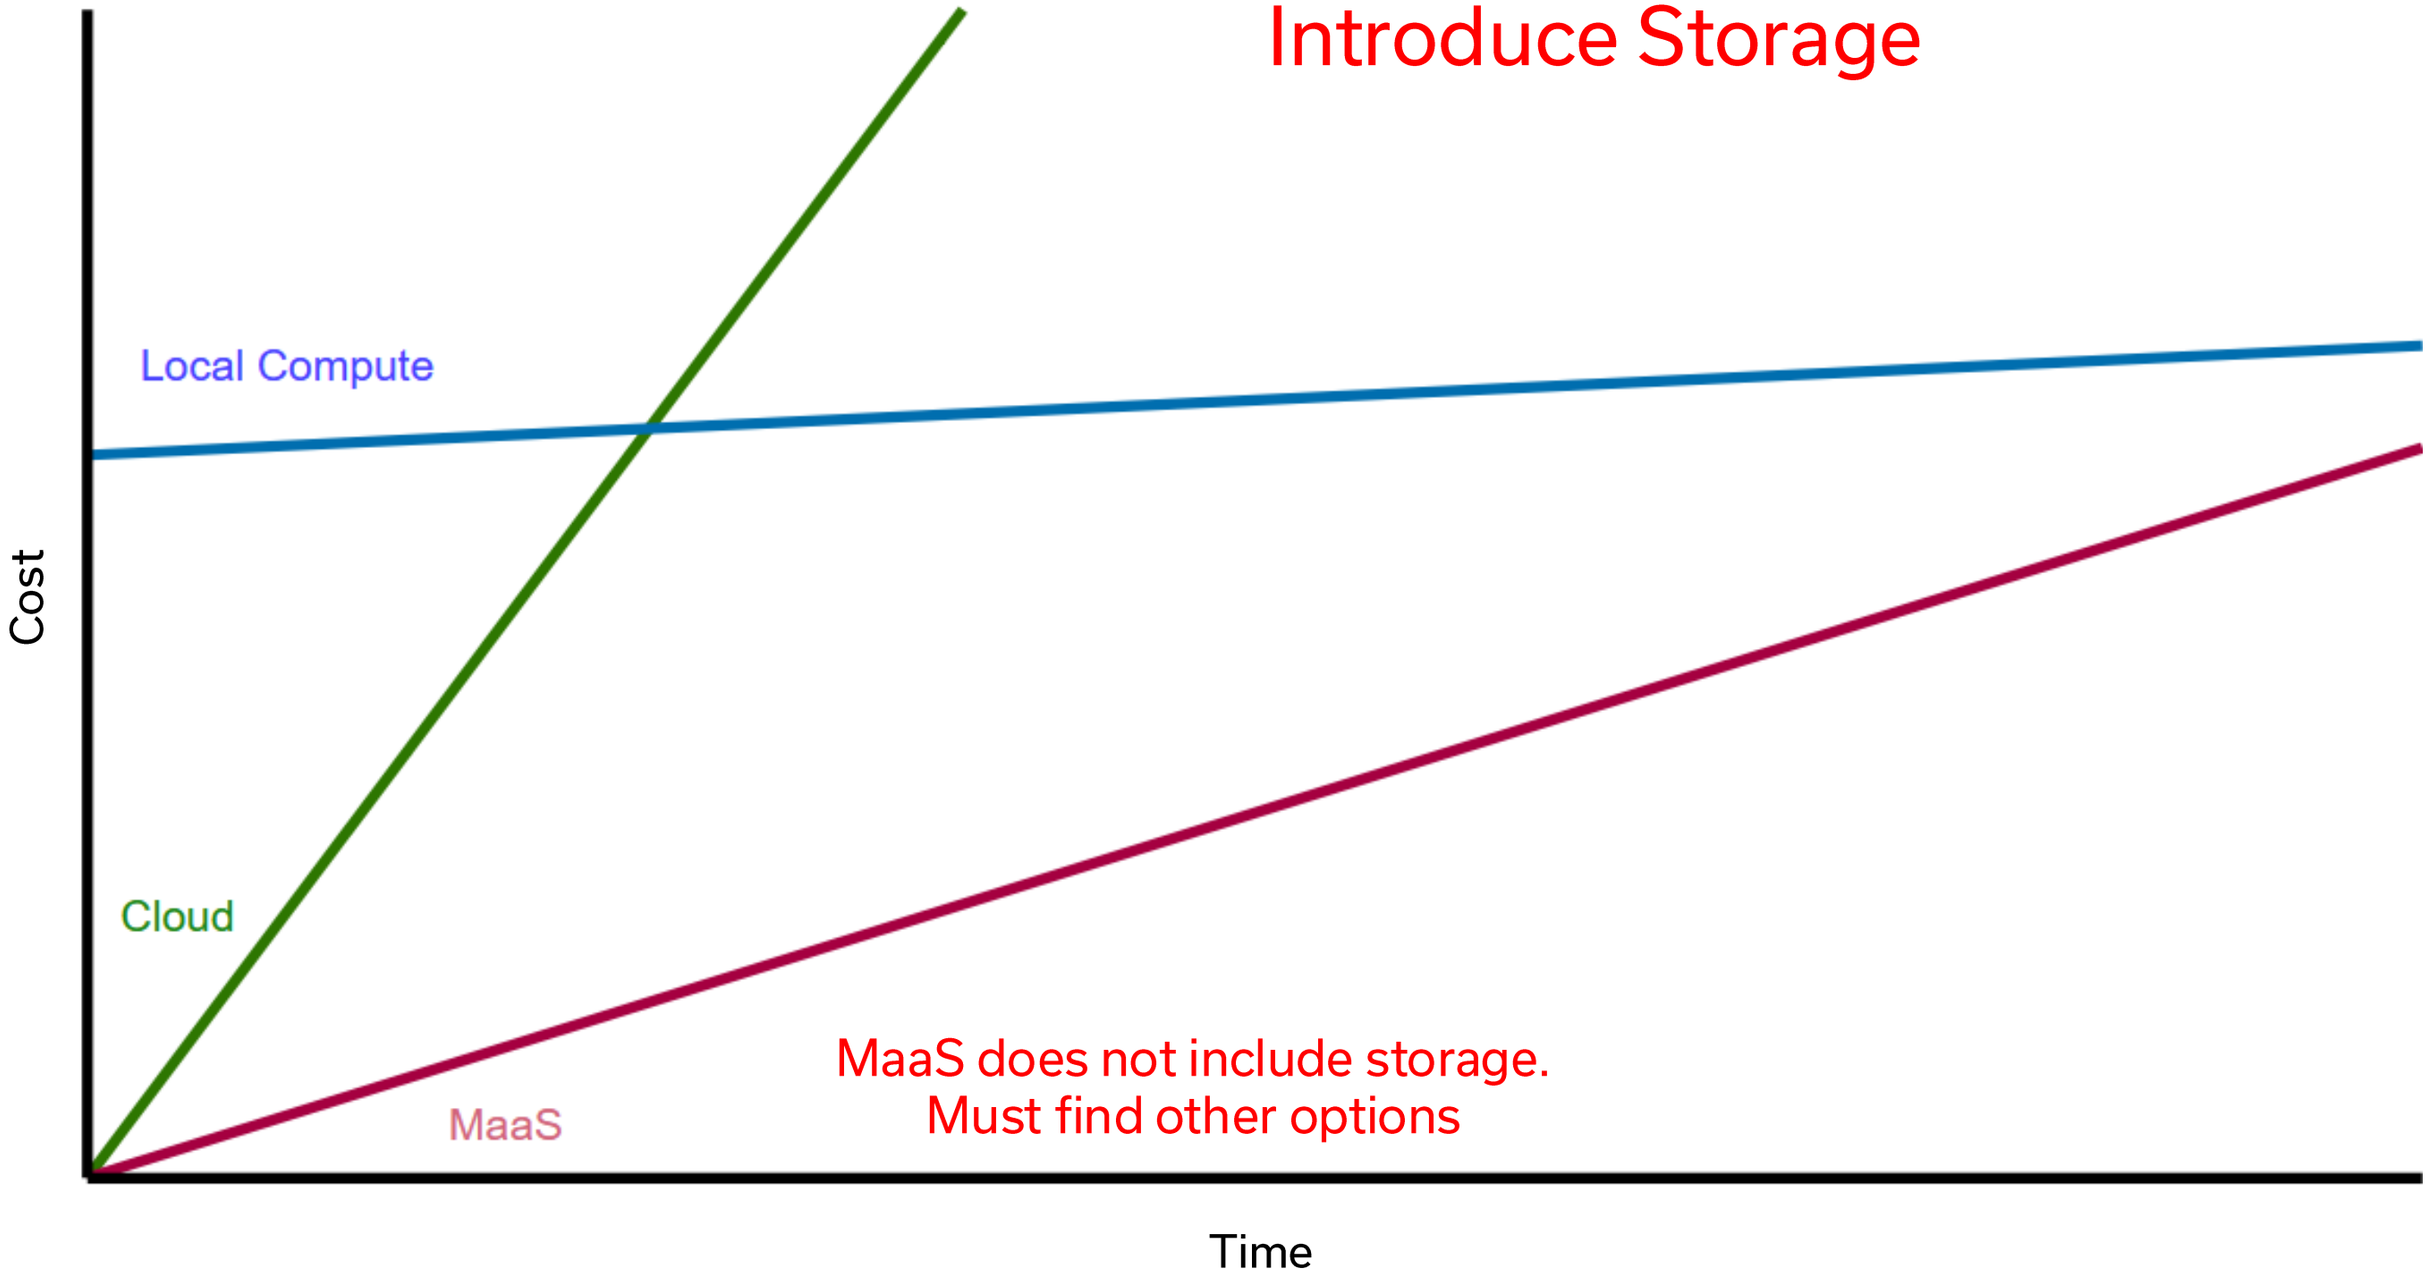
Resources

- Anthropic Model API
- OpenAI API
- SambaCloud (API inference services, Llama, DeepSeek, Qwen, etc.)

Deployment Strategies – Cost



Deployment Strategies – Cost



Applied AI Panel

CAAI Team

Hazard Community & Technical College
Hazard, Kentucky
March 26th, 2026

NAIRR Pilot

Panel: Share Your Work



- How has AI impacted your research? (Good and Bad!)
- What bold research question becomes possible because of AI?
- Have you incorporated AI solutions anywhere in your work?

THANK YOU

HAZARD COMMUNITY AND
TECHNICAL COLLEGE

NAIRR Pilot

NSF

ICPE
Higher Education Matters

HAZARD
COMMUNITY & TECHNICAL COLLEGE

University of Kentucky

Don’t Forget!

- Resources are available to you after this event
    - Office Hours
    - Online Forum
    - YouTube
    - Demos on GitHub
- NAIRR Pilot allocation requests are **open**!

1. nairrpilot.org/
2. caai.ai.uky.edu/about-us/news-and-events/nairr-pilot-expansion-project/

Post Event Survey

- Let us know how we did today
- Help us in our reporting
- Your responses will be used to shape future events

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

CPE Higher Education Matters

University of Kentucky

Thank you for attending!

Scan to access the webpage:

Scan to access post-event survey:

ai_workshops@uky.edu

NAIRR Pilot

KENTUCKY COMMUNITY & TECHNICAL COLLEGE SYSTEM

ICPE Higher Education Matters

University of Kentucky