الجامعة السورية الخاصة
SYRIAN PRIVATE UNIVERSITY

AI for Interdisciplinary Research

Part 4: Modern AI and Research Directions

From Embeddings to Transformers and Beyond

Dr. Riad Sonbol

April 2026

Road Map

- AI as a Black Box – Concepts and Foundations
- Opening the Box (1) – Traditional Machine Learning
- Opening the Box (2) – Deep Learning
- Modern AI and Research Directions

The Core Problem in AI

"How can computers understand language, images, and meaning?"

- Computers only understand numbers.
- So AI needs a way to convert:
  - words
  - images
  - sounds

into numerical representations.

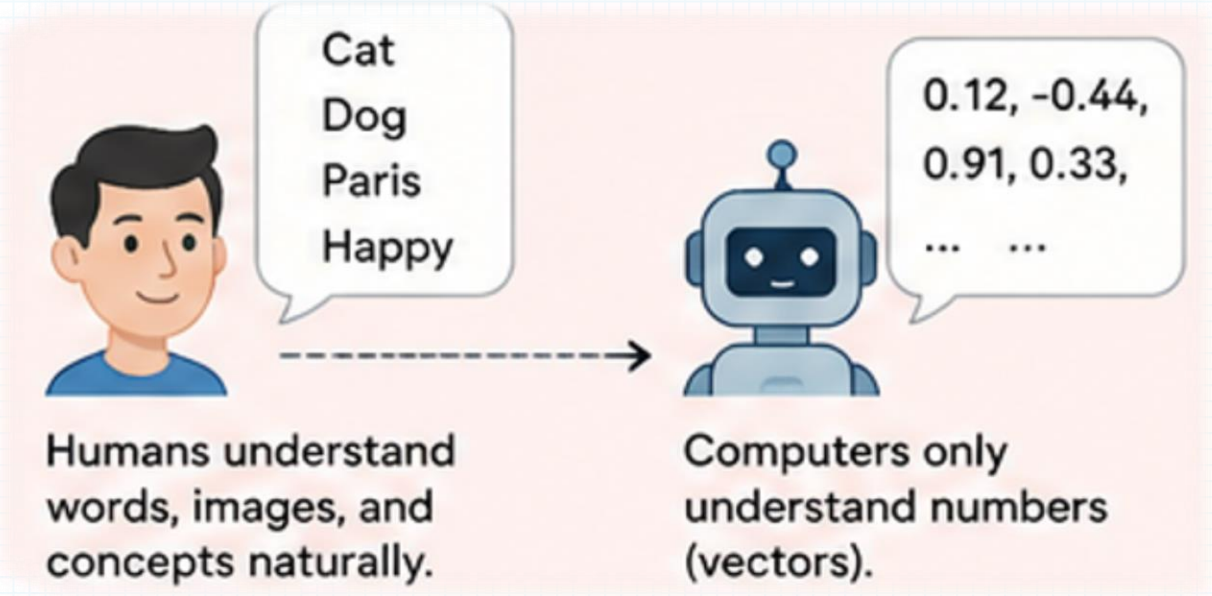

Embeddings: The Foundation of Modern AI

- Embeddings are numerical representations that capture meaning.

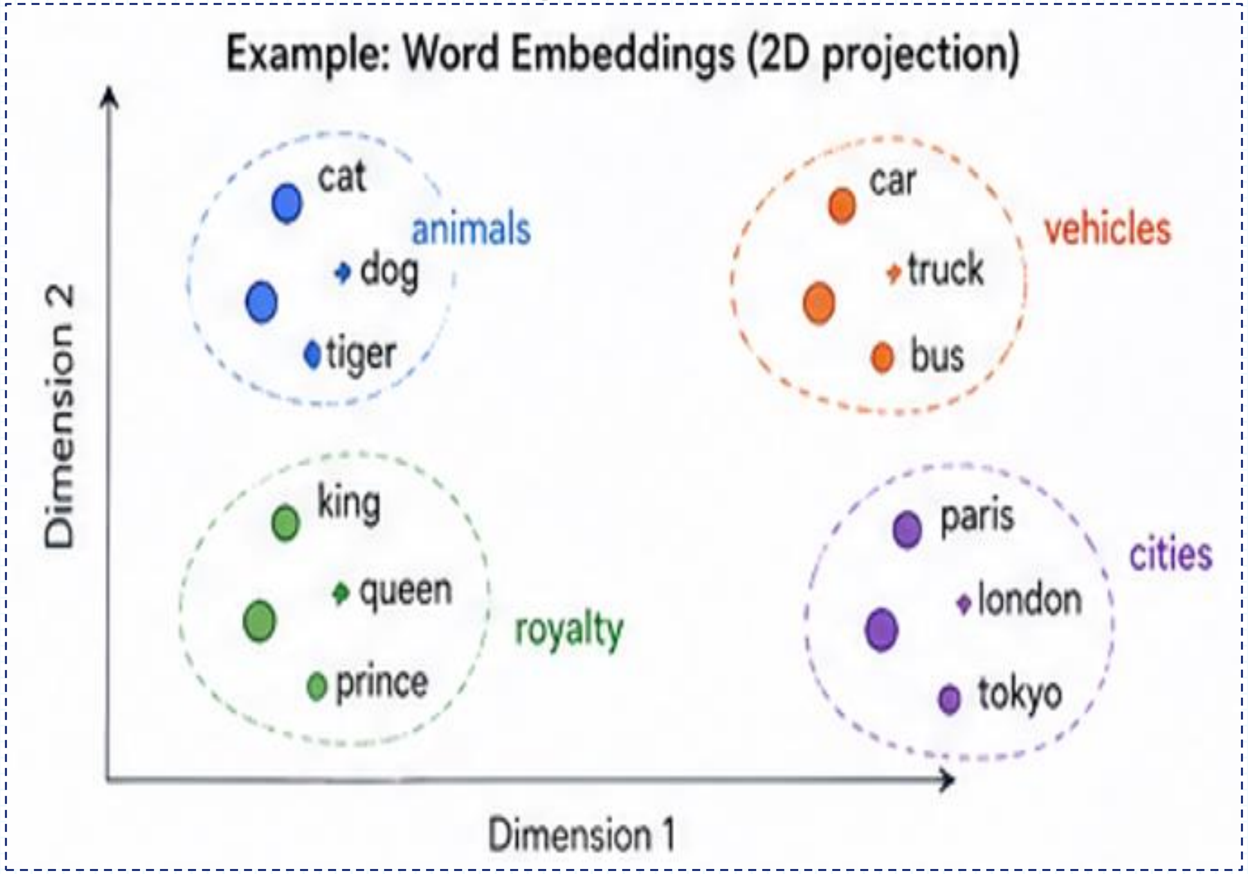

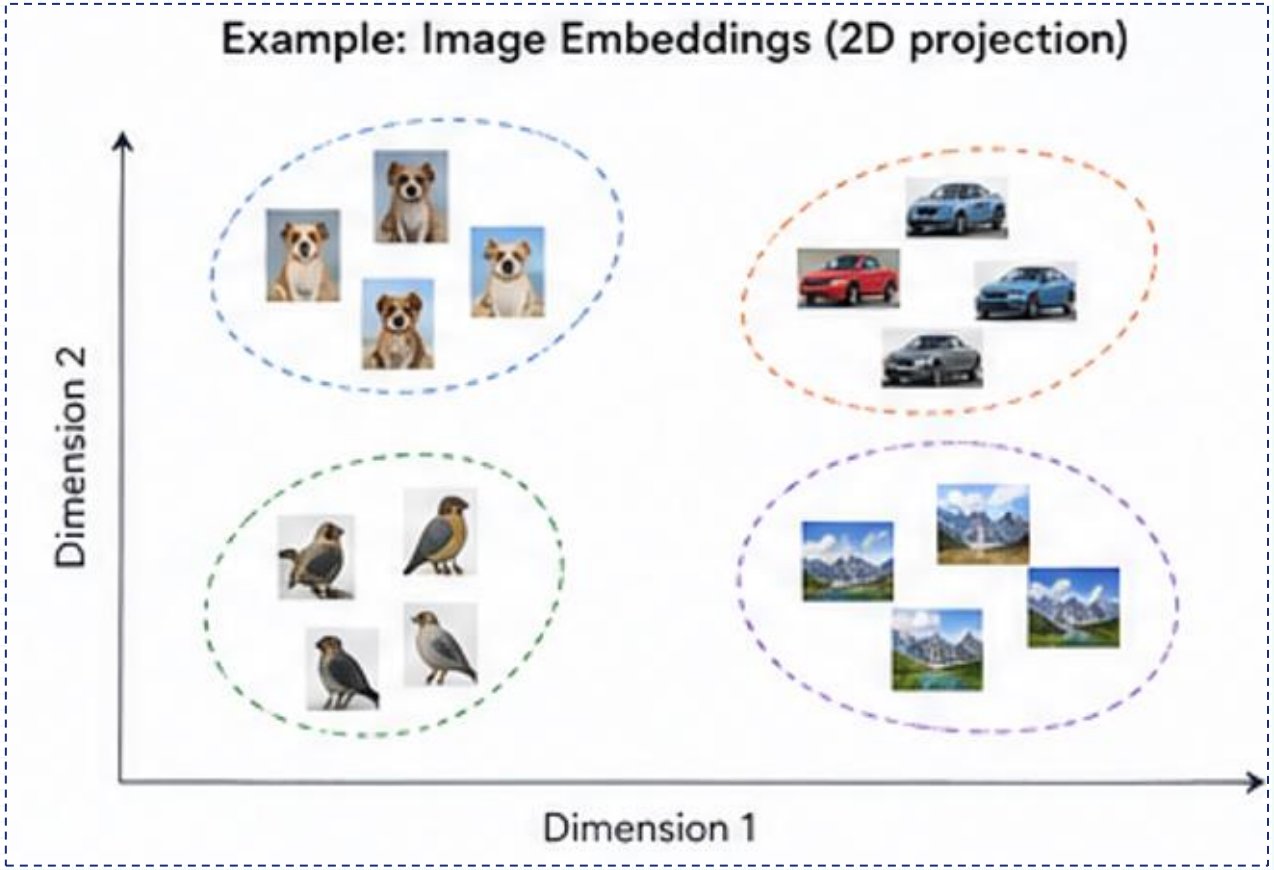

- Words with similar meanings are represented by nearby vectors.

The “Good” Embedding?

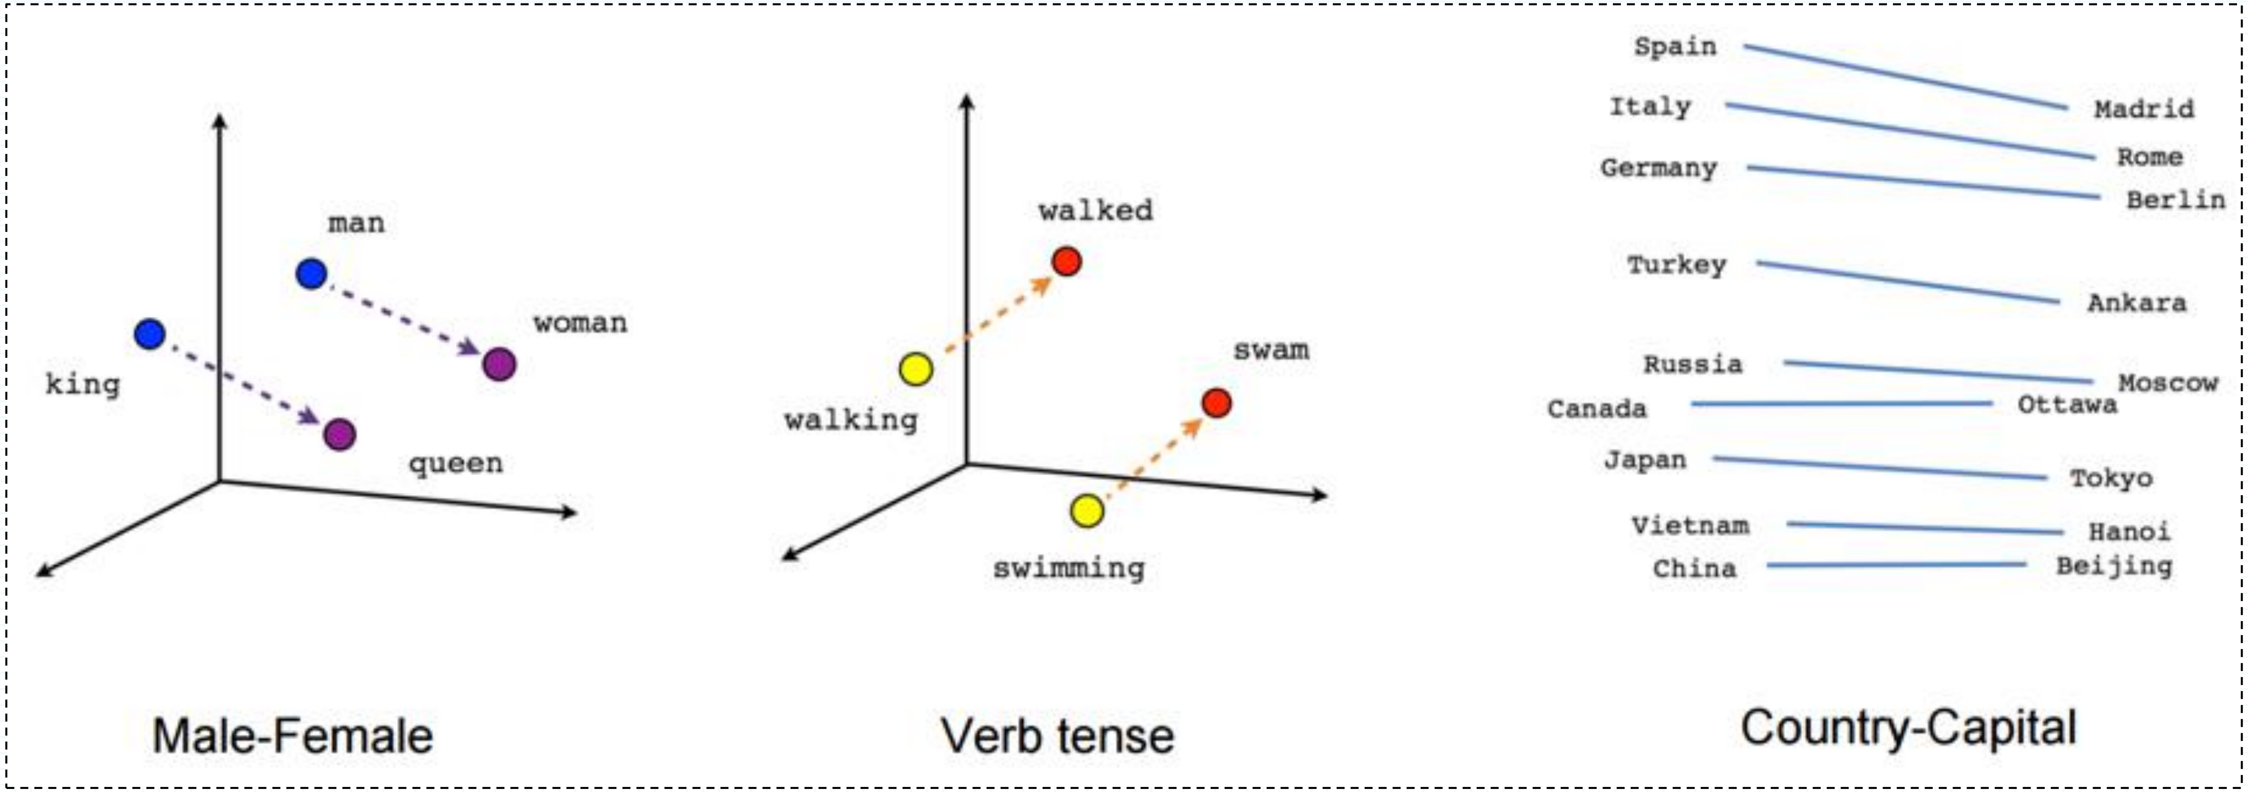

- vector[Queen] ≈ vector[King] - vector[Man] + vector[Woman]
- vector[Paris] ≈ vector[France] - vector[Italy] + vector[Rome]
  - This can be interpreted as “France is to Paris as Italy is to Rome”.

Embeddings are used for many types of data

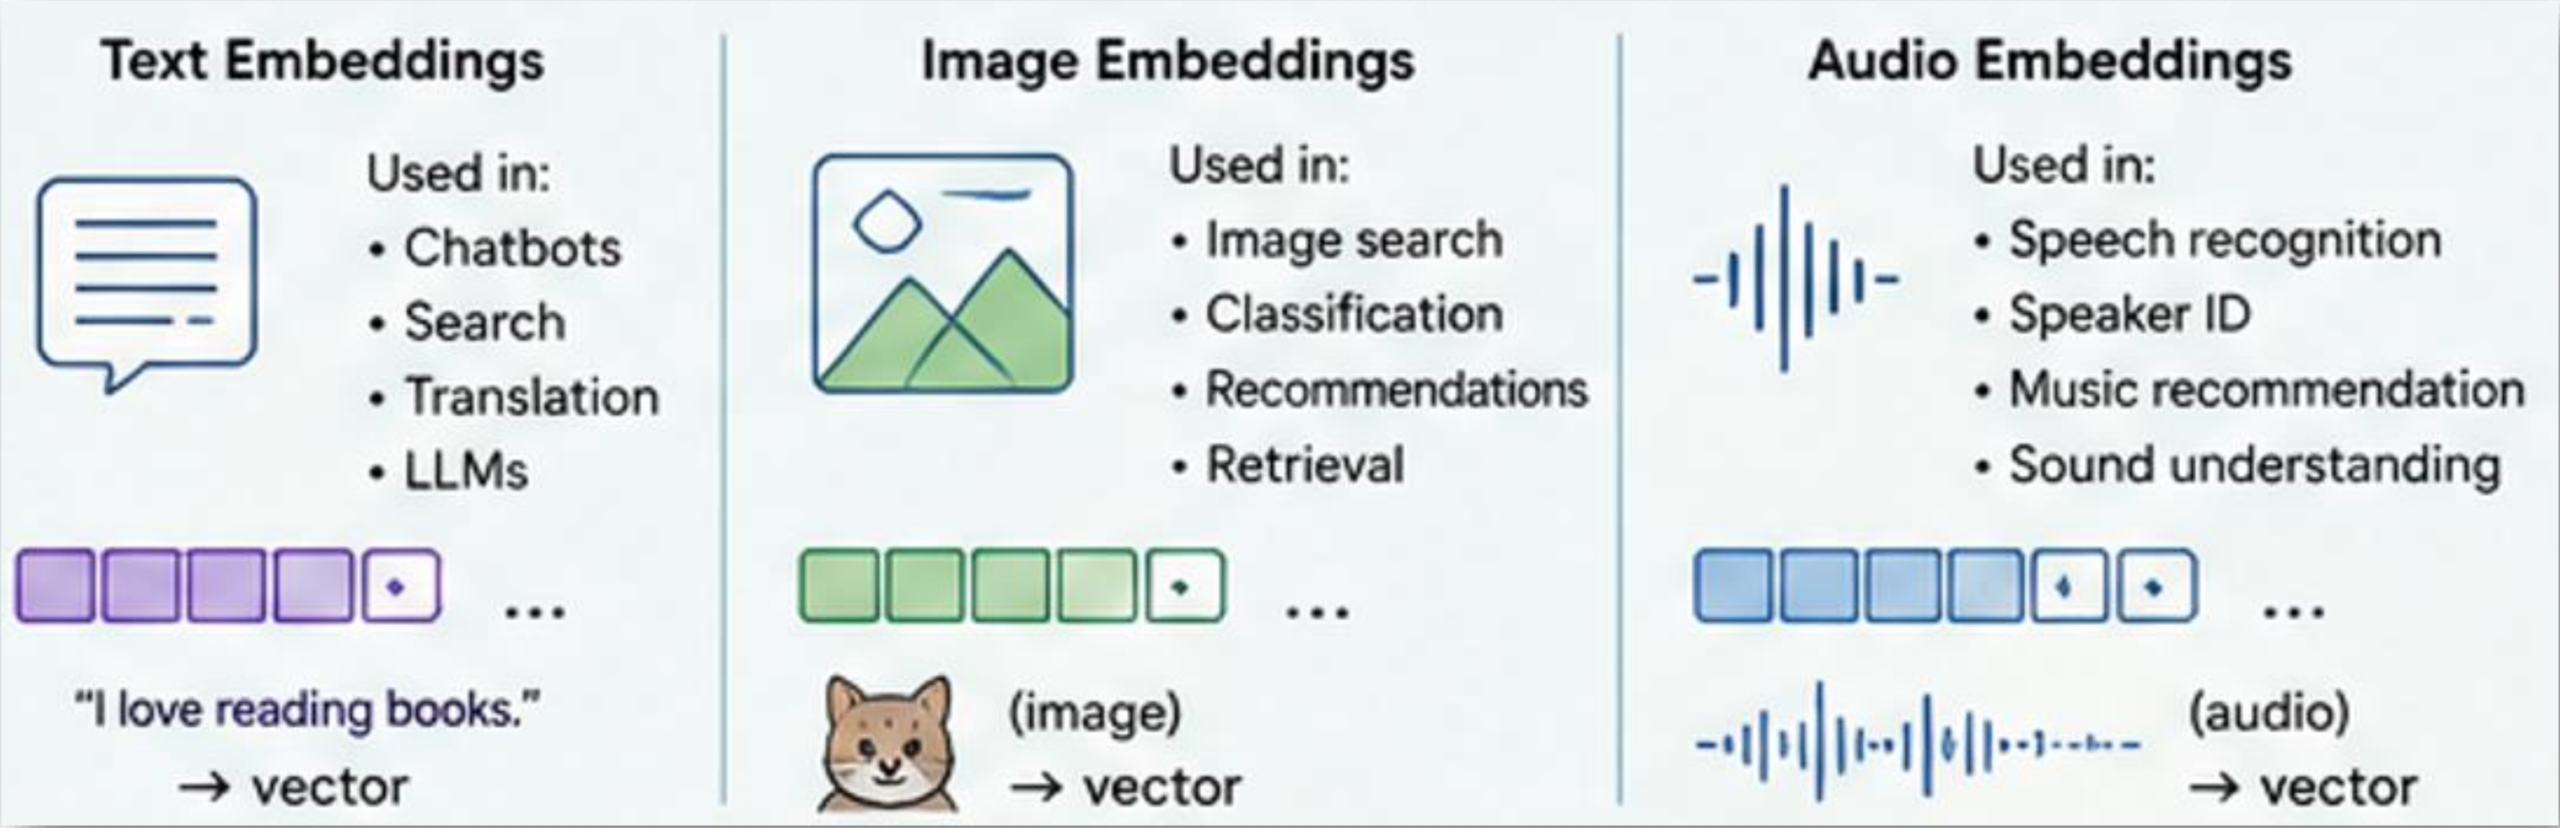

Embeddings aren't just for words/images!

In **biology**, an embedding could represent a **protein sequence**.

in **finance** , an embedding could represent a **Customer** or **Transaction**.

In **Civil Engineering**, an embedding could represent a **structural behavior**.

etc

Embeddings aren't just for words/images!

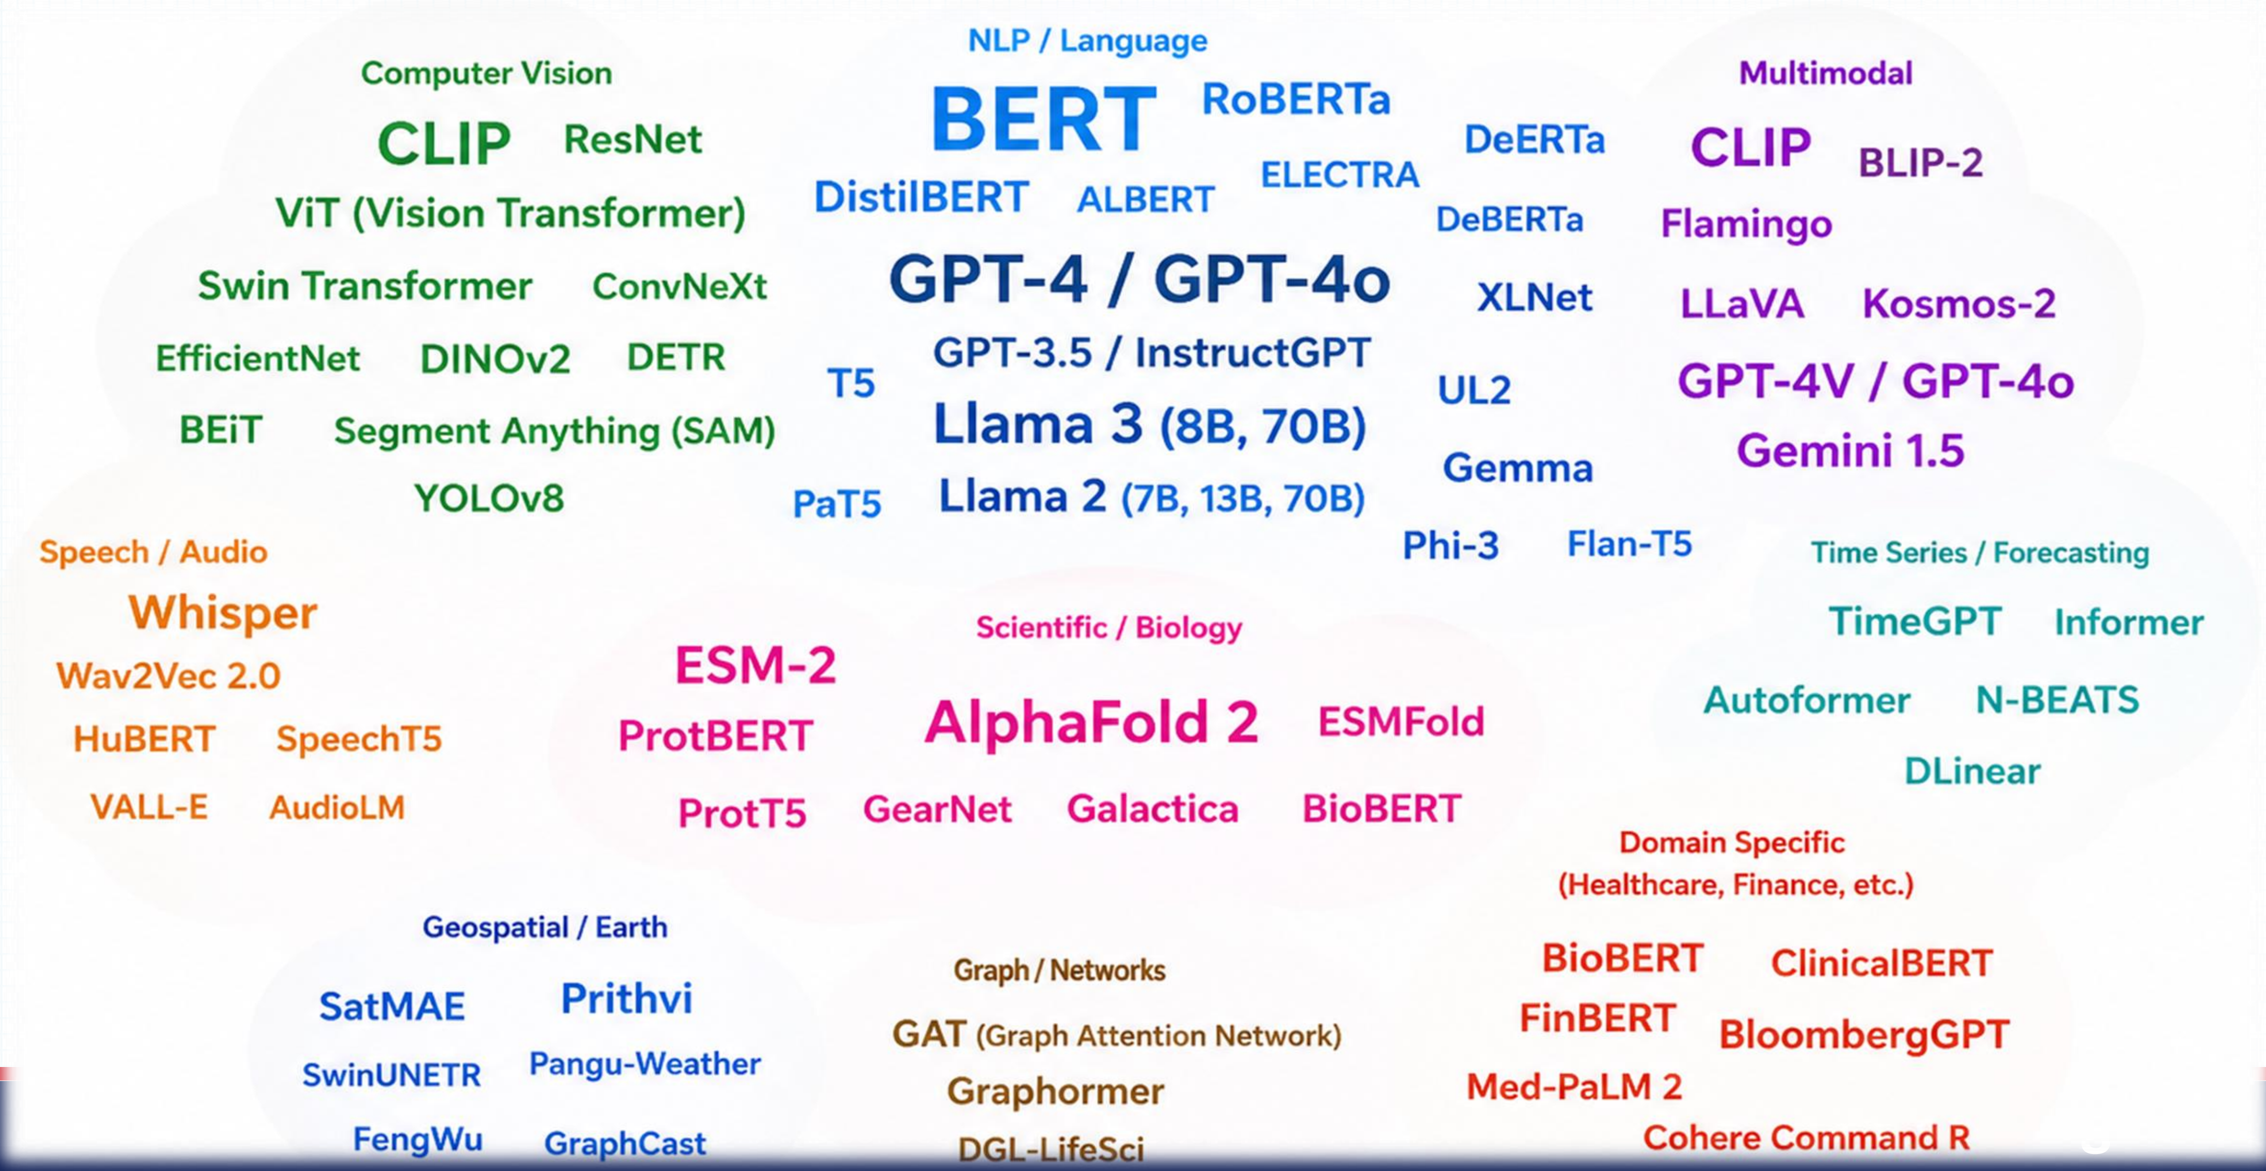

How to use these models in your research?

How to use these models in your research?

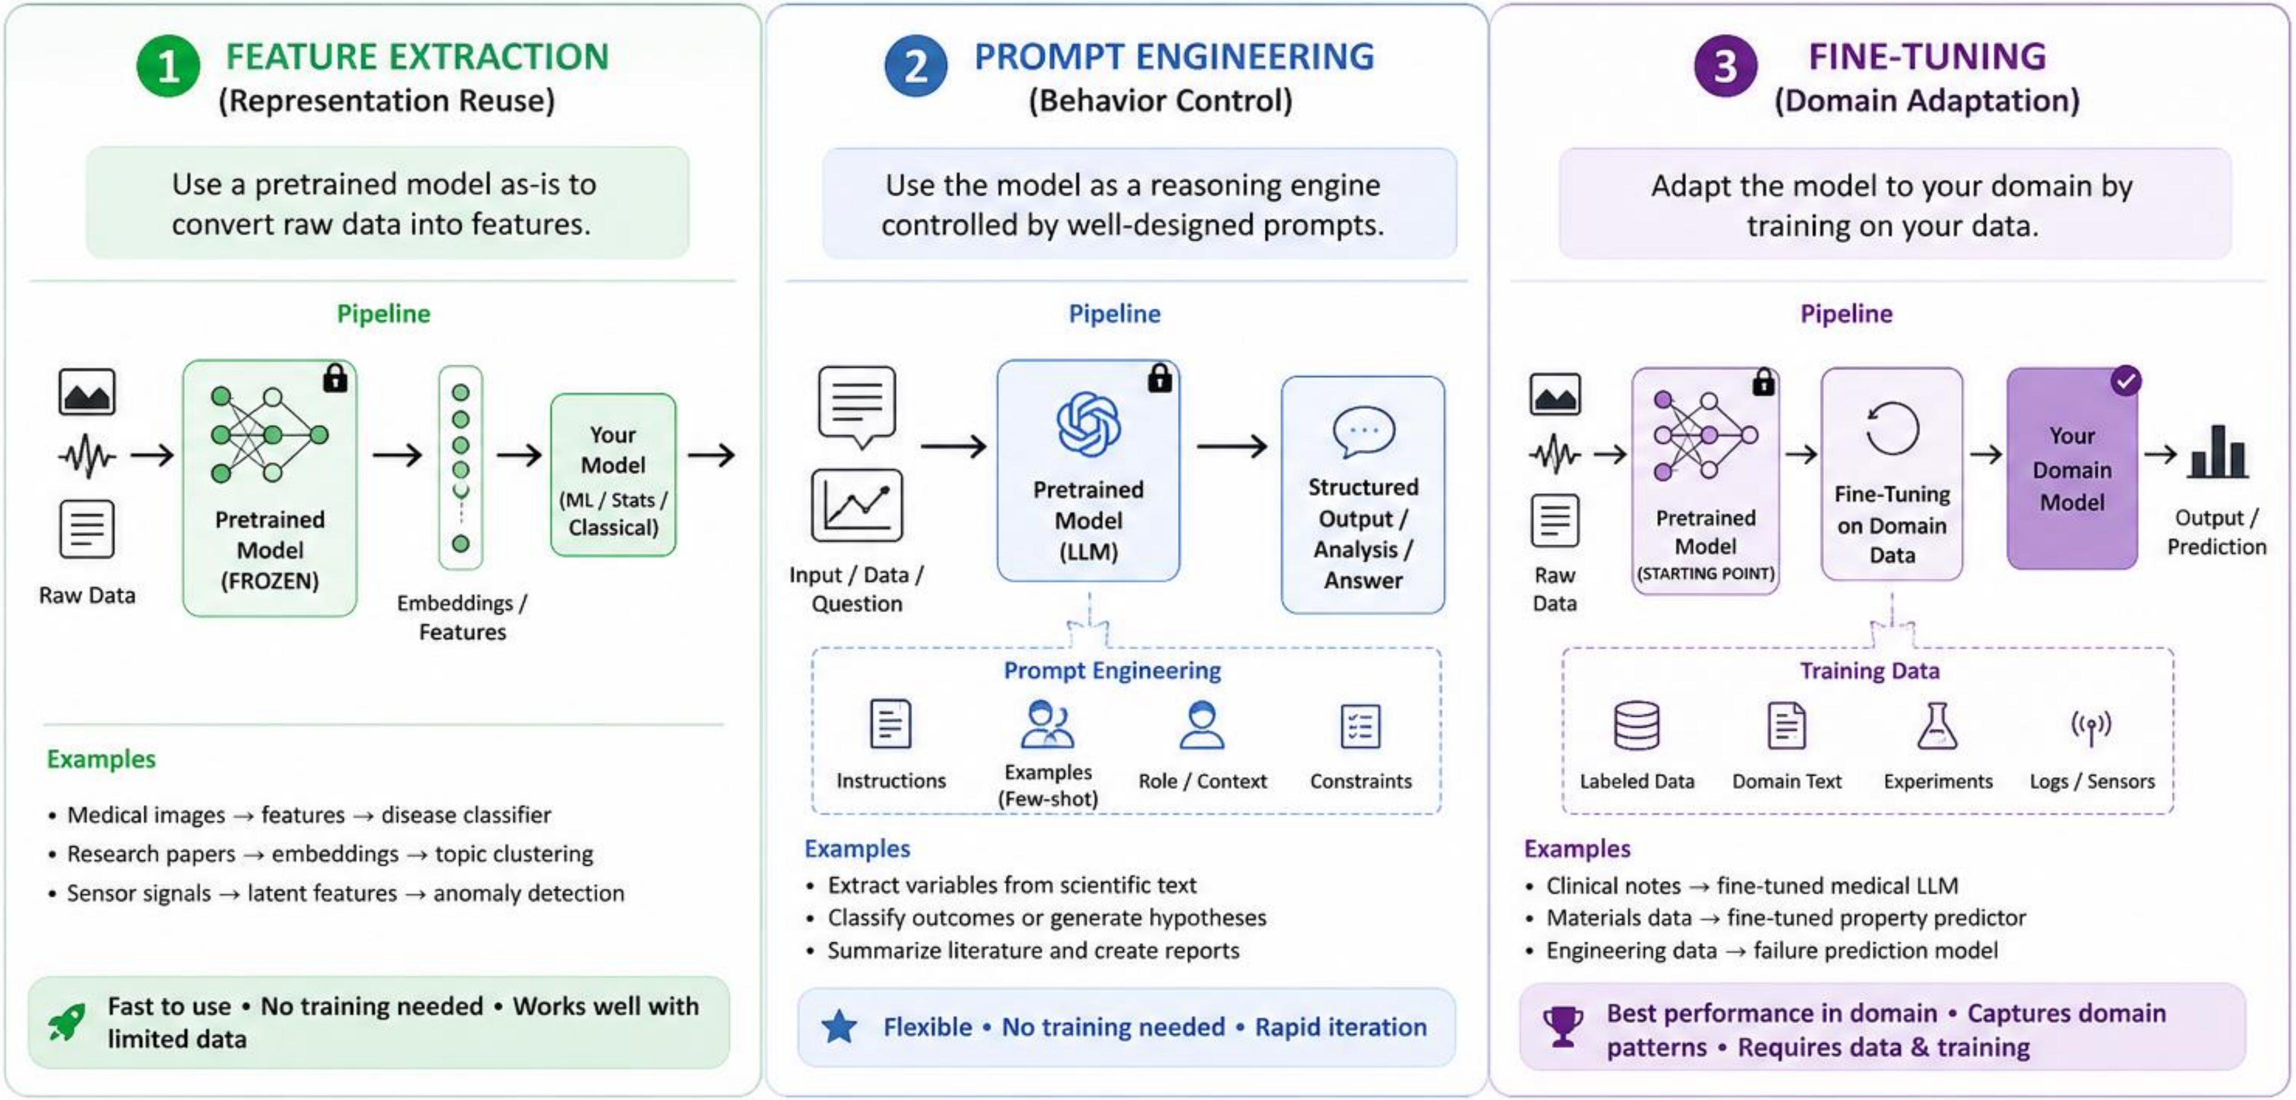

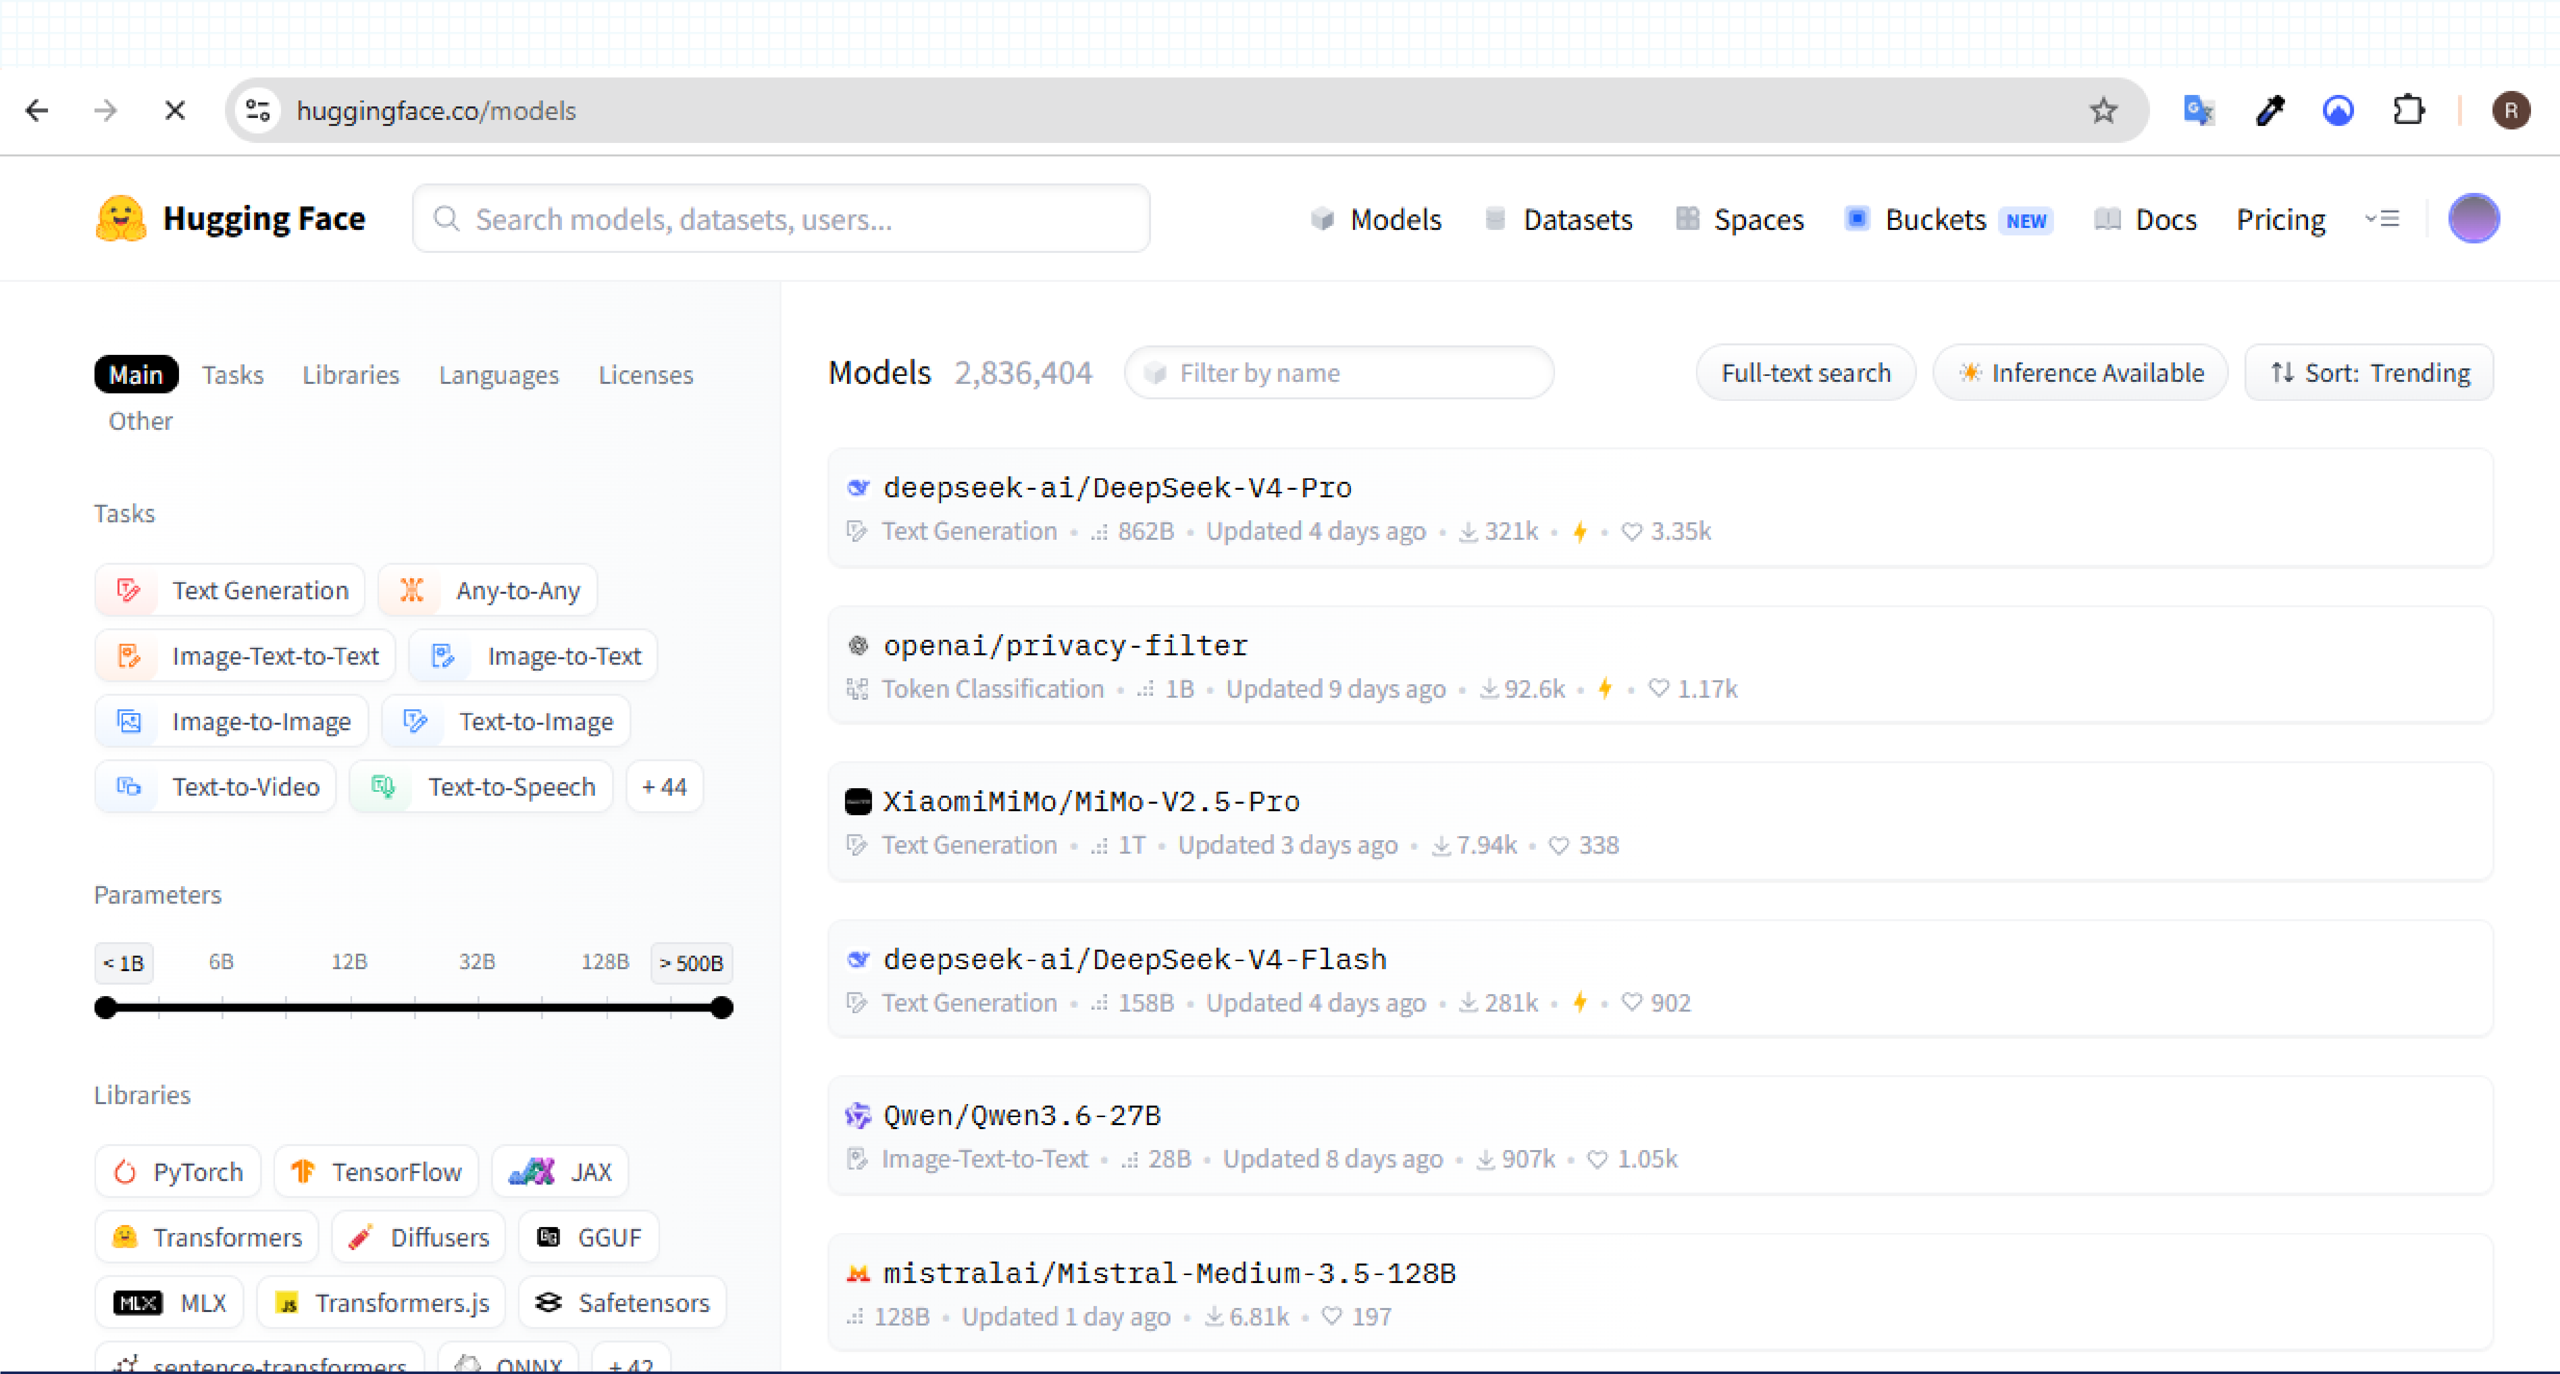
huggingface.co/models
Hugging Face
Search models, datasets, users...
Models
Datasets
Spaces
Buckets NEW
Docs
Pricing
Main
Tasks
Libraries
Languages
Licenses
Other
Tasks
Text Generation
Any-to-Any
Image-Text-to-Text
Image-to-Text
Image-to-Image
Text-to-Image
Text-to-Video
Text-to-Speech
+ 44
Parameters
< 1B
6B
12B
32B
128B
> 500B
Libraries
PyTorch
TensorFlow
JAX
Transformers
Diffusers
GGUF
MLX
Transformers.js
Safetensors
Models 2,836,404
Filter by name
Full-text search
Inference Available
Sort: Trending
deepseek-ai/DeepSeek-V4-Pro
Text Generation • 862B • Updated 4 days ago • 321k • 3.35k
openai/privacy-filter
Token Classification • 1B • Updated 9 days ago • 92.6k • 1.17k
XiaomiMiMo/MiMo-V2.5-Pro
Text Generation • 1T • Updated 3 days ago • 7.94k • 338
deepseek-ai/DeepSeek-V4-Flash
Text Generation • 158B • Updated 4 days ago • 281k • 902
Qwen/Qwen3.6-27B
Image-Text-to-Text • 28B • Updated 8 days ago • 907k • 1.05k
mistralai/Mistral-Medium-3.5-128B
128B • Updated 1 day ago • 6.81k • 197

huggingface.co/Rostlab/prot_bert

Hugging Face | Search models, datasets, users... | Models | Datasets | Spaces | Buckets NEW | Docs | Pricing

Rostlab/**prot_bert** | like 132 | Follow Rostlab 87

Fill-Mask | Transformers | PyTorch | Uniref100 | protein language model | protein

Model card | Files and versions xet | Community 9 | Deploy | Use this model

Edit model card

ProtBert model

Pretrained model on protein sequences using a masked language modeling (MLM) objective. It was introduced in this paper and first released in this repository. This model is trained on uppercase amino acids: it only works with capital letter amino acids.

Model description

ProtBert is based on Bert model which pretrained on a large corpus of protein sequences in a self-supervised fashion. This means it was pretrained on the raw protein sequences only, with no humans labelling them in any way (which is why it can use lots of publicly available data) with an automatic process to generate inputs

Downloads last month
462,675

Inference Providers NEW | HF Inference API

Fill-Mask | Examples

No mask token found for this model.

Your prompt here...

Generate

View Code Snippets | Maximize

1 FEATURE EXTRACTION
(Representation Reuse)

Use a pretrained model as-is to convert raw data into features.

Pipeline

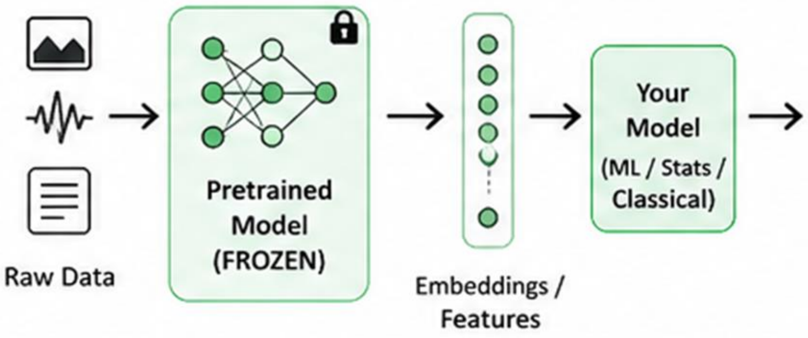

Examples

- Medical images → features → disease classifier
- Research papers → embeddings → topic clustering
- Sensor signals → latent features → anomaly detection

Fast to use • No training needed • Works well with limited data

```python
from sentence_transformers import SentenceTransformer
from sklearn.linear_model import LogisticRegression

# 1. Frozen pretrained encoder
encoder = SentenceTransformer("all-MiniLM-L6-v2")

# 2. Your research dataset
texts = [
    "Patient shows signs of pneumonia",
    "No evidence of lung disease",
    "Severe respiratory infection detected"
]
labels = [1, 0, 1]

# 3. Feature extraction (no training on encoder)
X = encoder.encode(texts)

# 4. Downstream model
clf = LogisticRegression()
clf.fit(X, labels)

# 5. Prediction
test = encoder.encode(["Possible lung infection"])
print(clf.predict(test))
```

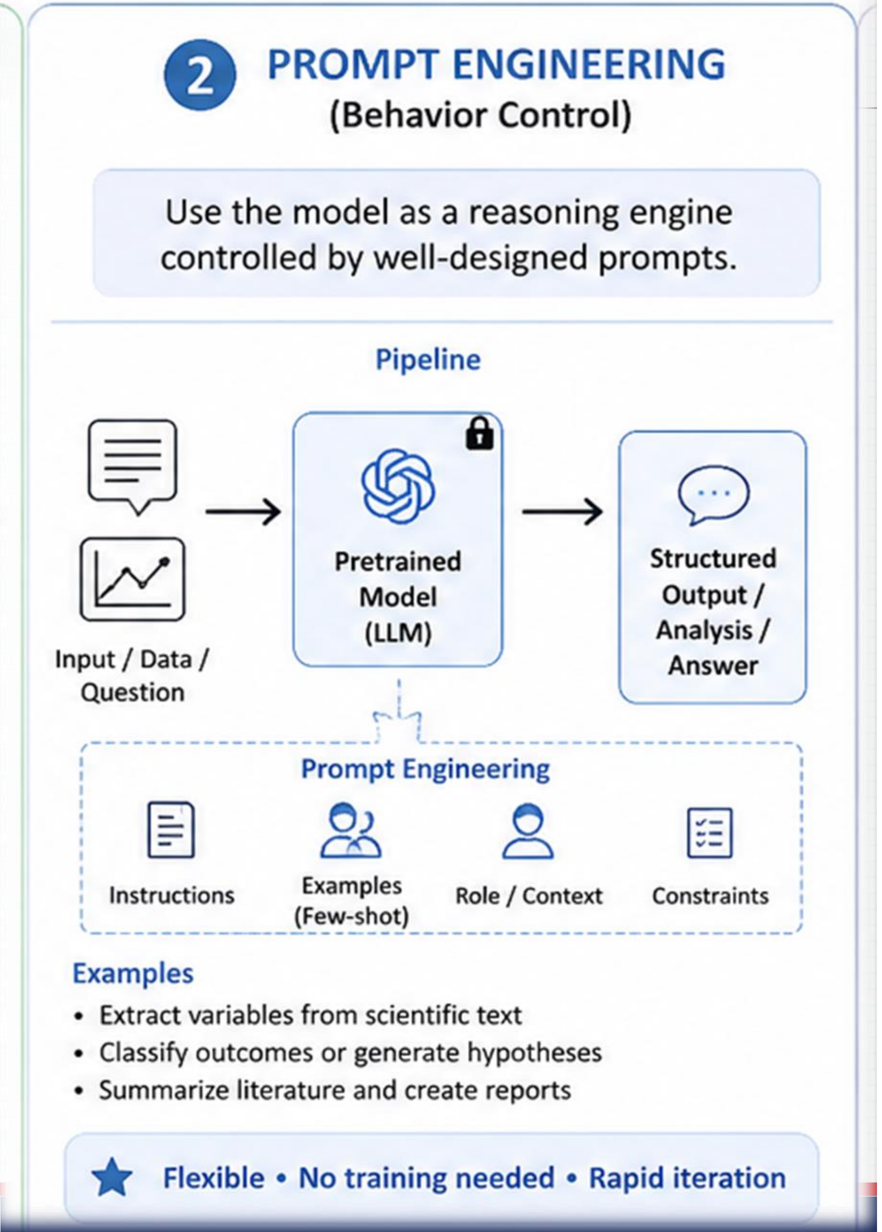


```python
from openai import OpenAI

client = OpenAI()

prompt = """
You are a biomedical research assistant.

Extract from the text:
- disease
- symptoms
- treatment

Return JSON only.
Text: The patient has fever and cough, treated with antibiotics.
"""


response = client.chat.completions.create(
    model="gpt-4.1-mini",
    messages=[{"role": "user", "content": prompt}]
)

print(response.choices[0].message.content)
```

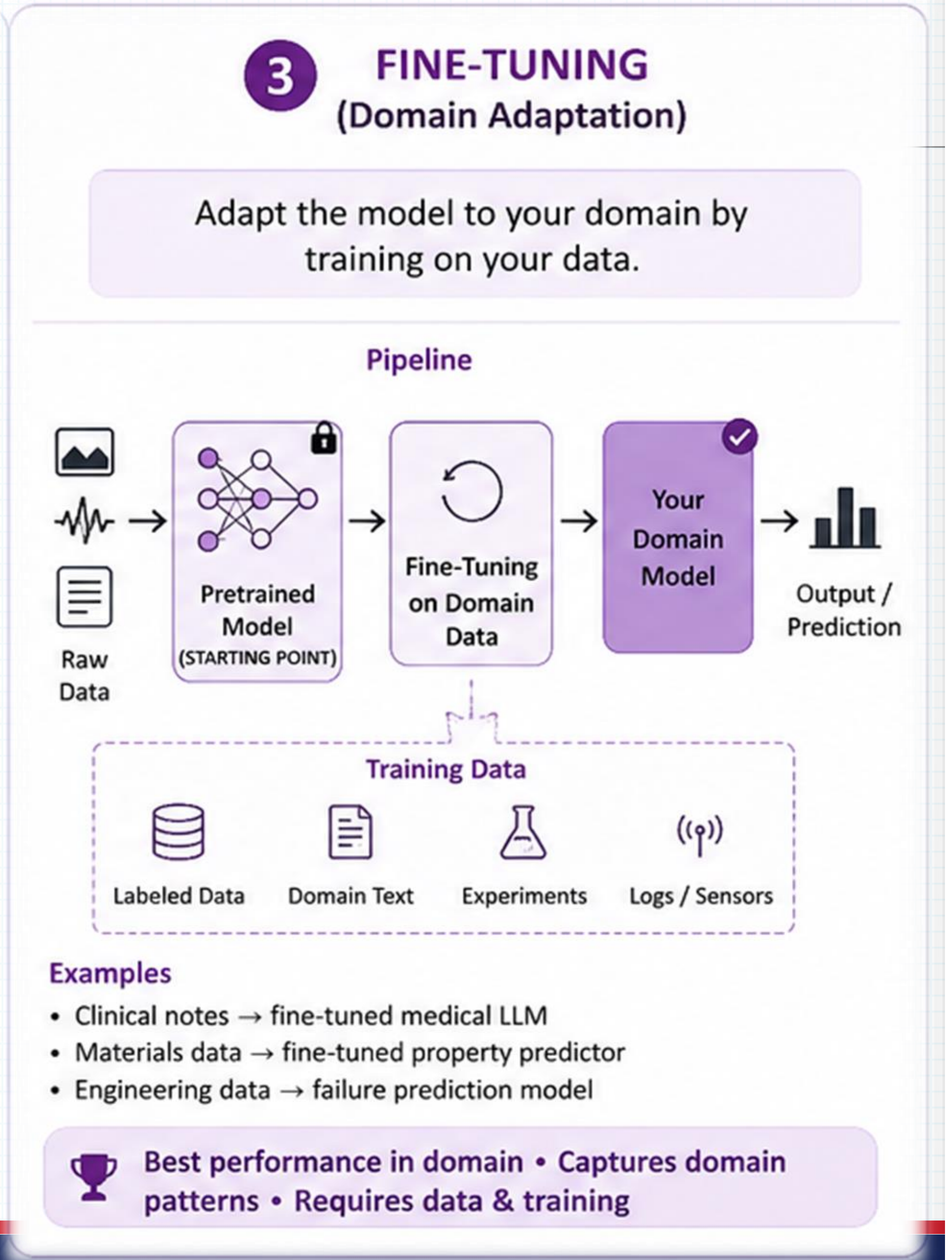


```python
# 1. Domain dataset
data = {
    "text": [
        "MRI shows brain tumor",
        "Normal CT scan",
        "Abnormal cardiac rhythm detected"
    ],
    "label": [1, 0, 1]
}

dataset = Dataset.from_dict(data)

# 2. Model
model = AutoModelForSequenceClassification.from_pretrained( "bert-base-uncased",

# 3. Training setup
args = TrainingArguments(
    output_dir="./model",
    num_train_epochs=3,
    per_device_train_batch_size=8
)

trainer = Trainer( model=model, args=args, train_dataset=dataset )

# 4. Fine-tune
trainer.train()

# 5. Now model is domain-specific
```

How to build these Embedding Vectors

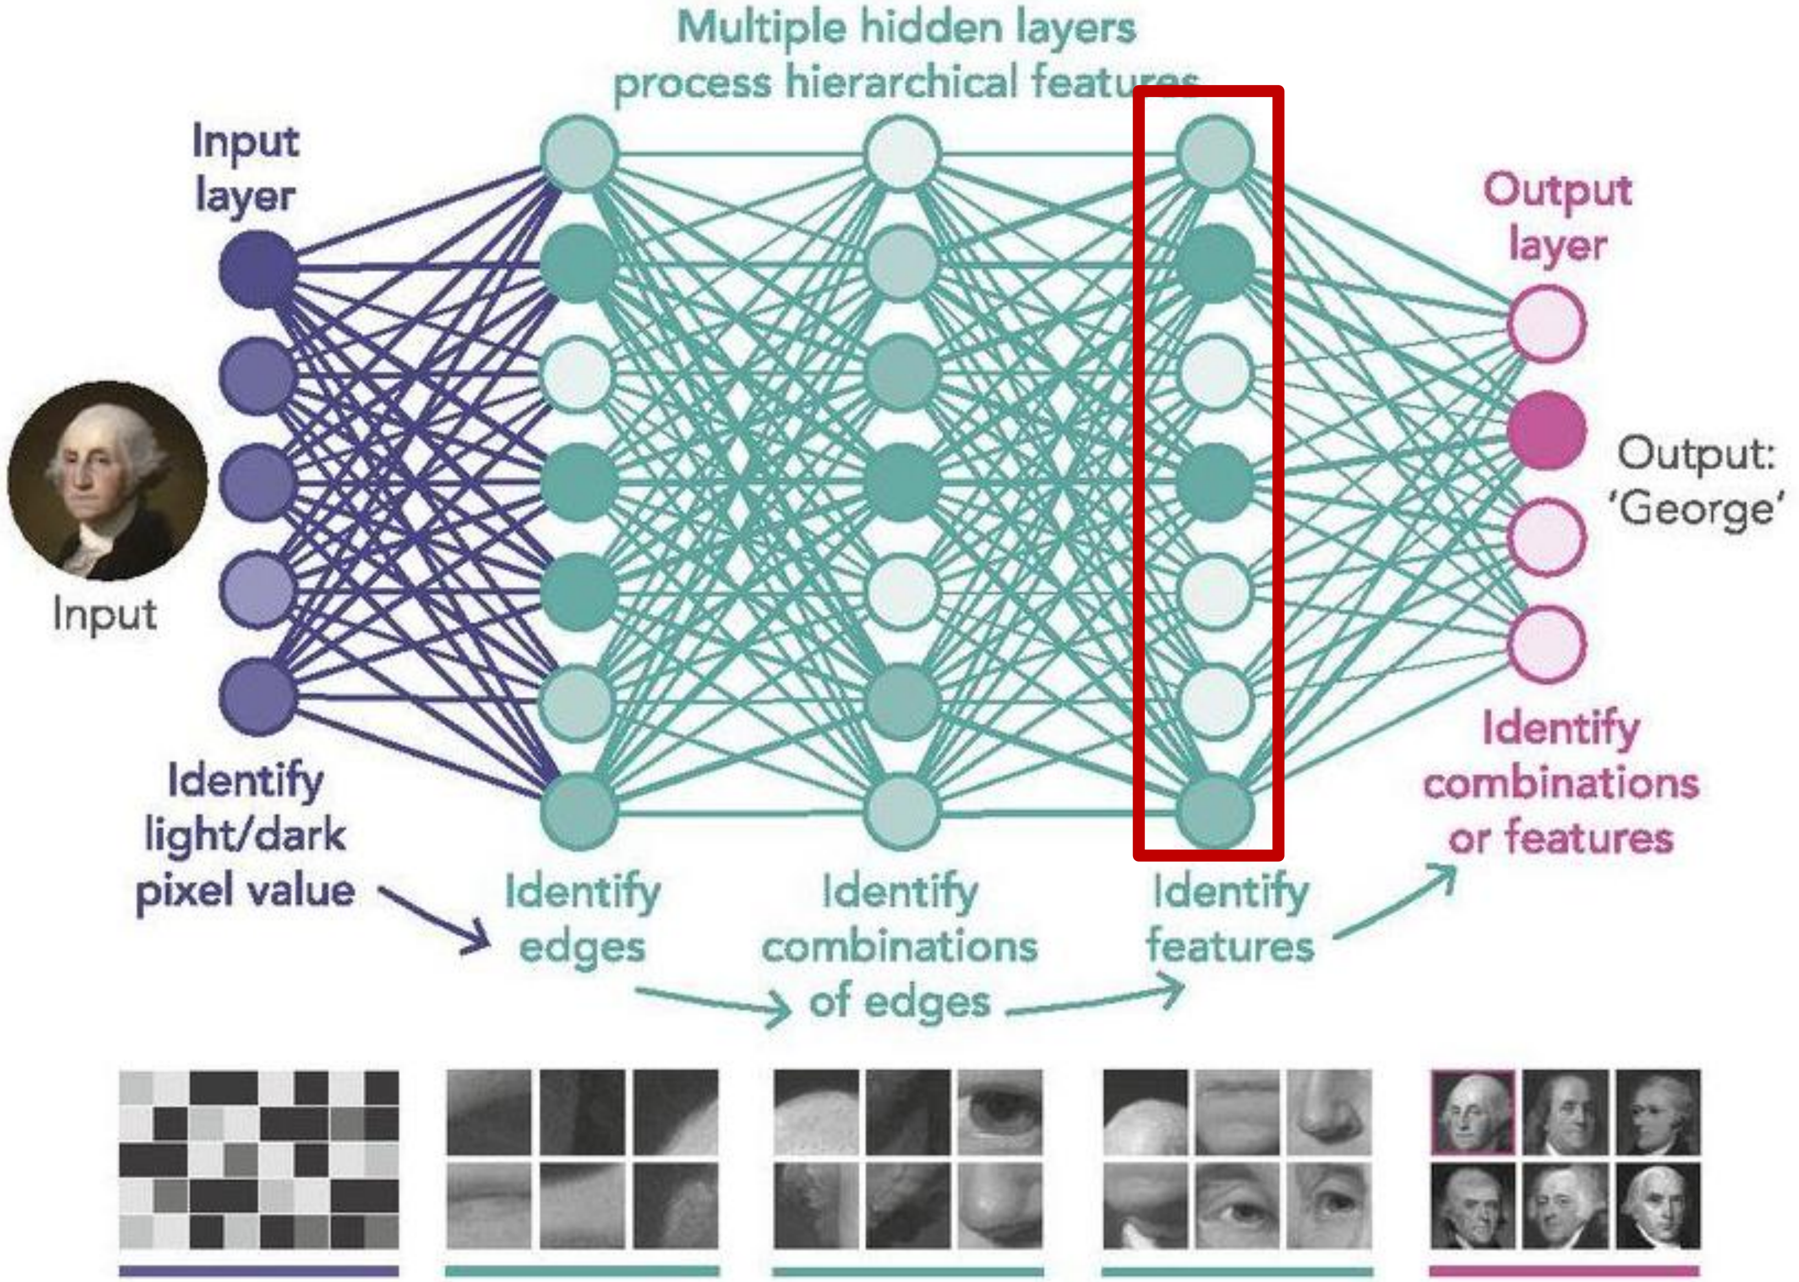
Multiple hidden layers
process hierarchical features
Input
layer
Input
Output
layer
Output:
'George'
Identify
light/dark
pixel value
Identify
edges
Identify
combinations
of edges
Identify
features
Identify
combinations
or features

Example: CNN

cat dog ……

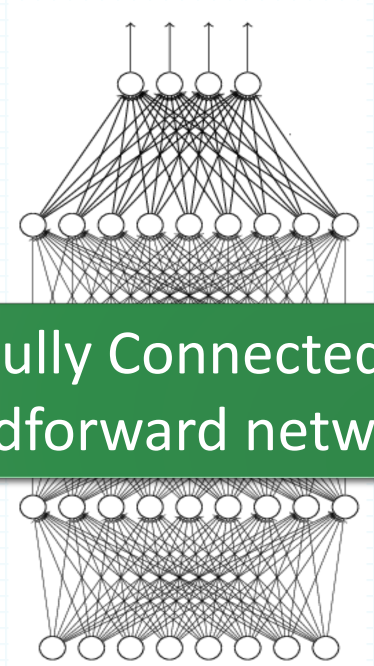

Fully Connected Feedforward network

Convolution

Max Pooling

A new image

Convolution

Max Pooling

A new image

Flatten

BUT…

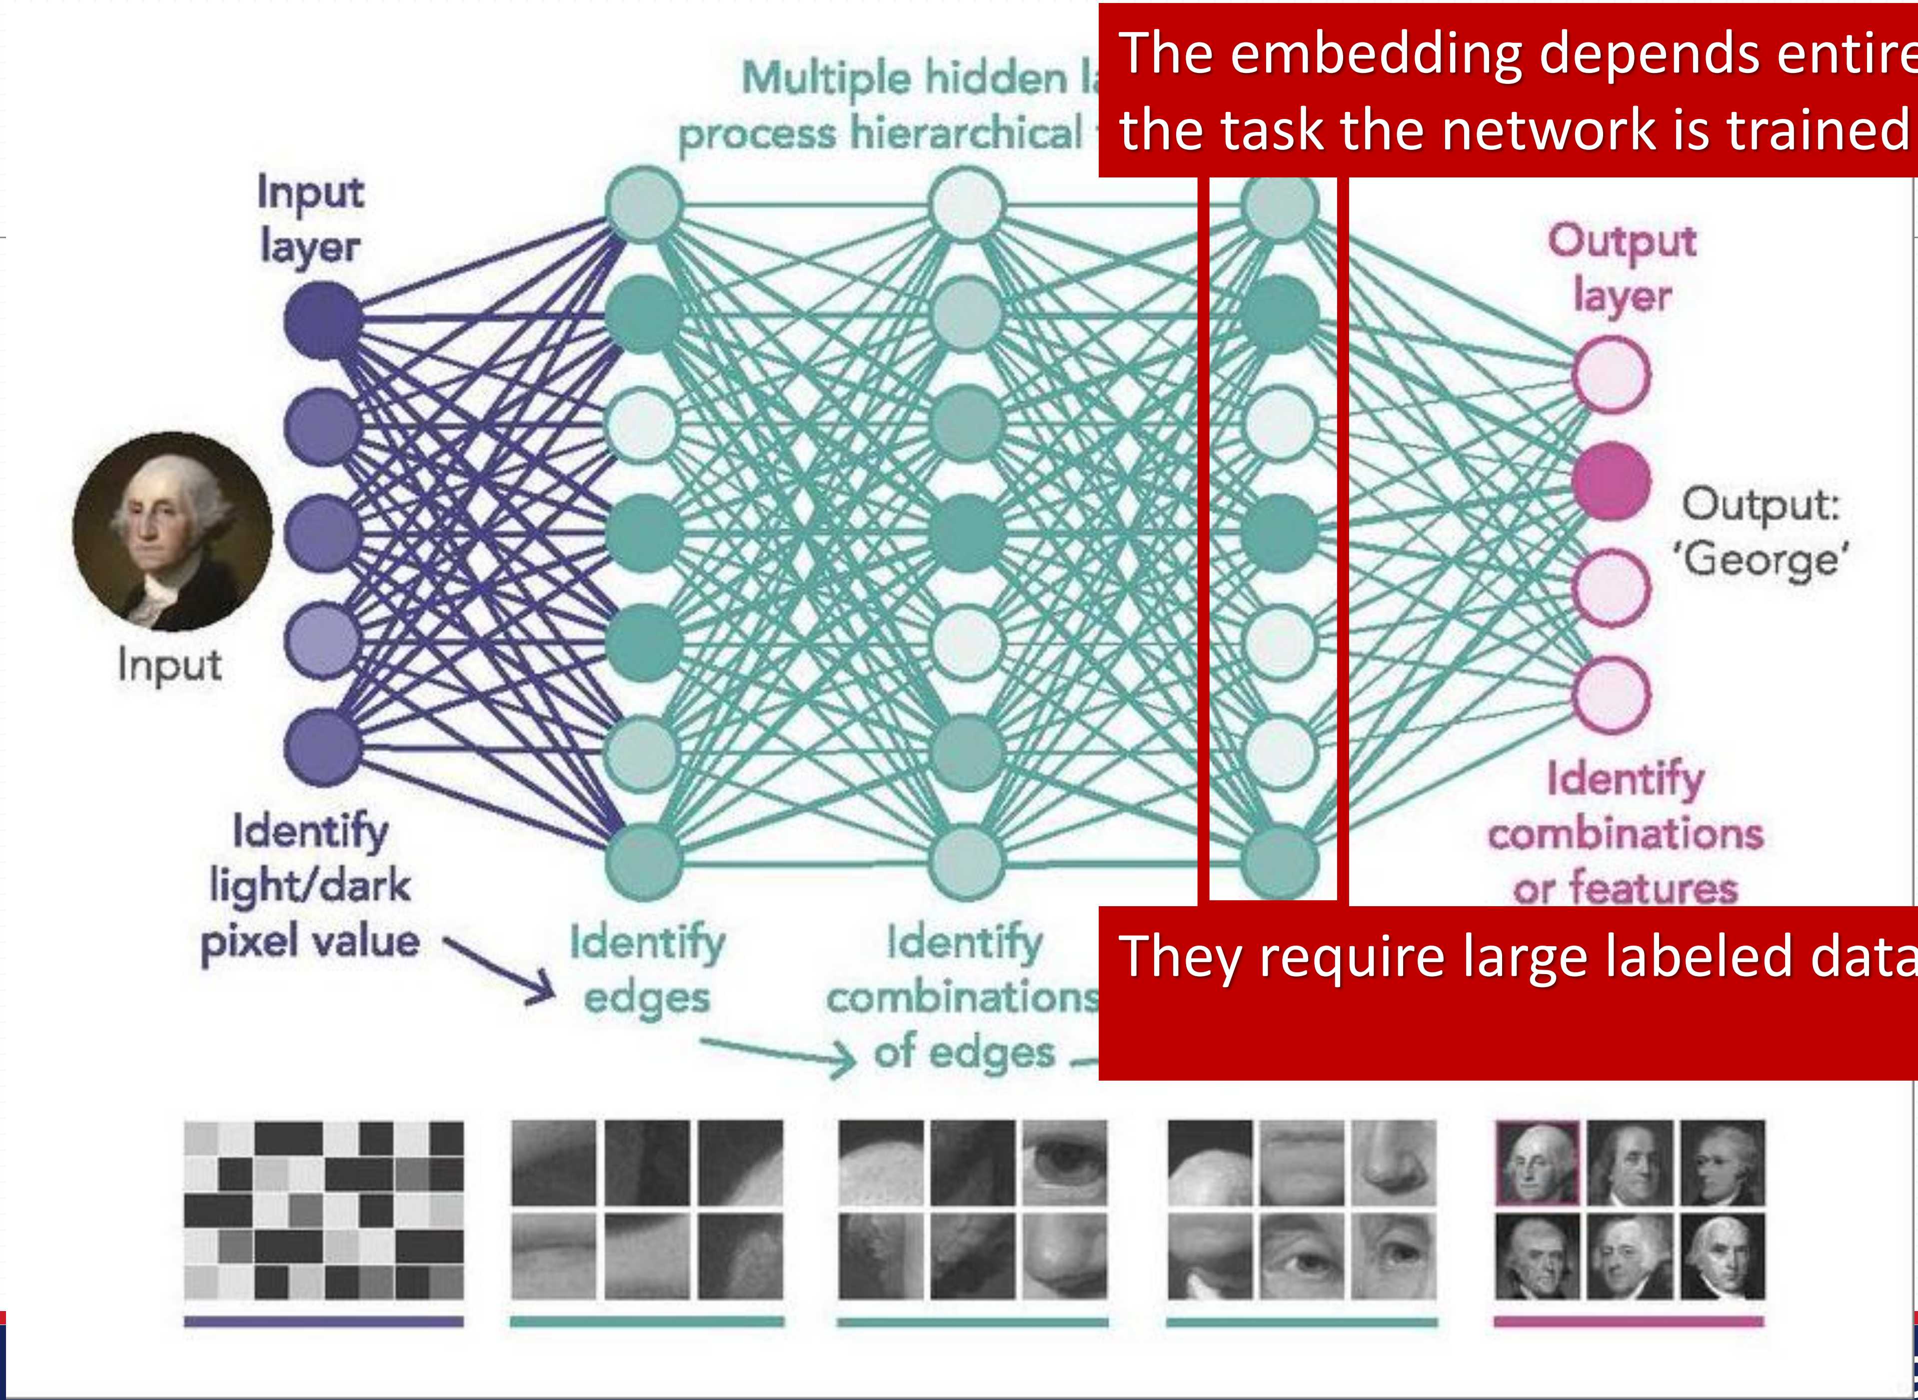

The embedding depends entirely on the task the network is trained on.

They require large labeled datasets!

Solution: Self-supervised learning

- This idea—**"given a word, predict the surrounding words"**

she had to pay using ___ instead of any digital or card-based method.

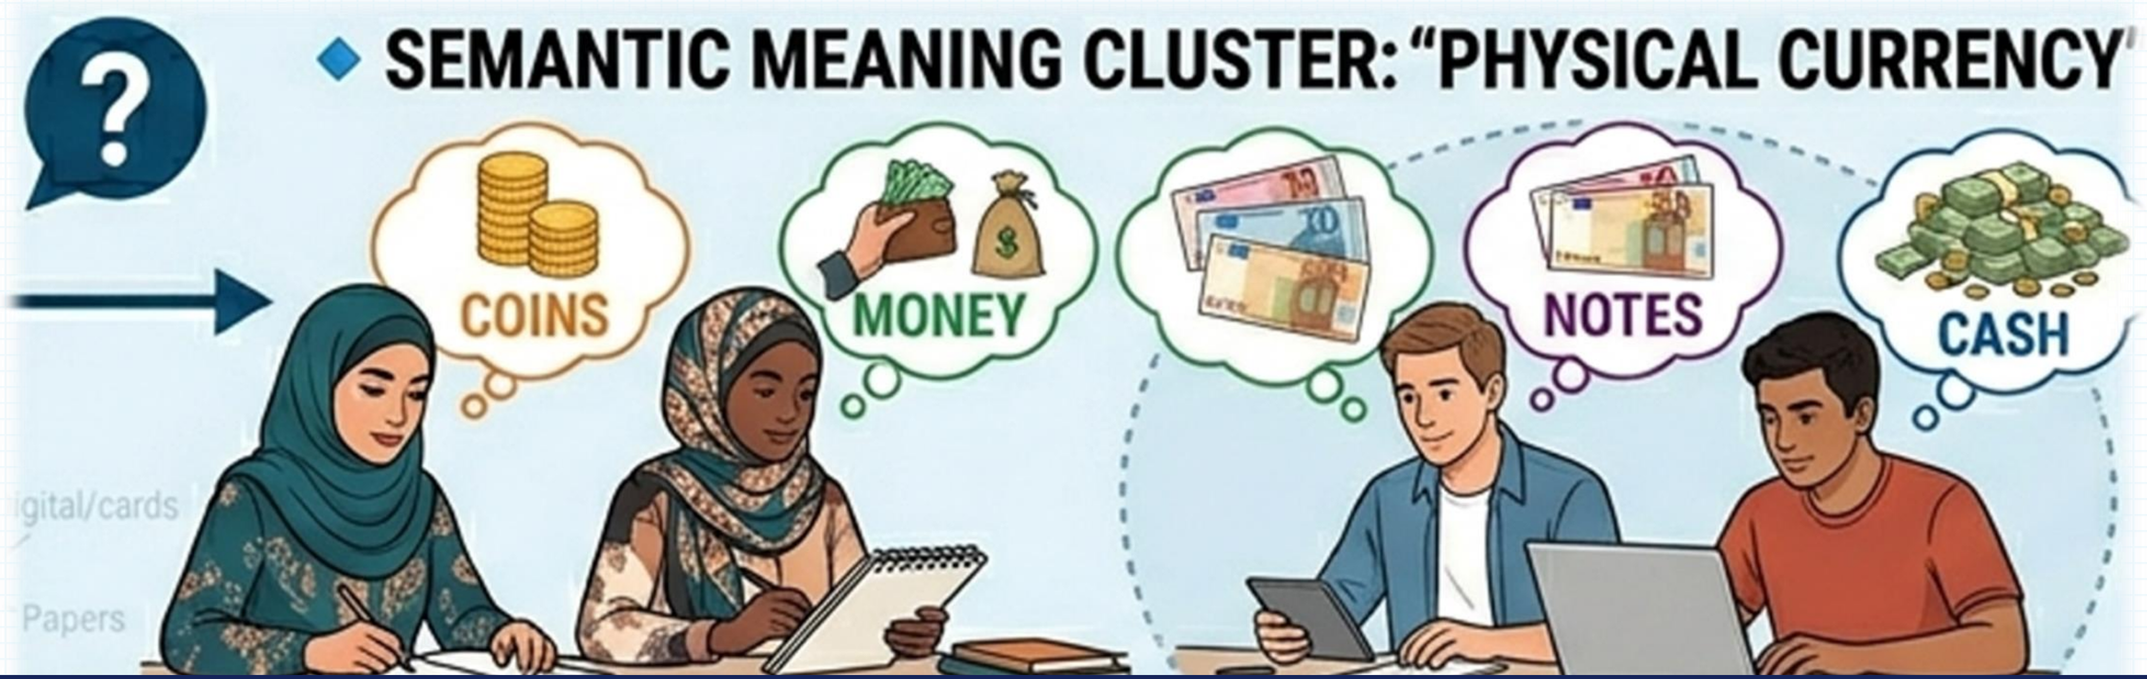

The model is not learning exact word replacement. It is learning a region of meaning.

word2vec

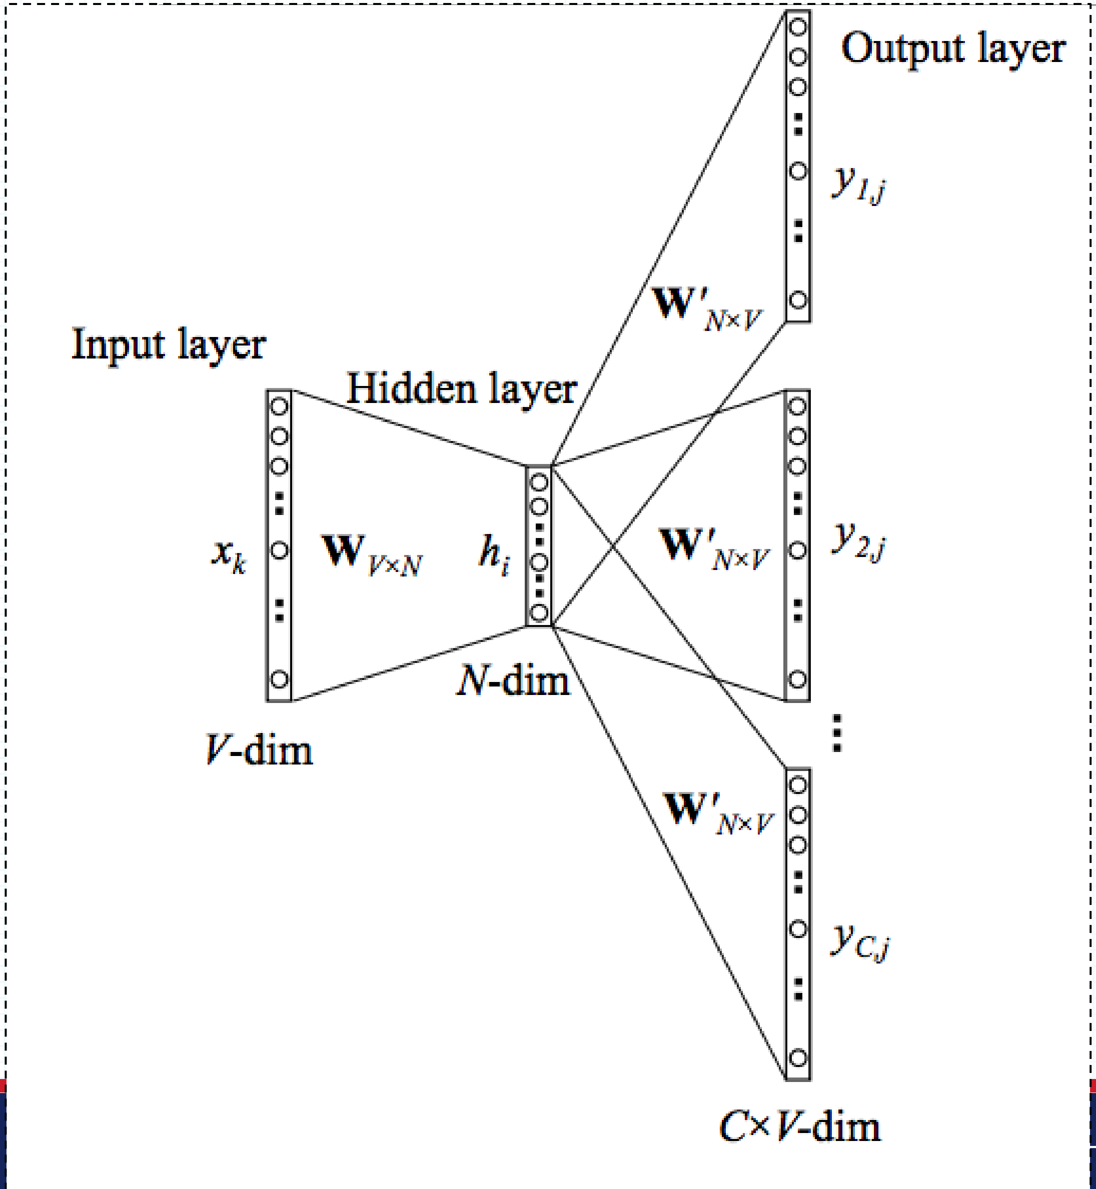
Output layer
Input layer
Hidden layer
x_k
$\mathbf{W}_{V\times N}$
h_i
N-dim
V-dim
$\mathbf{W}'_{N\times V}$
$\mathbf{W}'_{N\times V}$
$\mathbf{W}'_{N\times V}$
$y_{1,j}$
$y_{2,j}$
$y_{C,j}$
$C\times V$-dim

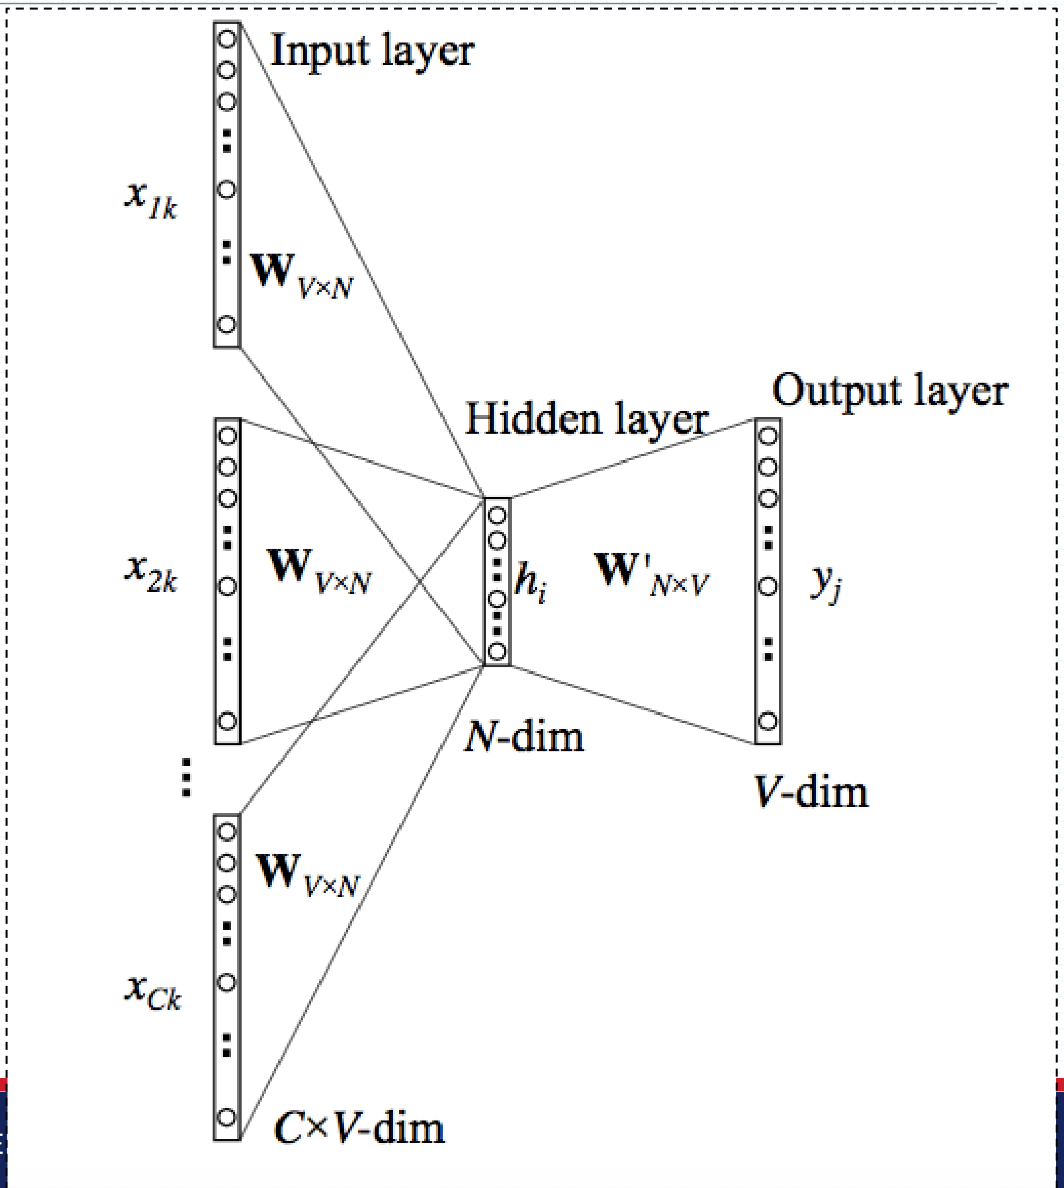
Input layer
x_{1k}
$\mathbf{W}_{V\times N}$
x_{2k}
$\mathbf{W}_{V\times N}$
x_{Ck}
$\mathbf{W}_{V\times N}$
$C\times V$-dim
Hidden layer
h_i
N-dim
$\mathbf{W}'_{N\times V}$
Output layer
y_j
V-dim

Embedding Vector

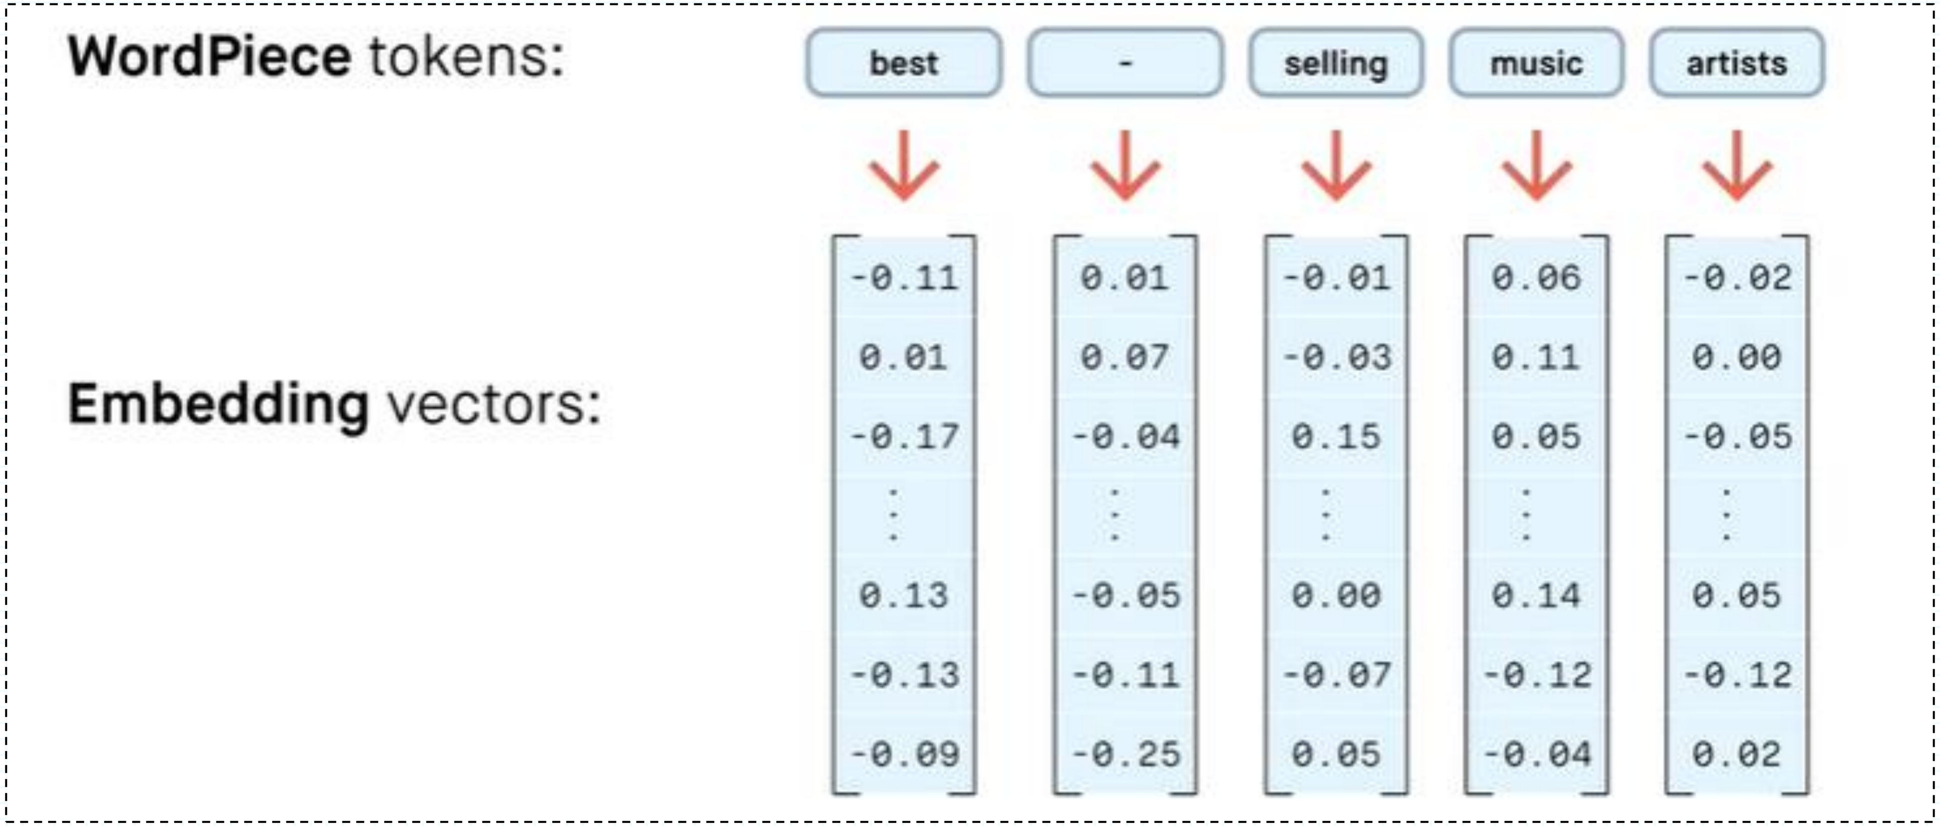

But… Contextual Embedding

Solution.. Attention

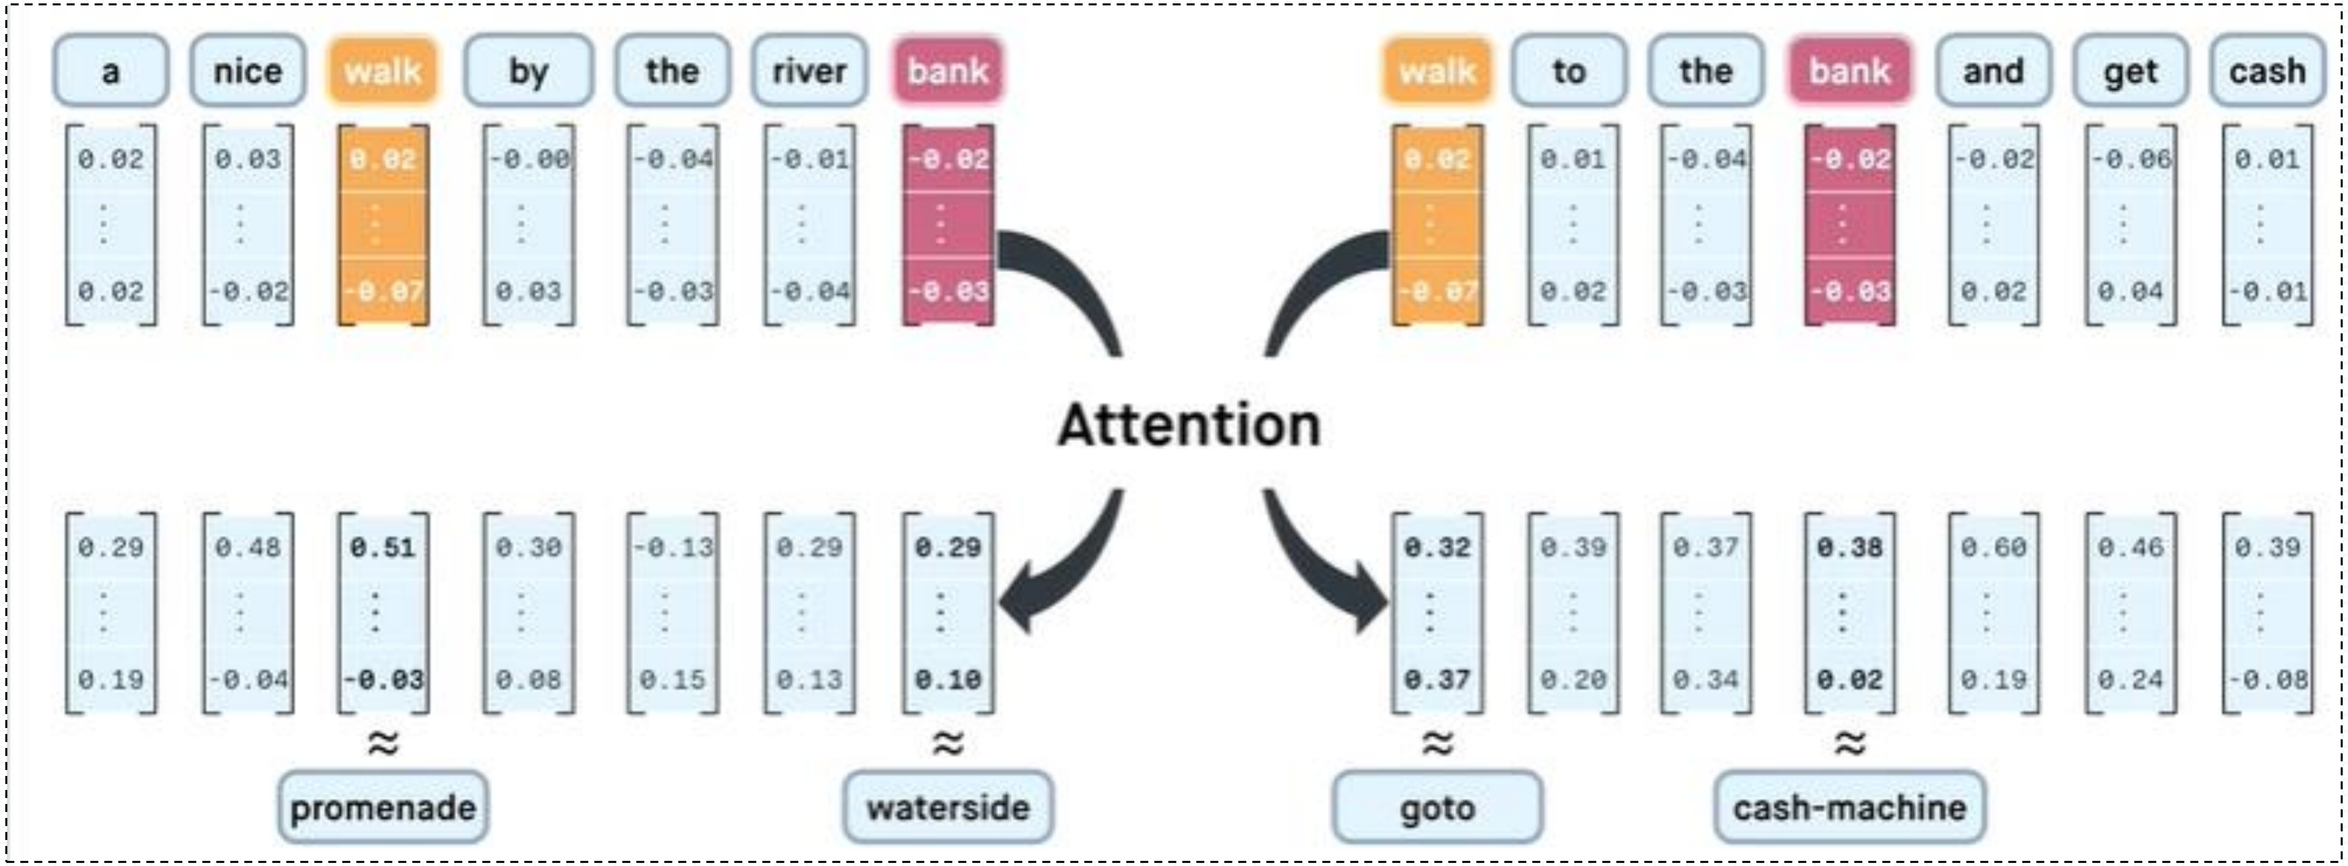

Transformers

- a Transformer architecture can be designed in **three different ways** depending on the task and objective:
- 1. **Encoder-Only Transformers**
- **2. Decoder-Only Transformers**
- 3. **Encoder-Decoder Transformers (Seq2Seq)**

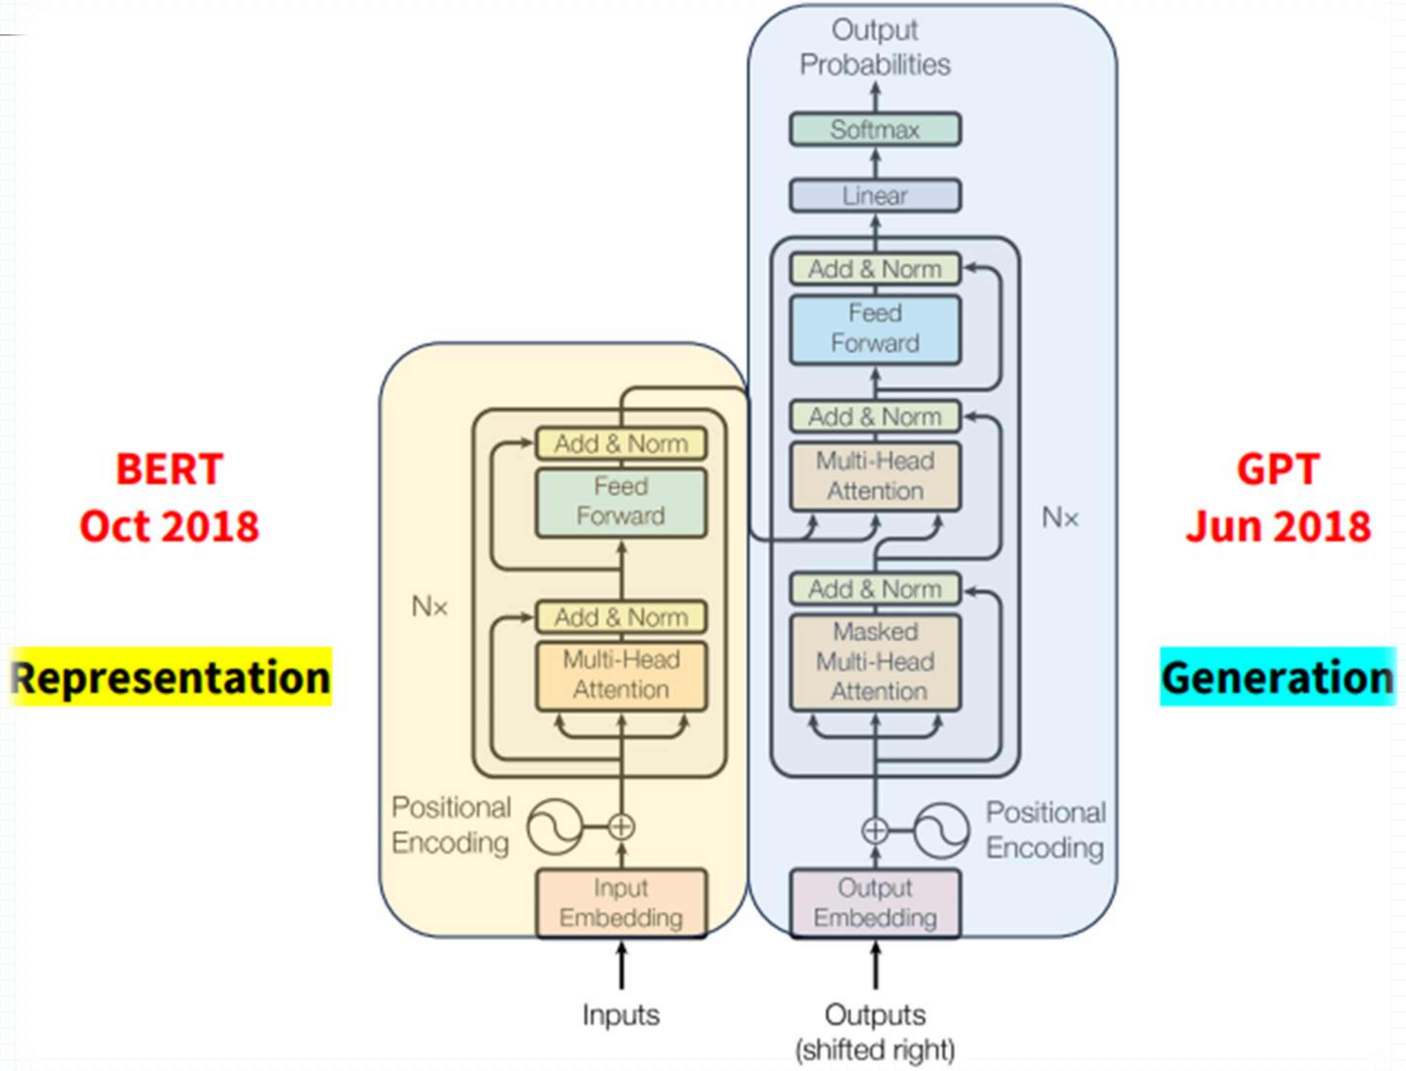

Transformers… The LLM Era

BERT – 2018
DistilBERT – 2019
RoBERTa – 2019
ALBERT – 2019
ELECTRA – 2020
DeBERTa – 2020
…

Representation

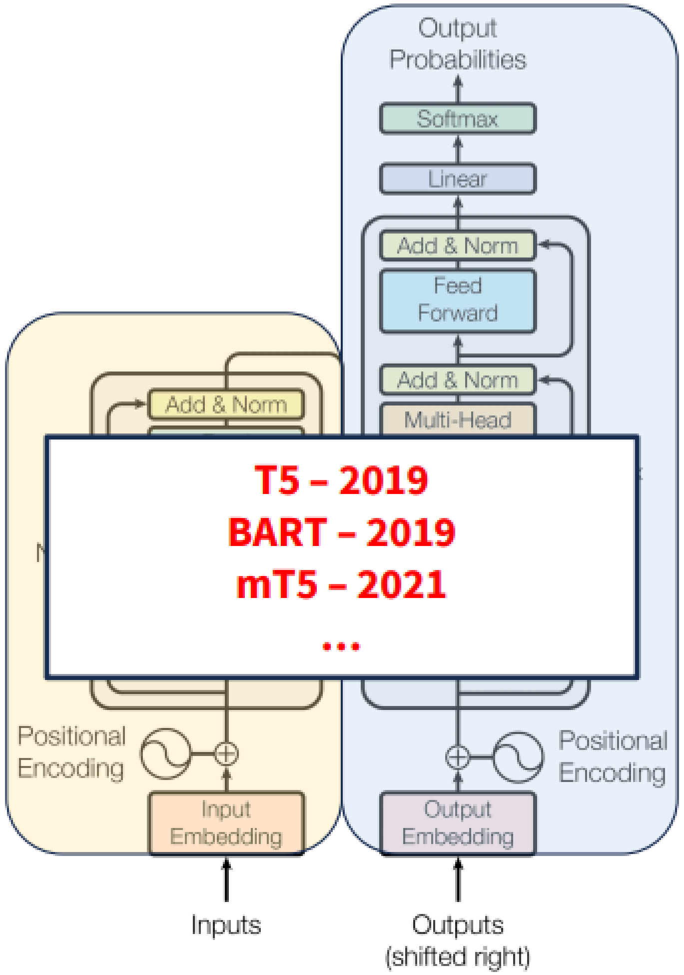

T5 – 2019
BART – 2019
mT5 – 2021
…

GPT – 2018
GPT-2 – 2019
GPT-3 – 2020
GPT-Neo – 2021
GPT-3.5 (ChatGPT) – 2022
LLaMA – 2023
GPT-4 – 2023
…

Generation

Major challenges & research directions in modern AI


CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Data Dependency and Data Quality
Generalization
Long-Term Planning


Self-Supervised Learning
Few-Shot Learning
Zero-Shot Learning
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Data Dependency and Data Quality
Generalization
Long-Term Planning


Transfer Learning
Domain Adaptation
Foundation Models
Meta-learning
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Data Dependency and Data Quality
Generalization
Long-Term Planning

Agentic AI
Reinforcement learning
Chain-of-Thought
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Generalization
Long-Term Planning
Data Dependency and Data Quality


Retrieval-Augmented Generation (RAG)
Fact Verification
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Data Dependency and Data Quality
Generalization
Long-Term Planning

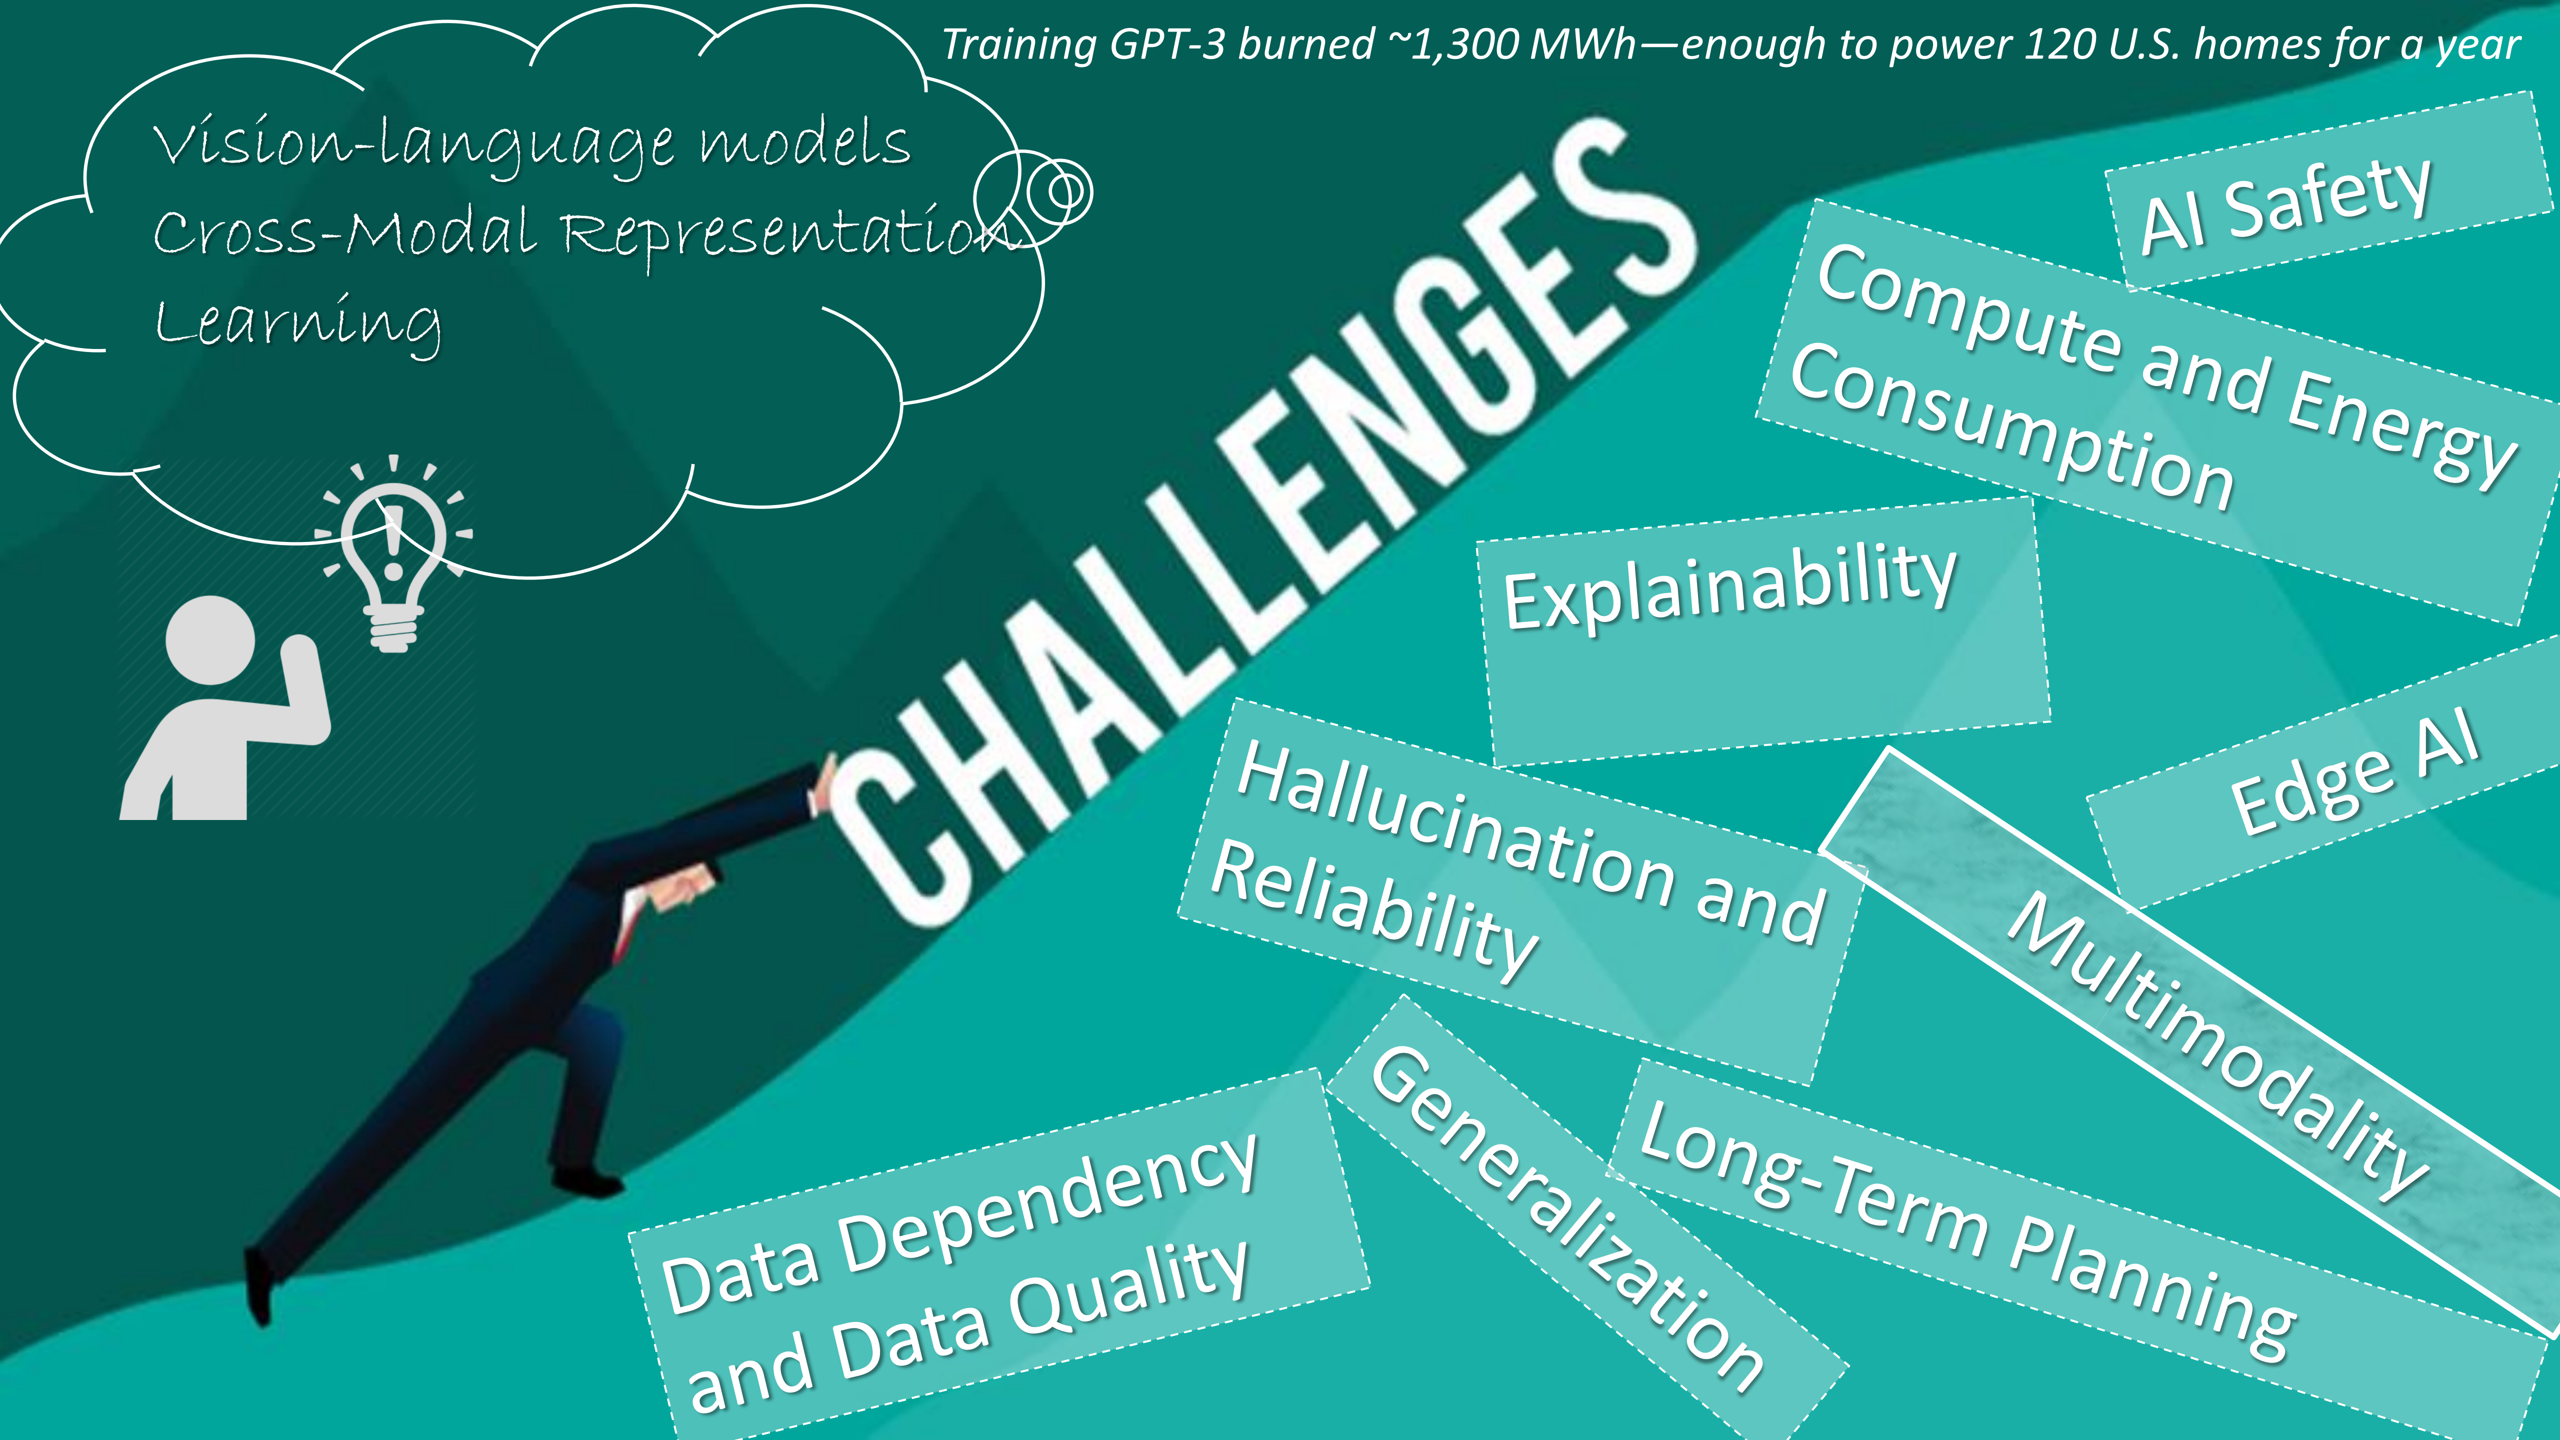
Training GPT-3 burned ~1,300 MWh—enough to power 120 U.S. homes for a year
Vision-language models
Cross-Modal Representation
Learning
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Data Dependency and Data Quality
Generalization
Long-Term Planning


TinyML
Efficient transformers
Model compression
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Generalization
Long-Term Planning
Data Dependency and Data Quality


Interpretable transformers
Attention analysis
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Data Dependency and Data Quality
Generalization
Long-Term Planning


Efficient architectures
Quantization
Distillation
LoRA / QLoRA
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Data Dependency and Data Quality
Generalization
Long-Term Planning


Adversarial Testing
AI Governance & Regulation
Constitutional AI
CHALLENGES
AI Safety
Compute and Energy Consumption
Explainability
Edge AI
Hallucination and Reliability
Multimodality
Data Dependency and Data Quality
Generalization
Long-Term Planning

Road Map

- AI as a Black Box – Concepts and Foundations
- Opening the Box (1) – Traditional Machine Learning
- Opening the Box (2) – Deep Learning
- Modern AI and Research Directions

Road Map

A

M

For researchers outside AI, this shift is not about becoming machine learning experts—it is about gaining a new instrument for their own field.

Workshop Files

https://rsonbol.github.io/workshops#AI4Researchers