Demystifying Generative AI for Legal Educators

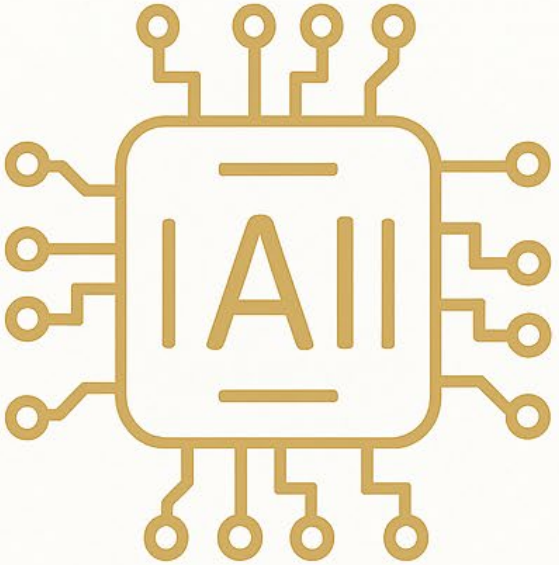

Tracy Norton & Susan Tanner

AALS

OVERVIEW

1

What is AI?

2

What is Generative AI and how does it work?

3

Introduction to Law Professor's Sandbox

4

Prompting

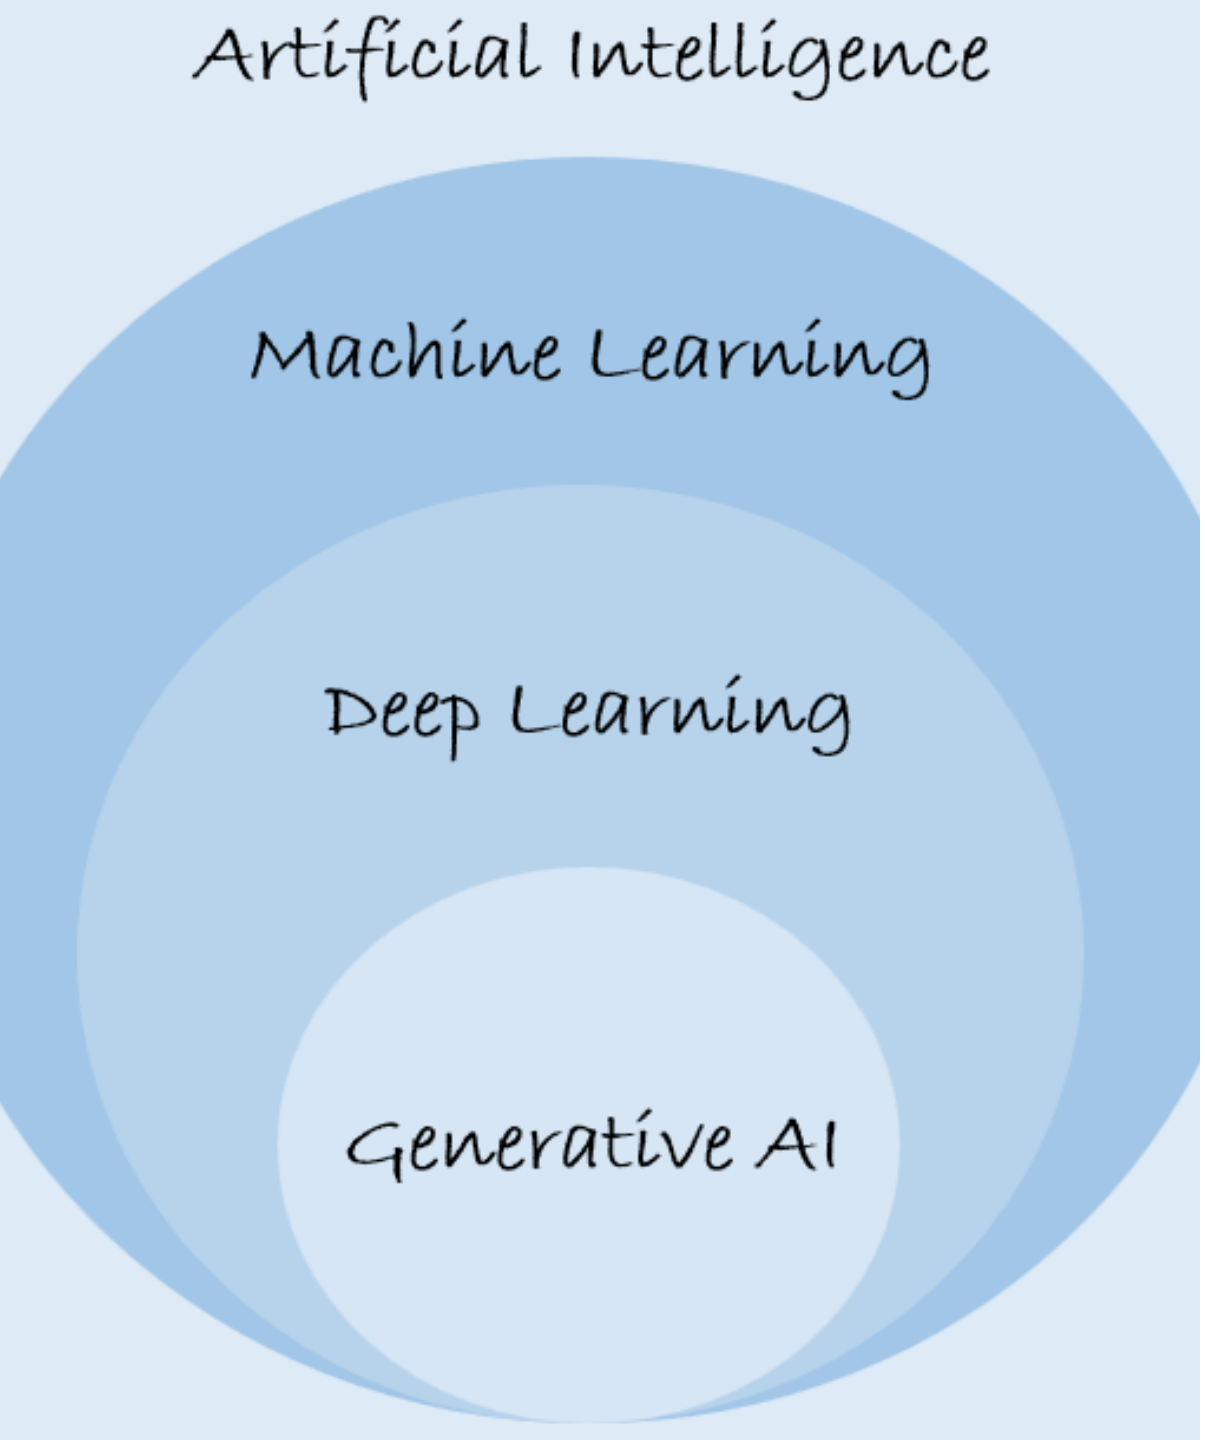

Decoding AI, Machine Learning, Learning, and Generative AI

What is AI?

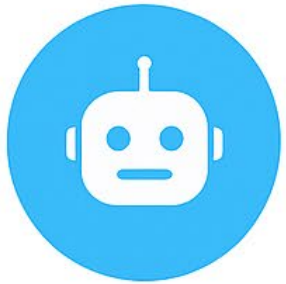

Machines performing human-like tasks

Often learning from data without explicit programming

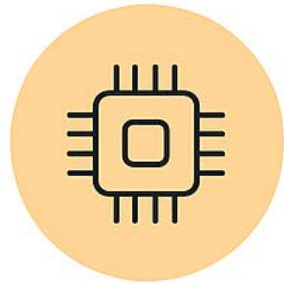

Powering a wide range of applications

Artificial Intelligence is the field of study that enables machines to simulate human intelligence and perform tasks that typically require human cognition.

WE KEEP MOVING THE TARGET.

- Whenever a once-astonishing capability becomes routine—chess grandmaster play (1997), spam filtering (2003), voice dictation (2011)—we quietly re-label it "just software" and set a new bar for what *counts* as AI. *
  - A. Gustafson, M. H. Goldberg & K. Gustafson, Why We Keep Moving the Goalposts on Artificial Intelligence (XandY Analytics 2024) ("AI means 'almost implemented.'").

AI in the Legal Industry Pre-2013

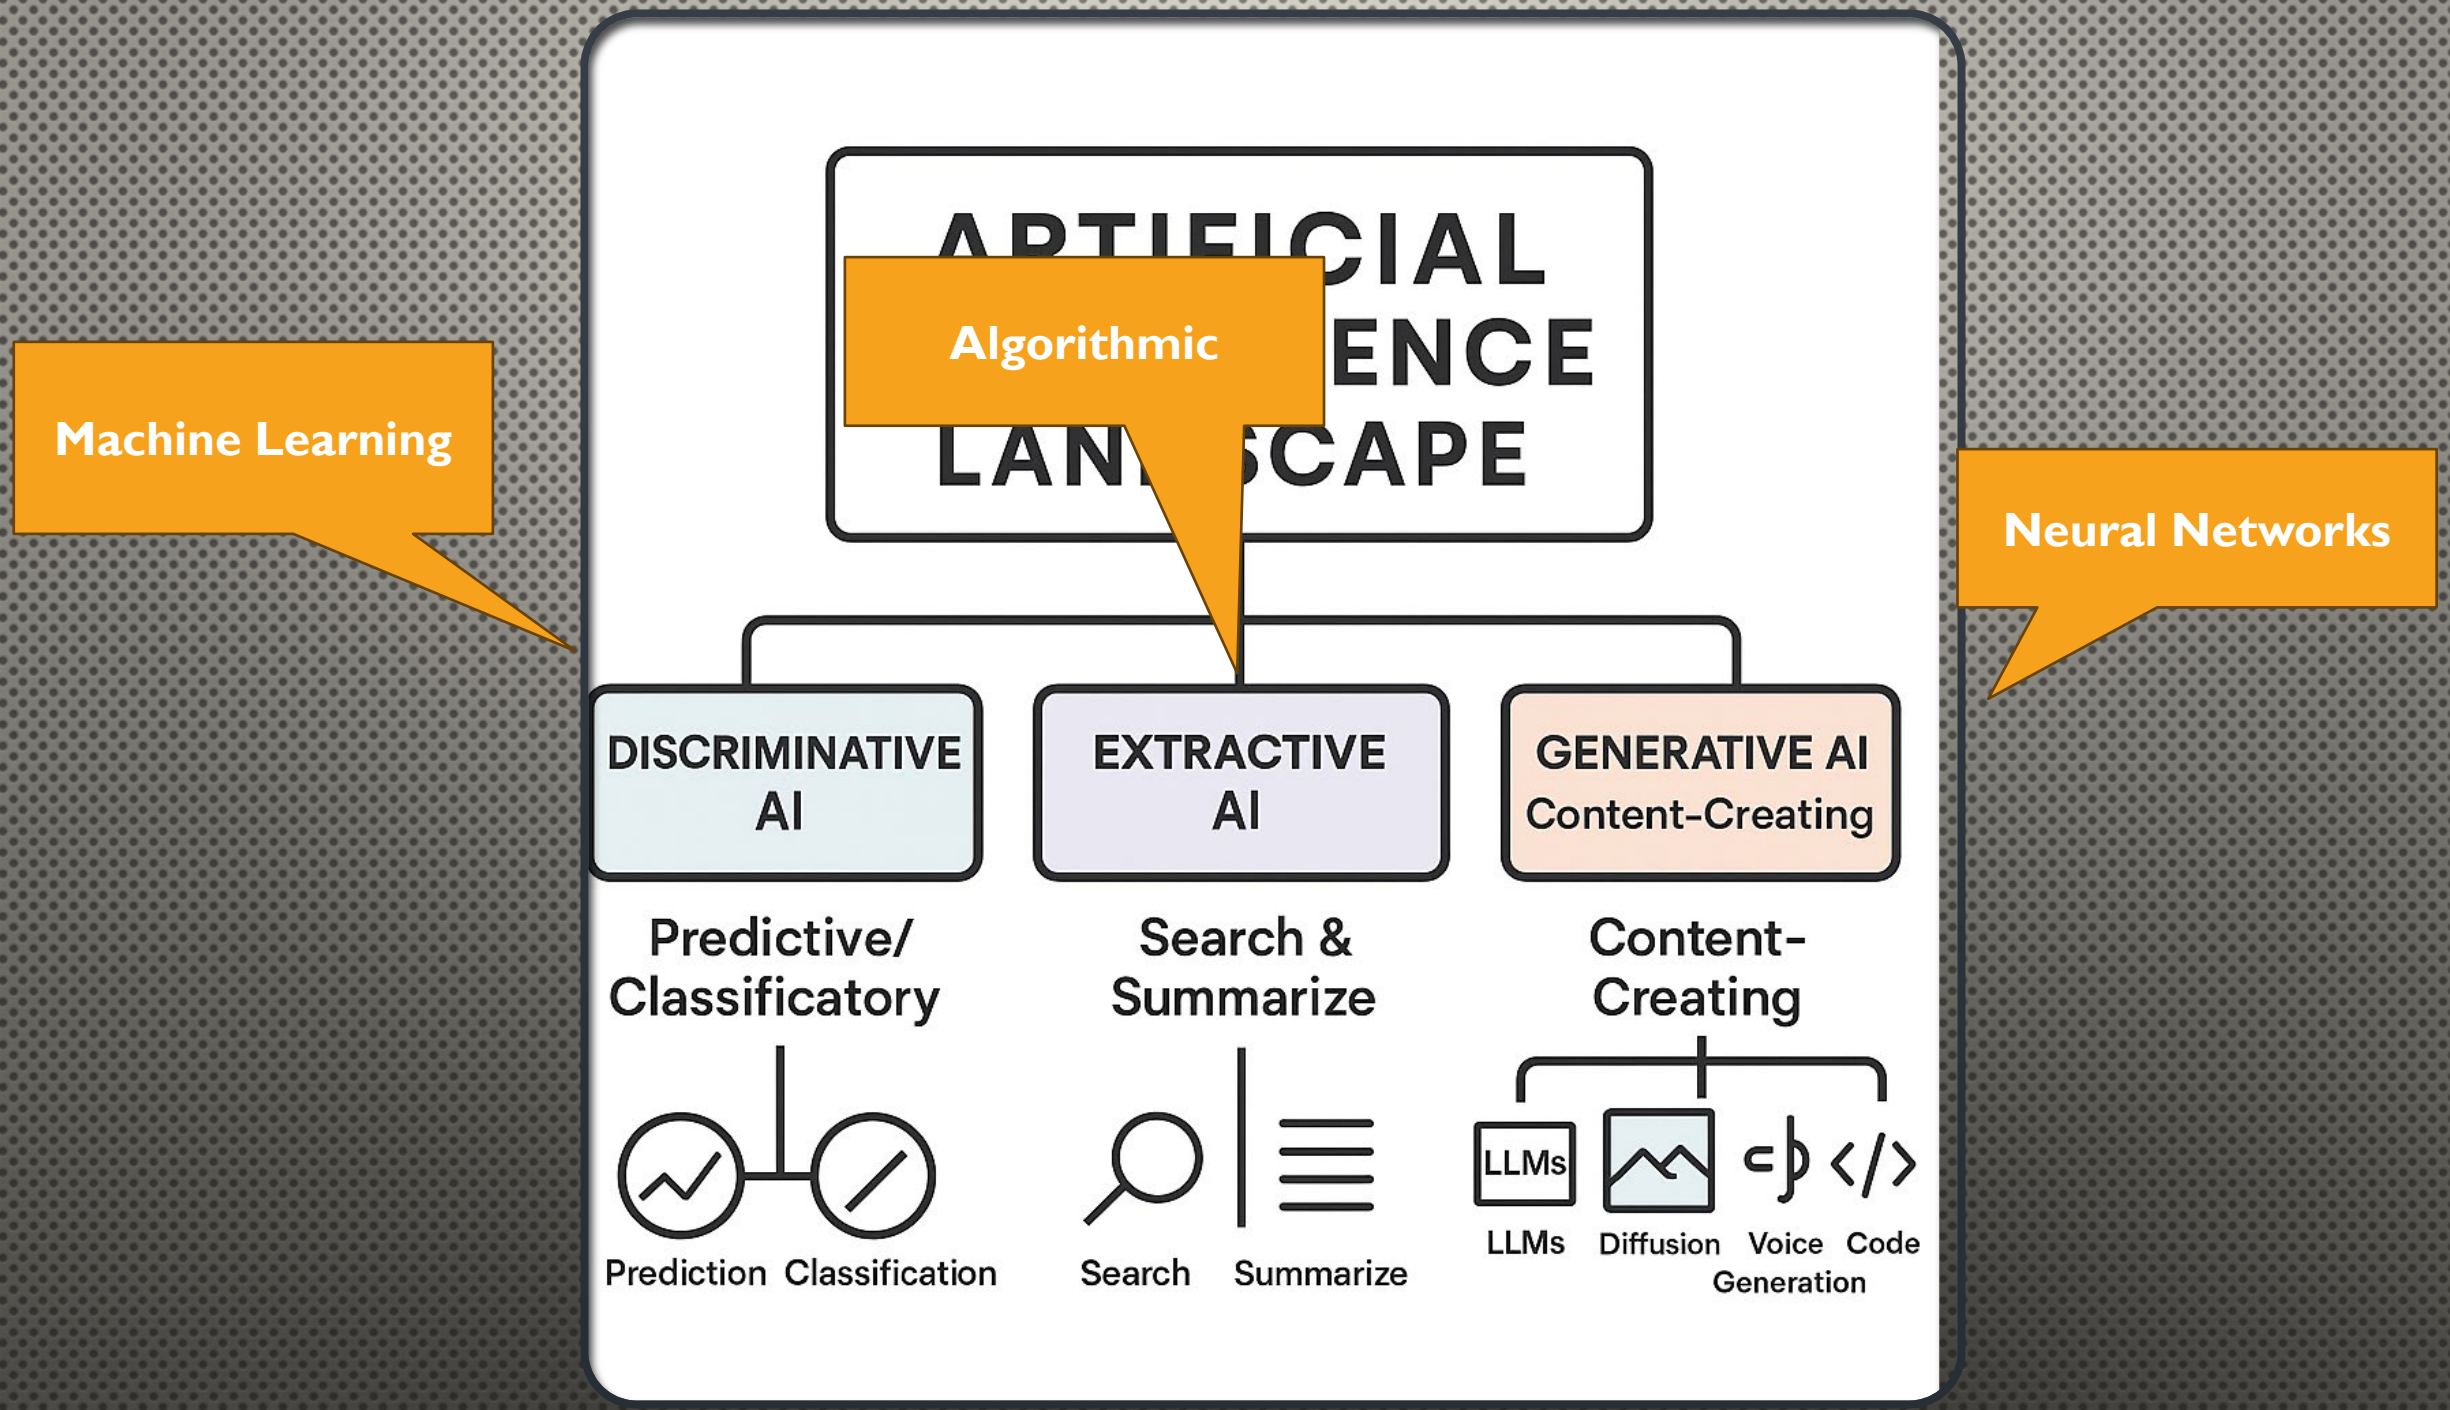
Machine Learning
Algorithmic
Neural Networks
DISCRIMINATIVE AI
EXTRACTIVE AI
GENERATIVE AI
Content-Creating
Predictive/ Classificatory
Search & Summarize
Content-Creating
Prediction
Classification
Search
Summarize
LLMs
LLMs
Diffusion
Voice Generation
Code

MUCH OF THE "MAGIC" BEHIND GENAI DOES NOT COME FROM WHAT WE THINK OF AS COMPUTER PROGRAMMING

Classical (Rule-Based) Software

Developers write explicit instructions ("if X, do Y") in languages like Python or Java.

Works beautifully for calculators or having a computer replicate our knowledge

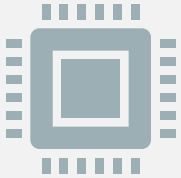

Data-Driven Machine Learning

You still write code, but most of your "logic" lives in the data.

Training: feed the model billions of text tokens; it learns statistical patterns rather than hard rules.

Deep Neural Networks

Layers of simple mathematical "neurons" transform raw tokens into higher-level representations.

Millions or billions of parameters are tuned via gradient-descent, not manually coded.

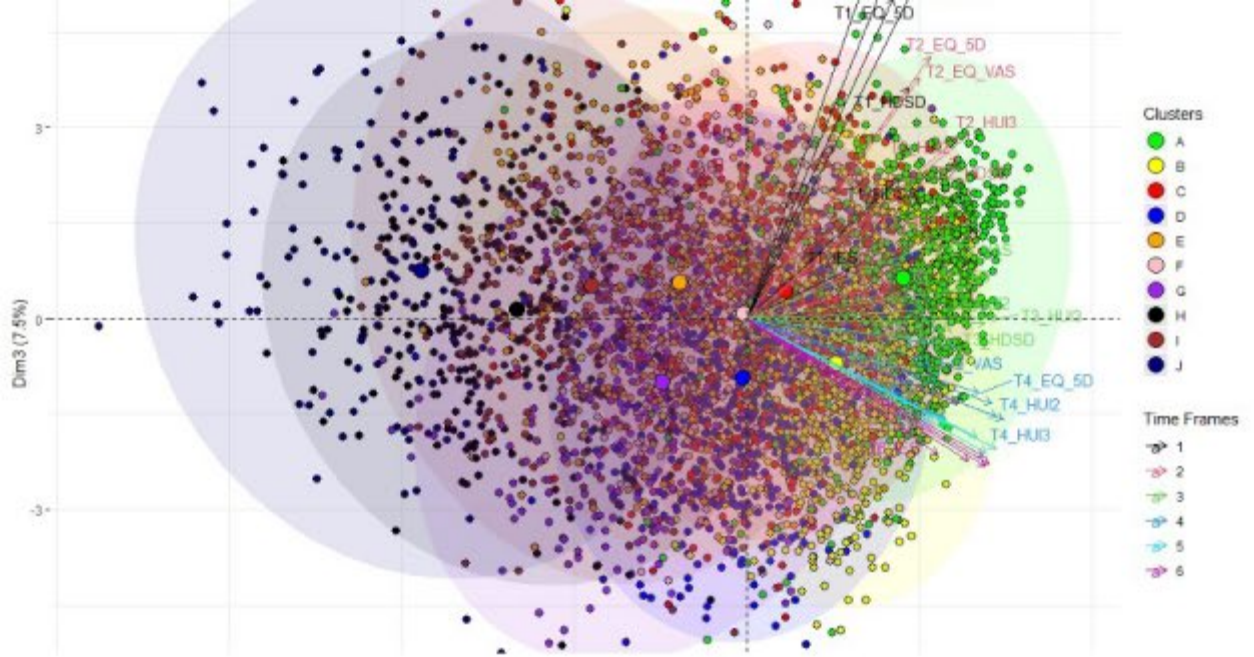

K-MEANS VS K-NN

UNDERSTANDING THE KEY DIFFERENCES

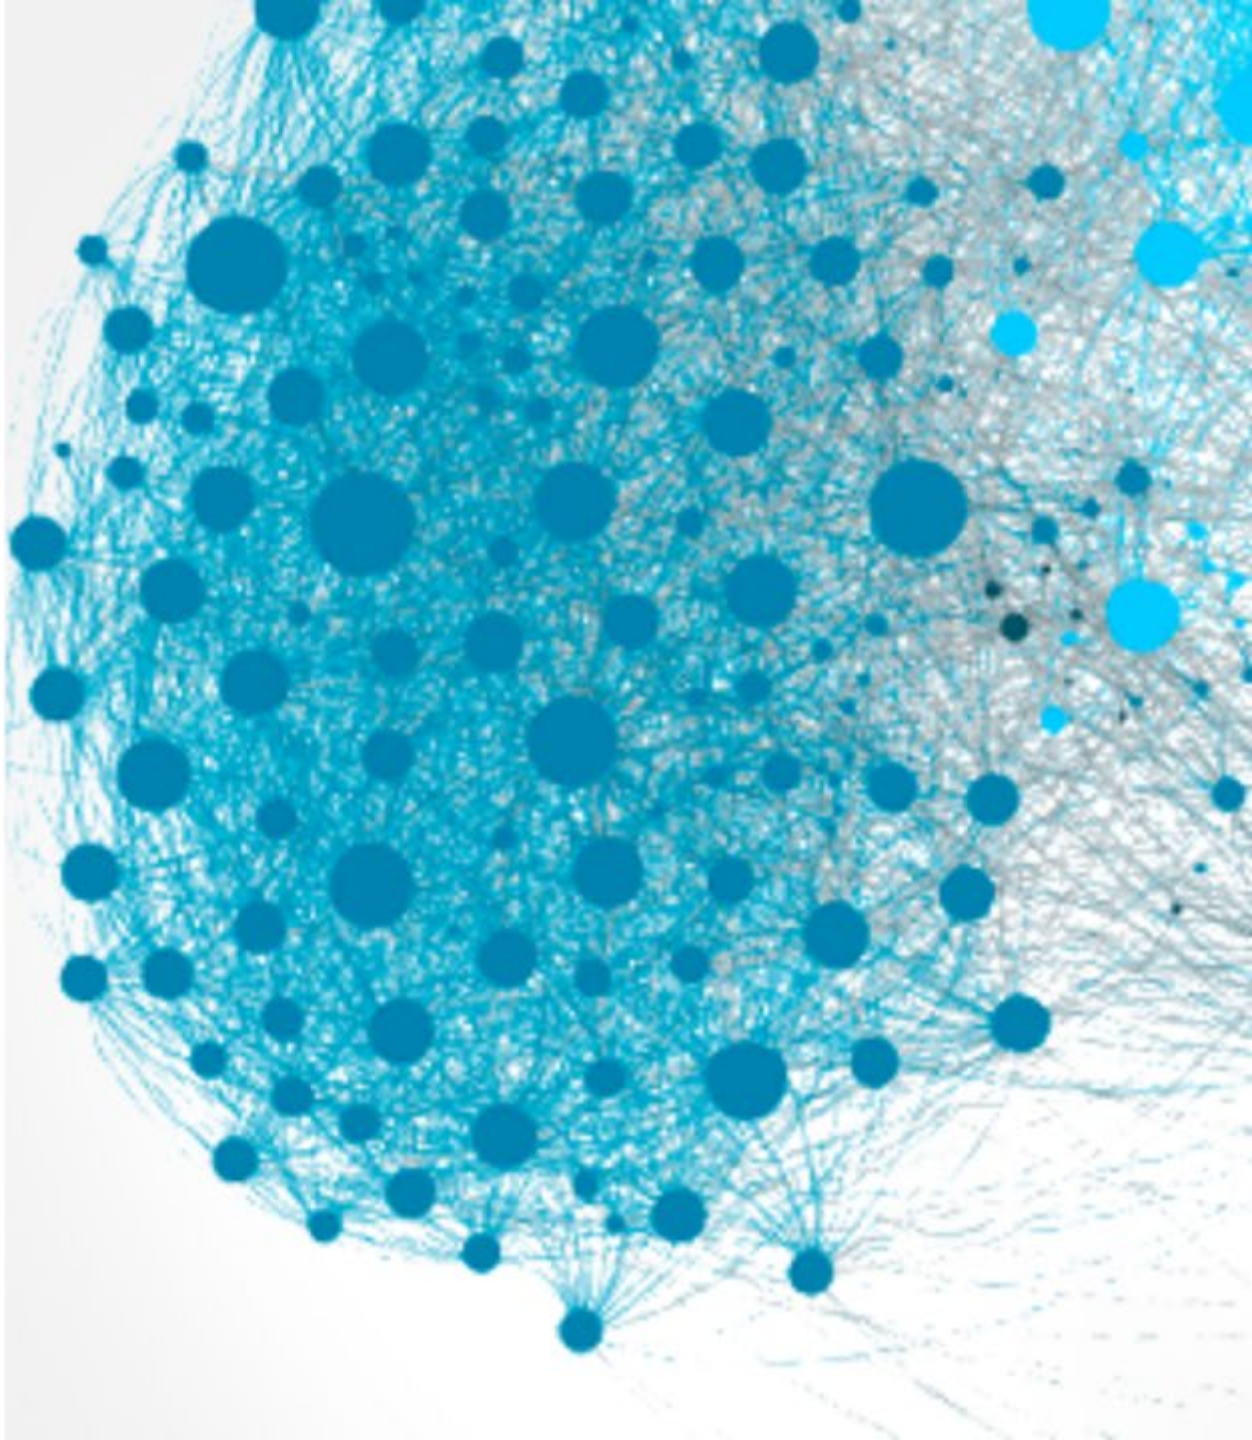

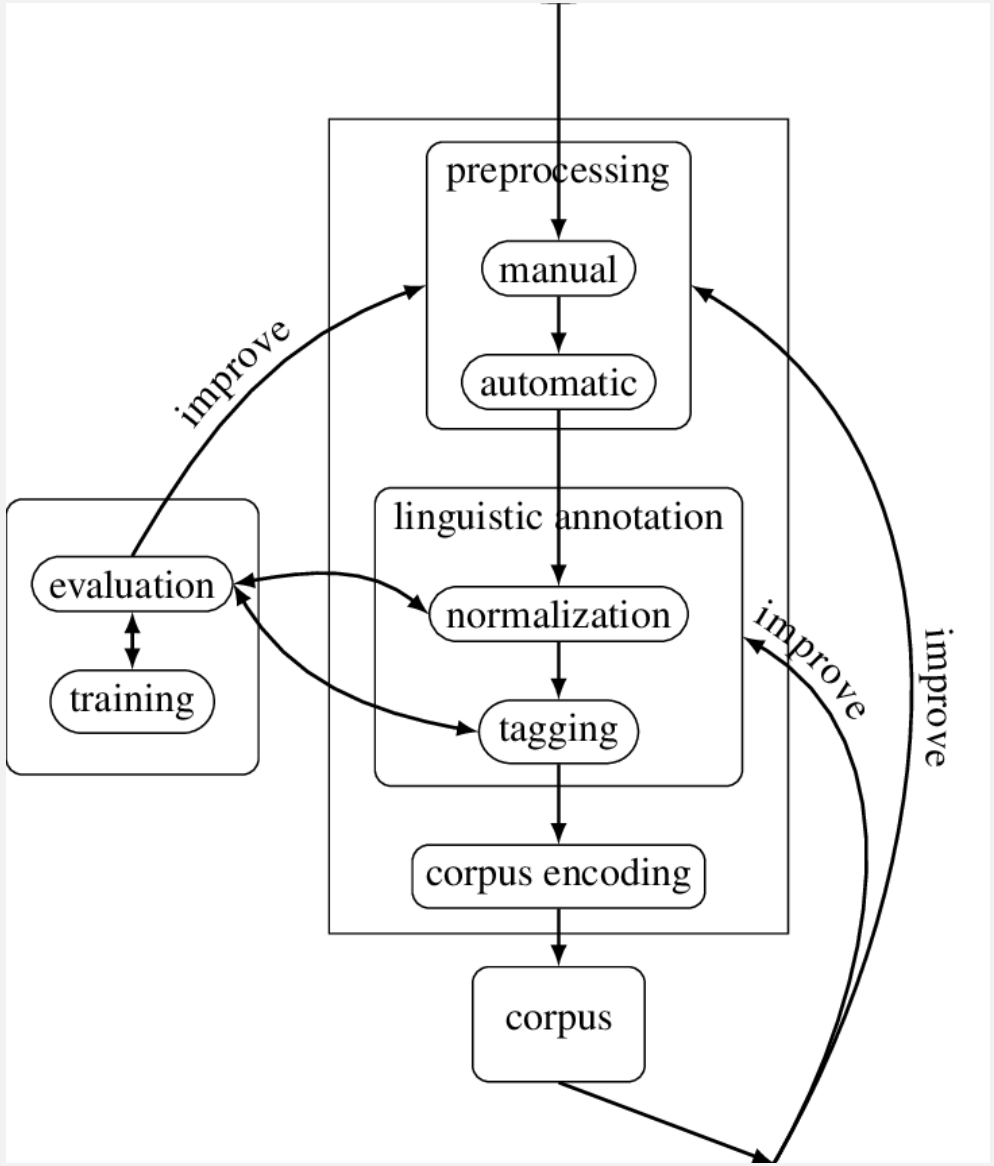

- Training data is treated differently than reinforcement learning
- Training data – a model for how to respond
- Reinforcement – did it get it right?
- Prompts: are a bit like a *run* command and a bit like reinforcement.

We “prompt” GenAI to do its thing – the “magic” happens behind the scenes and then it communicates its results with us.

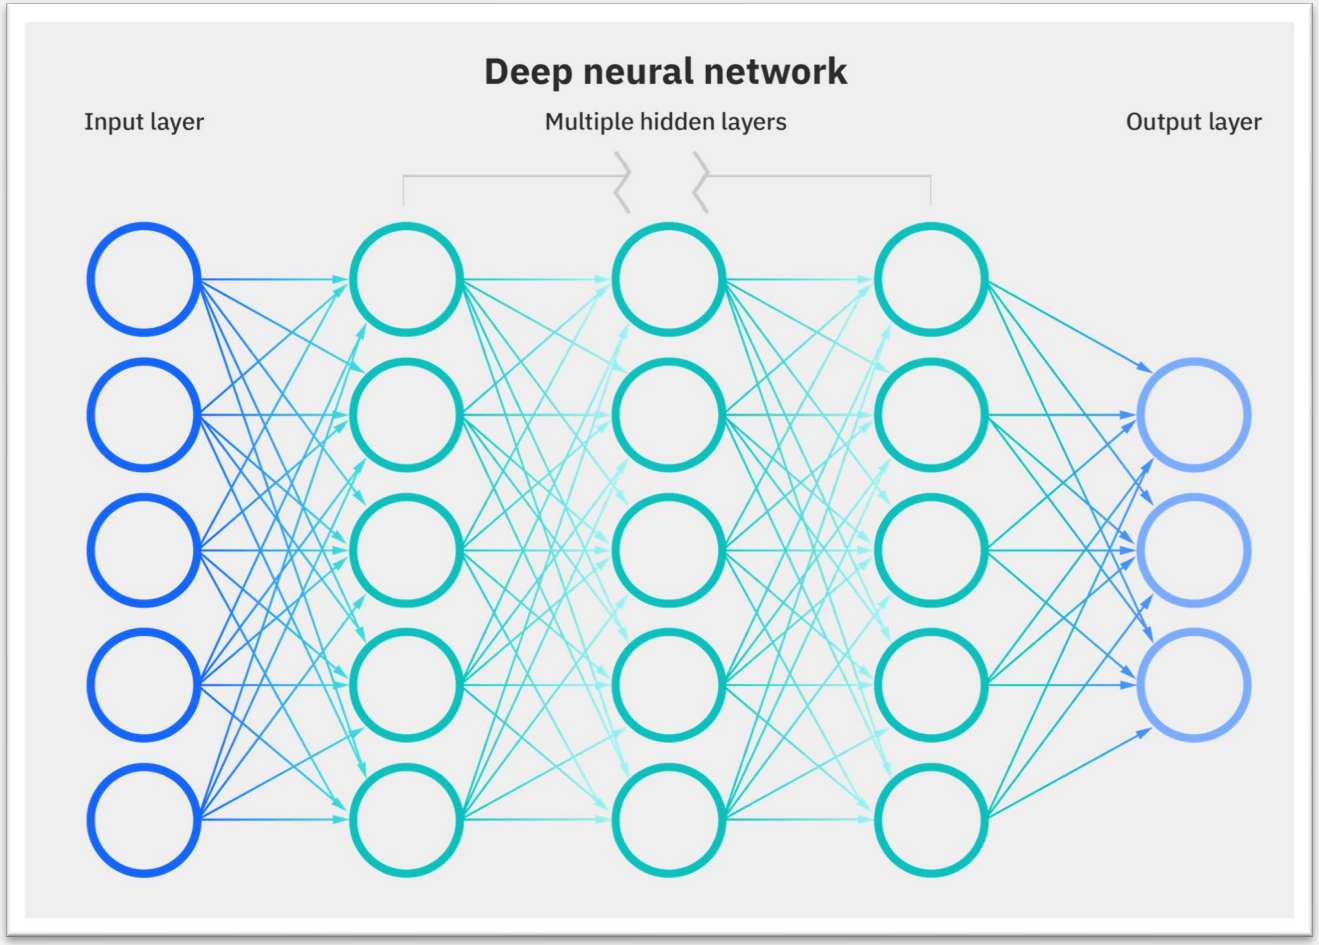

The OUTPUT LAYER is unique. It is not “extracted” or “retrieved” from the INPUT LAYER.

LARGE LANGUAGE MODELS ARE AS MUCH HUMANITIES AS THEY ARE COMPUTER SCIENCE

COMPUTATIONAL LINGUISTICS

NATURAL LANGUAGE PROCESSING

LLMS ARE STOCHASTIC PARROTS

e.g.,

next

ATTENTION IS…

- …all you need. (Vaswani et al., 2017 - Cited by 180,227)
- …the purest form of generosity. (Simone Weil)
- …the beginning of devotion. (Mary Oliver)
- …expensive to pay. (Song by Jhene Aiko)

HOW ARE LLMS TRAINED?

N-GRAM

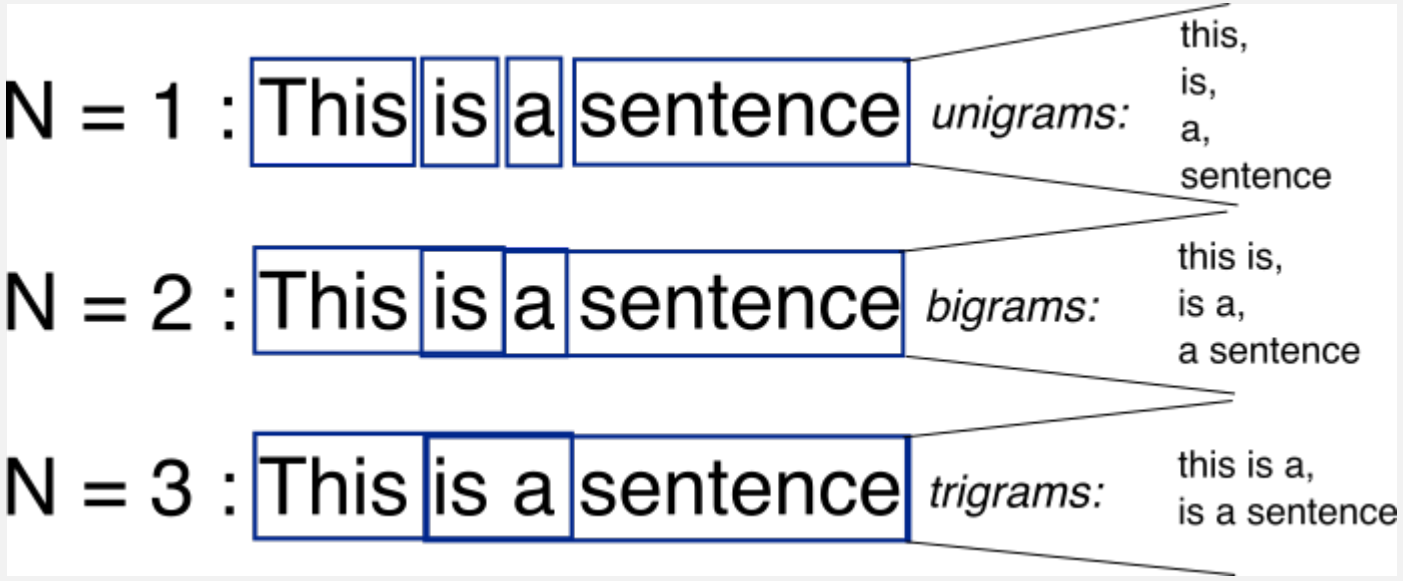

CONCORDANCE TABLE

TRAINING DATA PROCESSING & INFRASTRUCTURE

A. Preproce

- Sharding and ch
- Batching sequer
- Token-to-ID enc
- Mixed precision

B. Infrastru

- Distributed train
  model parallelis
- Hardware: GPUs (NVIDIA A100), TPUs
- Frameworks: PyTorch, TensorFlow, JAX, DeepSpeed, Megatron-LM

And on…

essing
g & Evalutic

Tasks

words or

Tasks

ization
e generation
nsfer

MACHINE LEARNING & DEEP LEARNING PIPELINE

A. Model Architecture Design

ecture (Vaswani et al. 2011)
ention

Variants:
- Decoder-only (GPT)
- Encoder-decoder (T5)
- Encoder-only (BERT)

). Encoder-decoder (T5)
T)

aining Process

ion (causal LM) or masked
LM)
s-entropy loss
ient descent (sudam or AdamW)
pout, weight decay, label

2. Fine-tuning
- Task-specitic adjustment on smaller, high-quality

3. Reinforcement Learning from Human Feedback
- Reward model trained on human preferences,
- Policy optimization (eg. PPO) to adjust gen behavio

C. Evaluation Metrics
- Perplexity
- BLEU, ROUGE (for summarization)
- Truthfulness (eg. TruthfulQA)
- Hallucination and factuality metrics

AT BEST, ASKING GENERATIVE AI IS LIKE ASKING YOUR AUNT BETTY FOR RESEARCH

- It's not necessarily going to be wrong, but it's not a trustworthy source.
- It will probably get the common knowledge parts right.

AT WORST, IT IS LIKE ASKING YOUR 4-YEAR-OLD NEPHEW.

- It will try to please you.
- If it doesn't know the answer, it will make one up.

Law Profs' Gen AI Sandbox

Welcome to the GenAI Sandbox

Empowering & Inspiring Law Professors in the Age of Generative AI

This sandbox is your resource for exploring how you might integrate artificial intelligence into law school instruction. Whether you've never tried generative AI or you're already building your own bots, our sandbox provides tailored resources to help you learn to use AI to enhance your teaching and support your students' success.

Discover Your AI Persona: Take the Disposition Survey