

GENERATIVE AI SYSTEMS

JAKUB M. TOMCZAK |

Group Leader, Associate Professor
Generative AI Group & Dept. M&CS

TU/e

**EINDHOVEN
UNIVERSITY OF
TECHNOLOGY**

Jakub M. Tomczak, Ph.D.

Associate Professor at TU/e

Group Leader of the Generative AI group

Founder Amsterdam AI Solutions

+15y experience in ML/AI and GenAI

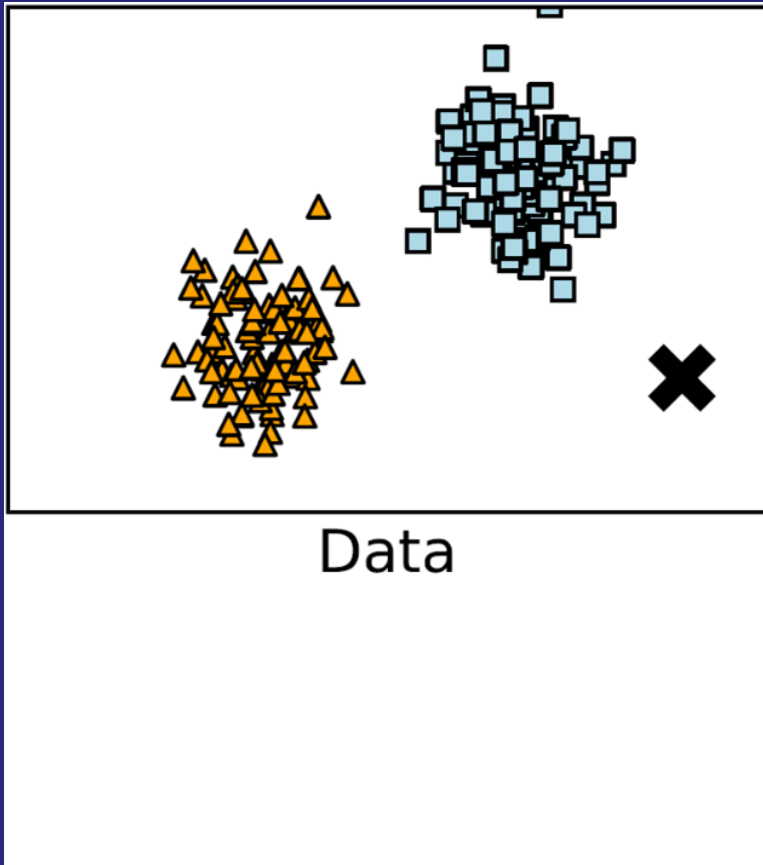
The author of “Deep Generative Modeling”



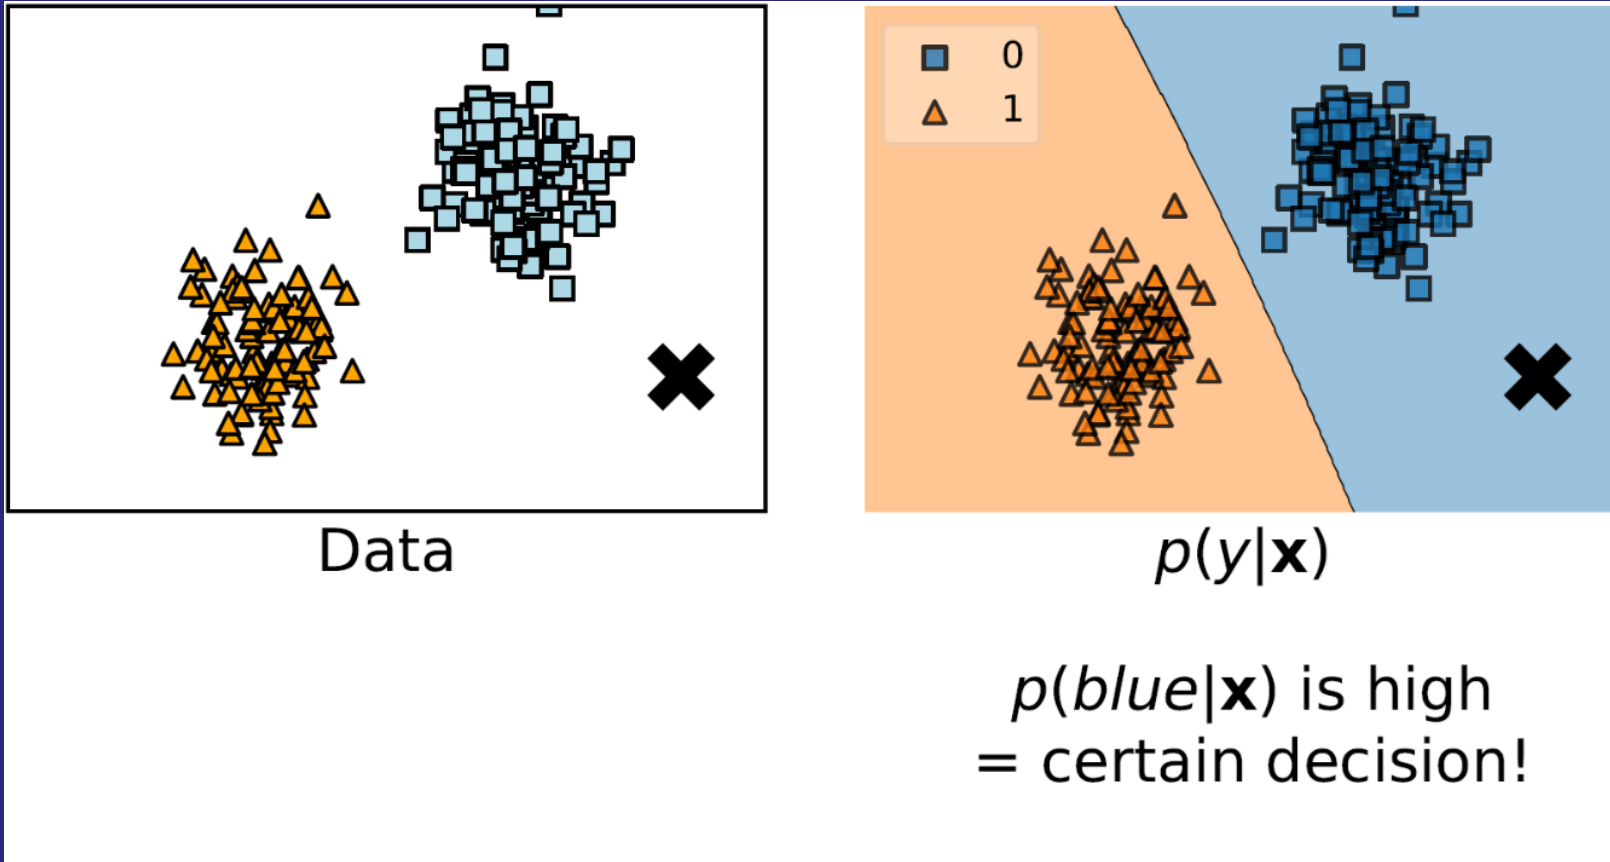
Generativ/e

**Amsterdam
AI Solutions**

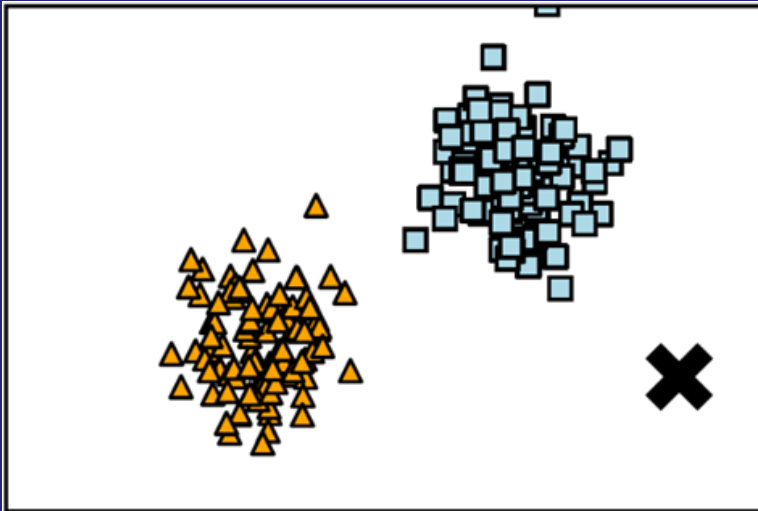
Modeling: Discriminative vs. Generative



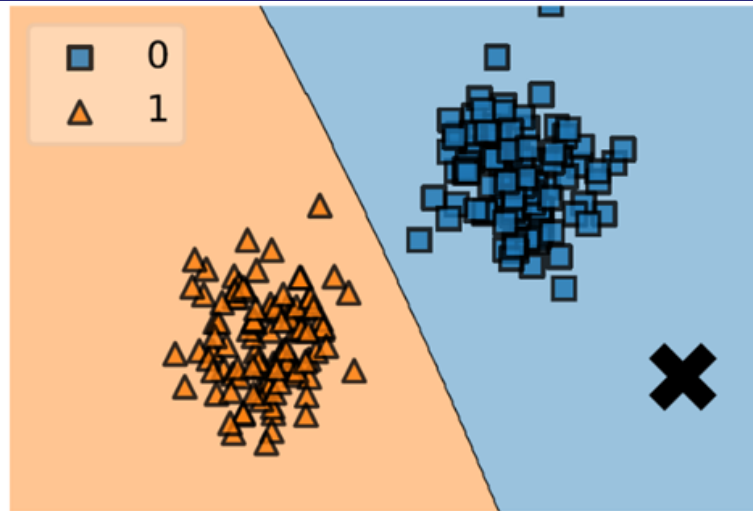
Modeling: Discriminative vs. Generative



Modeling: Discriminative vs. Generative

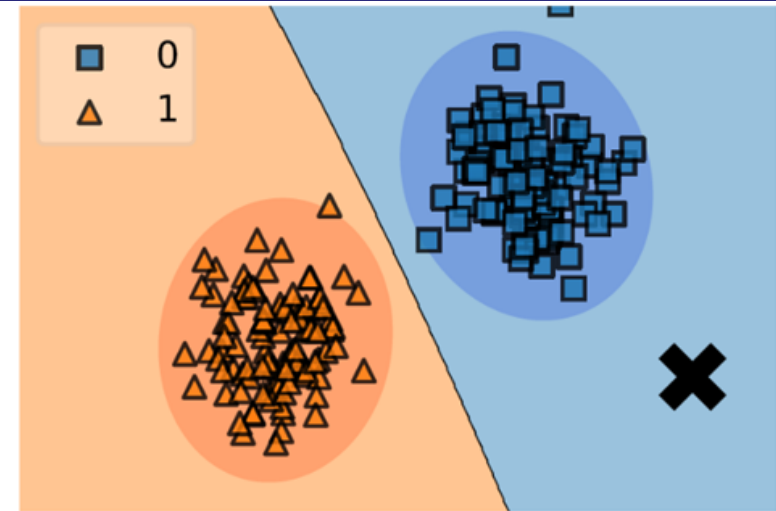


Data



$p(y|\mathbf{x})$

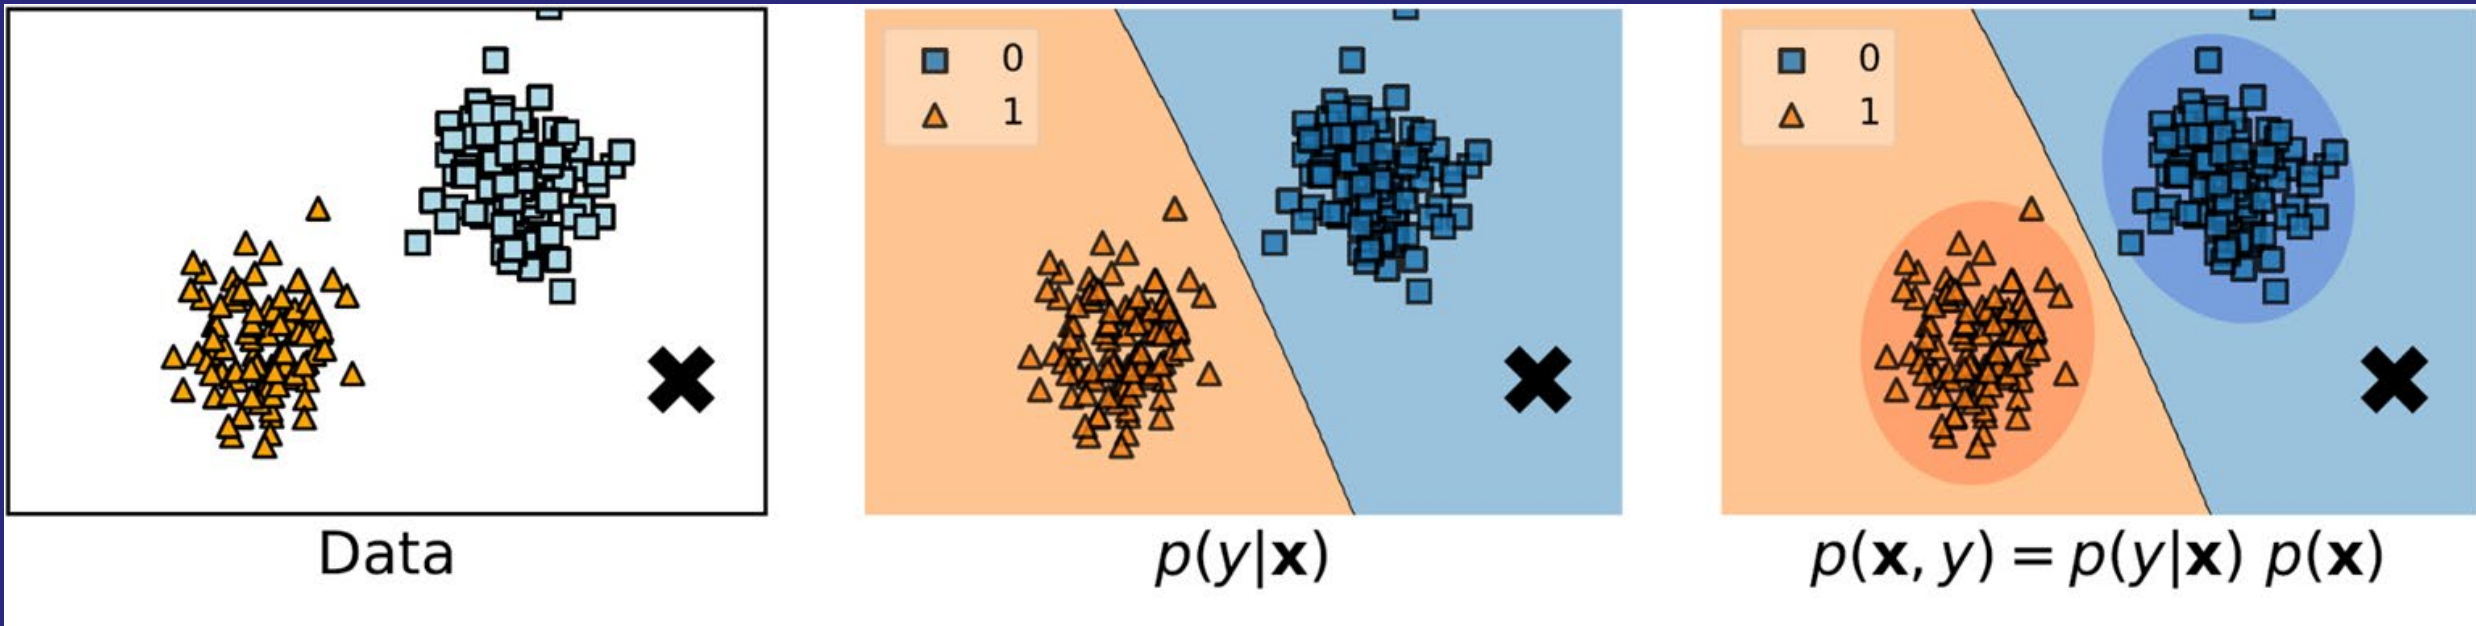
$p(\text{blue}|\mathbf{x})$ is high
= certain decision!



$p(\mathbf{x}, y) = p(y|\mathbf{x}) p(\mathbf{x})$

$p(\text{blue}|\mathbf{x})$ is high
and $p(\mathbf{x})$ is low
= uncertain decision!

Modeling: Discriminative vs. Generative



Knowing the generative process allows us to understand phenomena
and
synthesize (a.k.a. generate) new data.

Generative AI

Probabilistic modeling

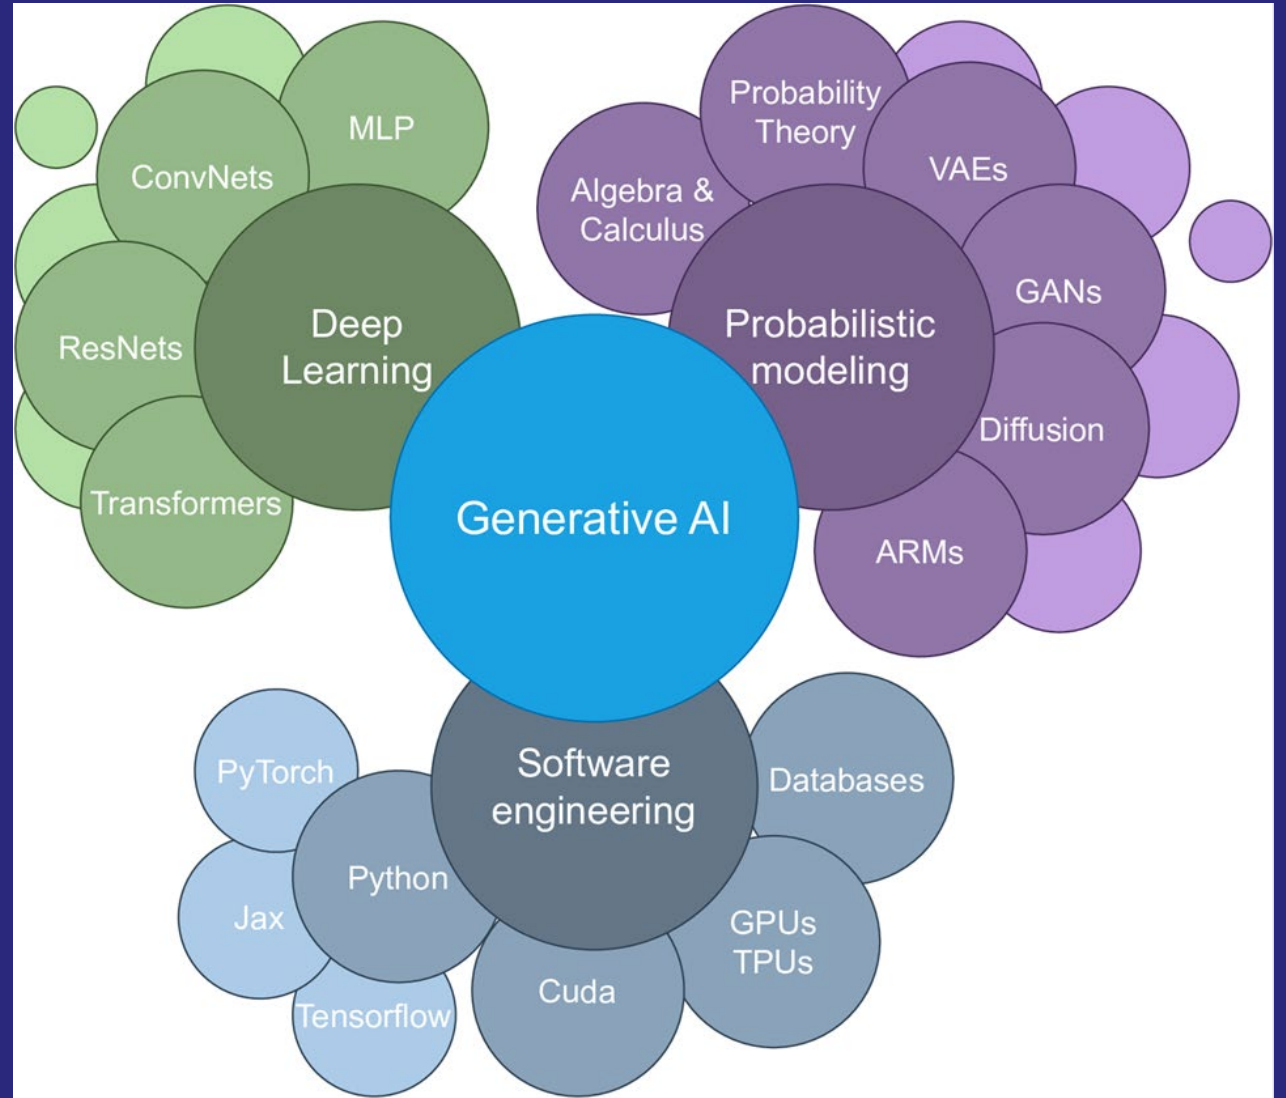
(principles of building models)

Deep Learning

(parameterizations of distributions)

Software engineering

(effective and efficient implementation of Generative AI)



Generative AI

Probabilistic modeling

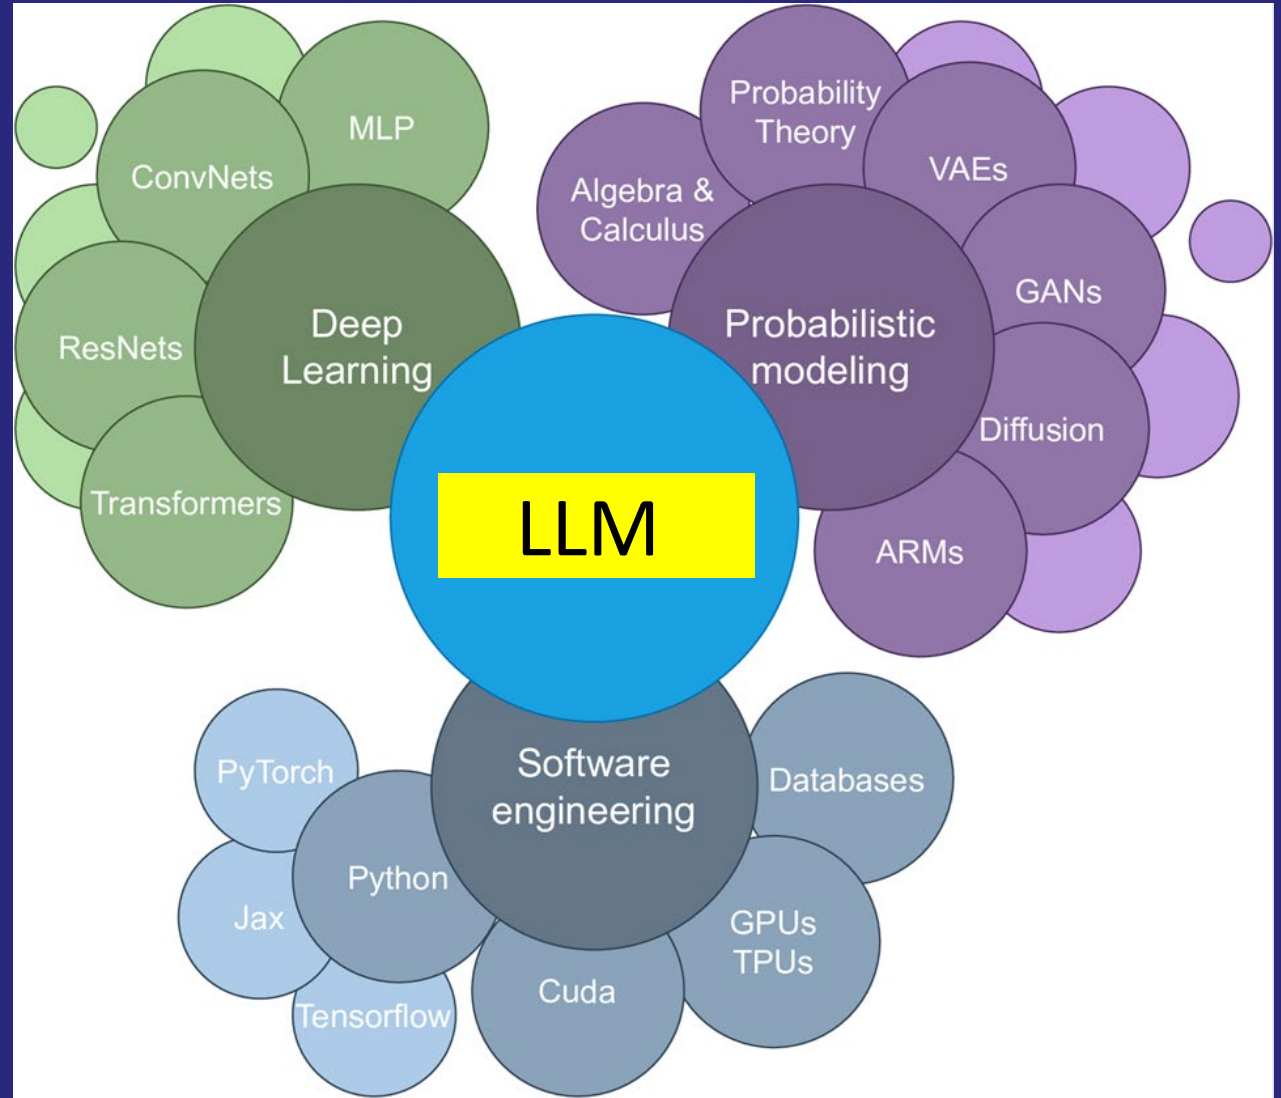
(principles of building models)

Deep Learning

(parameterizations of distributions)

Software engineering

(effective and efficient implementation of Generative AI)



Generative AI

Probabilistic modeling

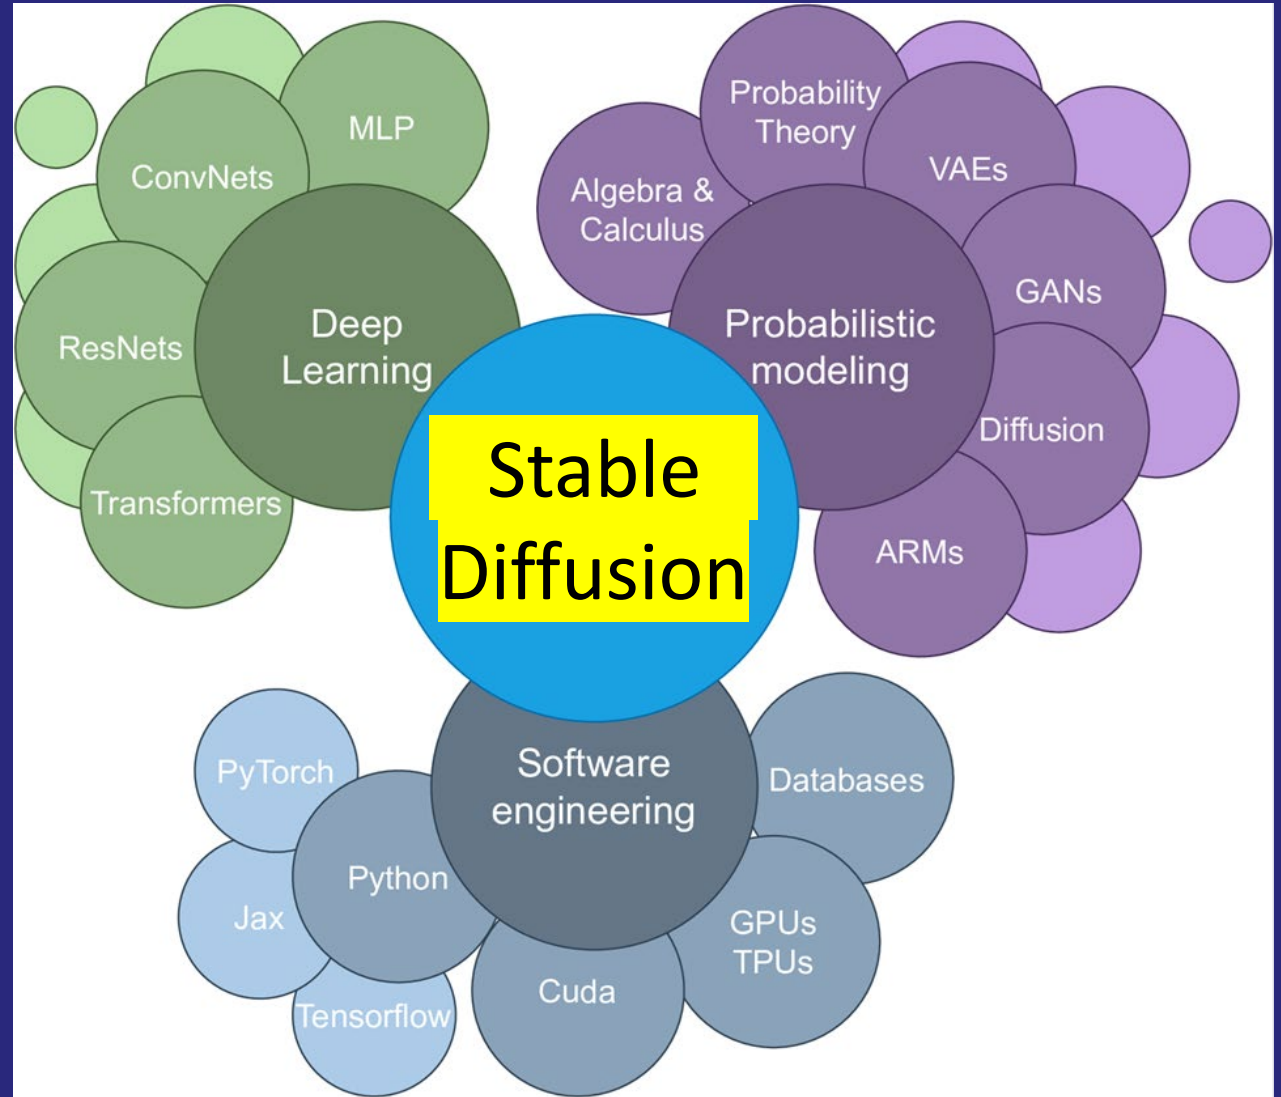
(principles of building models)

Deep Learning

(parameterizations of distributions)

Software engineering

(effective and efficient implementation of Generative AI)



How to build Generative AI?

Why Deep Generative Modeling?

Jakub M. Tomczak
Pages 1–13

NEW!

Probabilistic Modeling: From Mixture Models to Probabilistic Circuits

Jakub M. Tomczak
Pages 15–36

Updated!

Autoregressive Models

Jakub M. Tomczak
Pages 37–62

Flow-Based Models

Jakub M. Tomczak
Pages 63–92

Latent Variable Models

Jakub M. Tomczak
Pages 93–167

Updated!

Hybrid Modeling

Jakub M. Tomczak
Pages 169–181

Energy-Based Models

Jakub M. Tomczak
Pages 183–199

Generative Adversarial Networks

Jakub M. Tomczak
Pages 201–215

NEW!

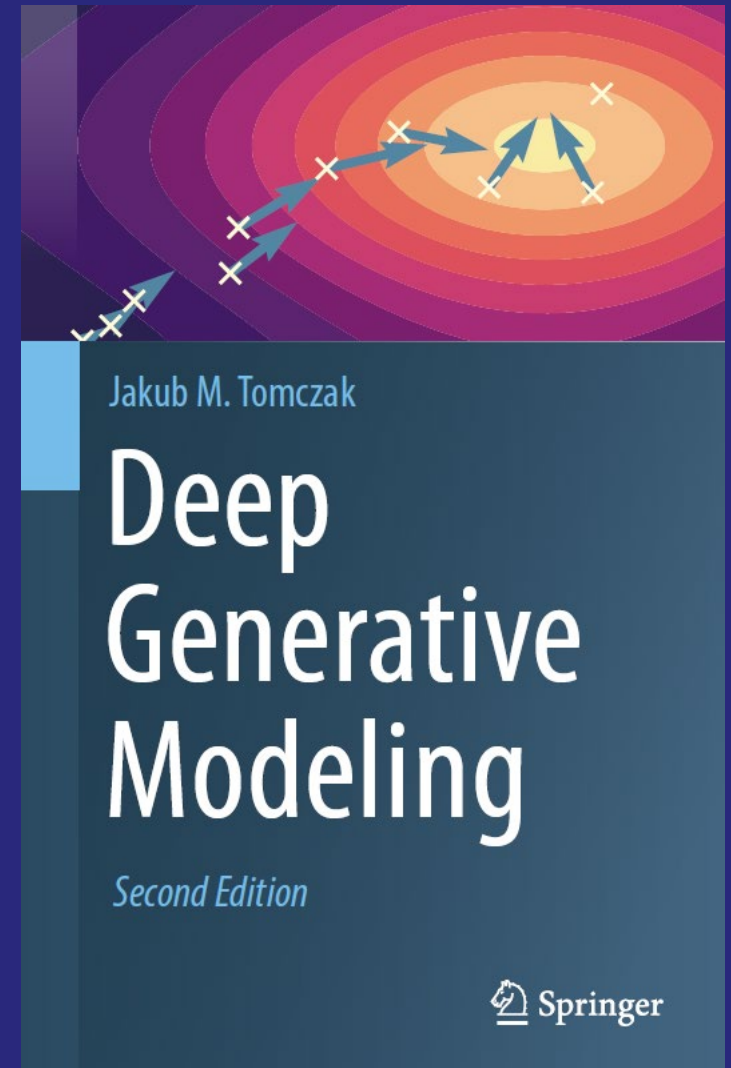
Score-Based Generative Models

Jakub M. Tomczak
Pages 217–257

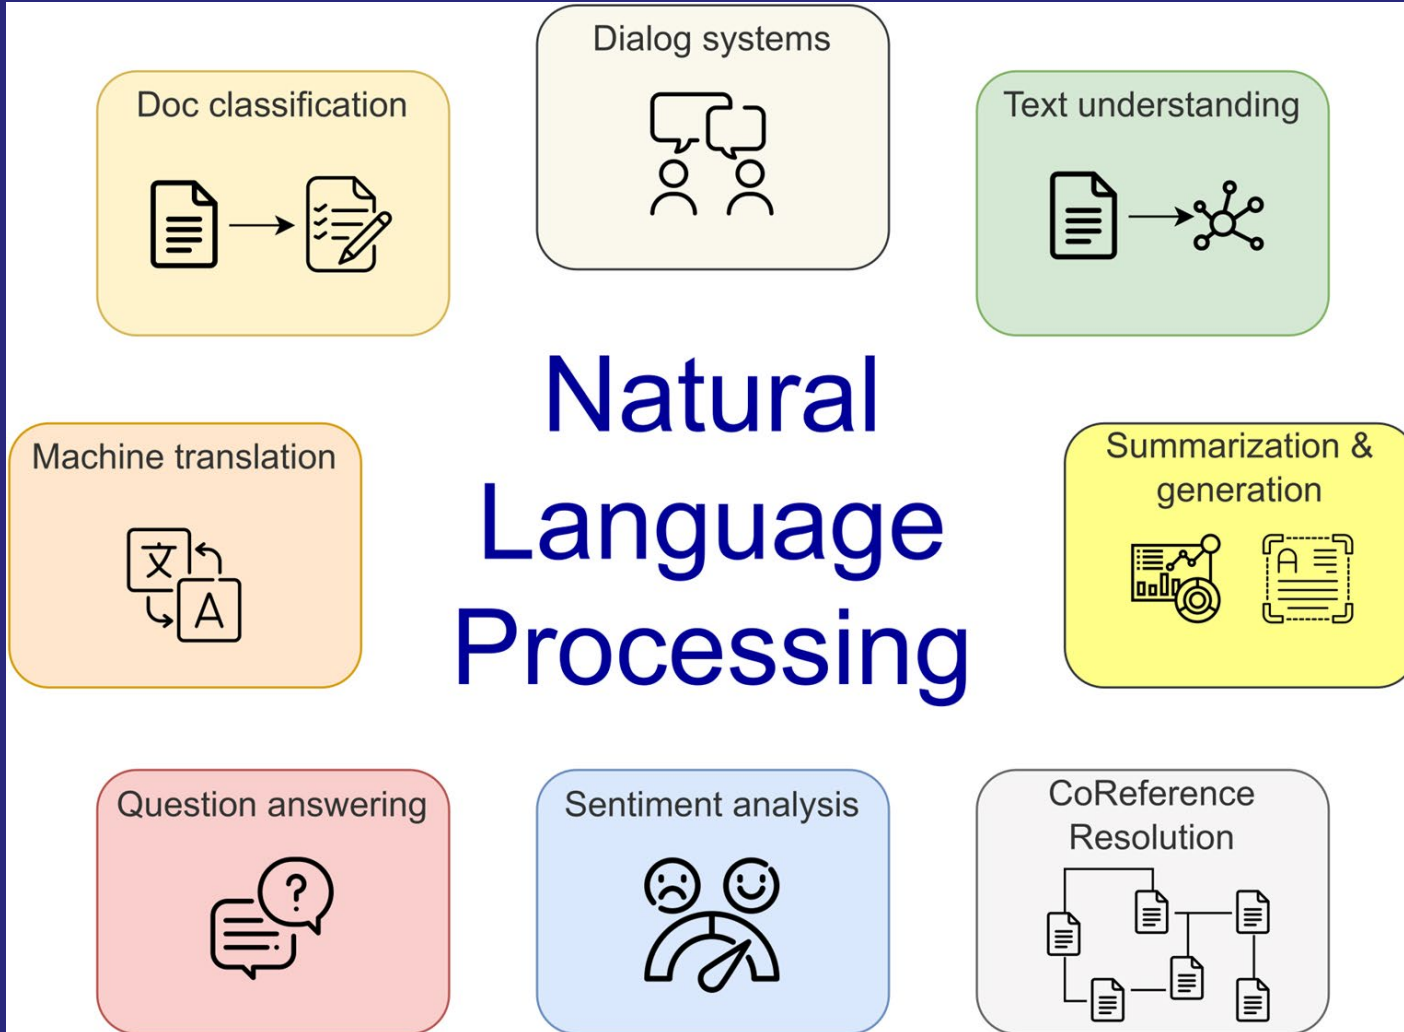
From Large Language Models to Generative AI Systems

Jakub M. Tomczak
Pages 277–302

NEW!



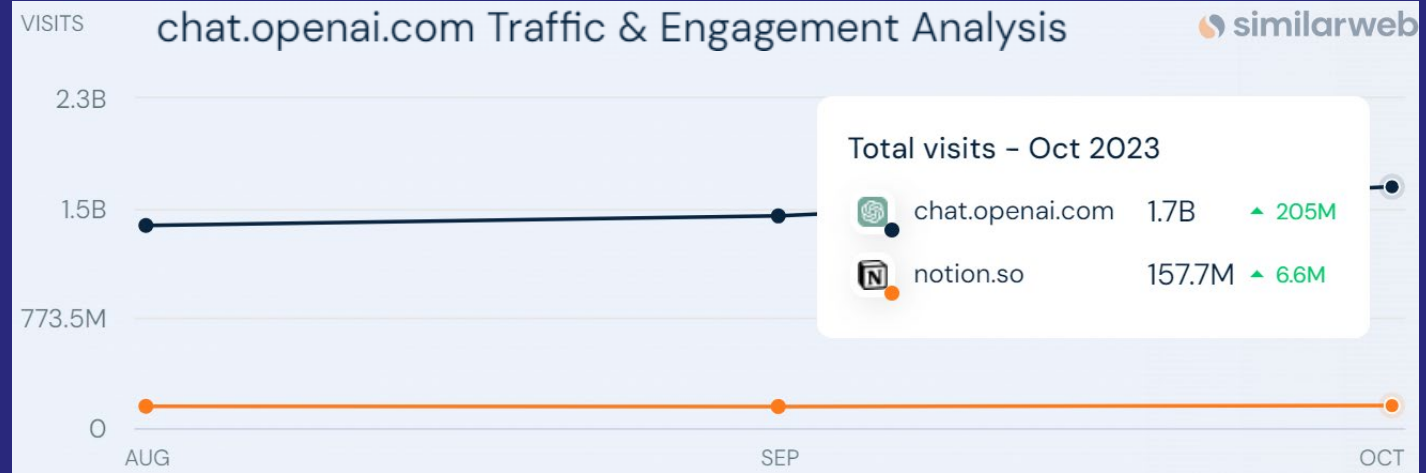
Large Language Models





OpenAI Chat GPT

Launched November 30, 2022



Published 11:06 AM EDT, Wed October 23, 2024

ChatGPT can be tricked into telling people how to commit crimes, a tech firm finds



Newsroom

01/08/24 | Wolfsburg | Models | Press Release

World premiere at CES: Volkswagen integrates ChatGPT into its vehicles



These lawyers used ChatGPT to save time. They got fired and fined.

Artificial intelligence is changing how law is practiced, but not always for the better

By Pranshu Verma and Will Oremus

Updated November 16, 2023 at 10:39 a.m. EST | Published November 16, 2023 at 6:00 a.m. EST



Dutch ChatGPT

ChatGPT Online



Hallo mens, ik ben een door GPT aangedreven AI-chatbot. Vraag mij alles!



Hoi! Hoe gaat het met jij?



Hallo! Het gaat goed, bedankt voor het vragen. Hoe gaat het met jou?



Prima! IK maak een presentatie voor studenten in Eindhoven. Weet jij wat ik presenteren moeten zou?



Wat leuk dat je een presentatie gaat geven! Misschien is het handig om te presenteren over een onderwerp waar je gepassioneerd over bent of waar je veel kennis over hebt. Denk bijvoorbeeld aan recente ontwikkelingen in jouw vakgebied, interessante onderzoeken of inspirerende projecten. Veel succes met je presentatie!

Type een bericht





	Claude 3.5 Sonnet (new)	Claude 3.5 Haiku	Claude 3.5 Sonnet	GPT-4o*	GPT-4o mini*	Gemini 1.5 Pro	Gemini 1.5 Flash
Graduate level reasoning <i>GPQA (Diamond)</i>	65.0% 0-shot CoT	41.6% 0-shot CoT	59.4% 0-shot CoT	53.6% 0-shot CoT	40.2% 0-shot CoT	59.1% 0-shot CoT	51.0% 0-shot CoT
Undergraduate level knowledge <i>MMLU Pro</i>	78.0% 0-shot CoT	65.0% 0-shot CoT	75.1% 0-shot CoT	—	—	75.8% 0-shot CoT	67.3% 0-shot CoT
Code <i>HumanEval</i>	93.7% 0-shot	88.1% 0-shot	92.0% 0-shot	90.2% 0-shot	87.2% 0-shot	—	—
Math problem-solving <i>MATH</i>	78.3% 0-shot CoT	69.2% 0-shot CoT	71.1% 0-shot CoT	76.6% 0-shot CoT	70.2% 0-shot CoT	86.5% 4-shot CoT	77.9% 4-shot CoT
High school math competition <i>AIME 2024</i>	16.0% 0-shot CoT	5.3% 0-shot CoT	9.6% 0-shot CoT	9.3% 0-shot CoT	—	—	—
Visual Q/A <i>MMMU</i>	70.4% 0-shot CoT	—	68.3% 0-shot CoT	69.1% 0-shot CoT	59.4% 0-shot CoT	65.9% 0-shot CoT	62.3% 0-shot CoT
Agentic coding <i>SWE-bench Verified</i>	49.0%	40.6%	33.4%	—	—	—	—
Agentic tool use <i>TAU-bench</i>	Retail 69.2% Airline 46.0%	Retail 51.0% Airline 22.8%	Retail 62.6% Airline 36.0%	—	—	—	—

* Our evaluation tables exclude OpenAI's o1 model family as they depend on extensive pre-response computation time, unlike typical models. This fundamental difference makes performance comparisons difficult.



GitHub Copilot

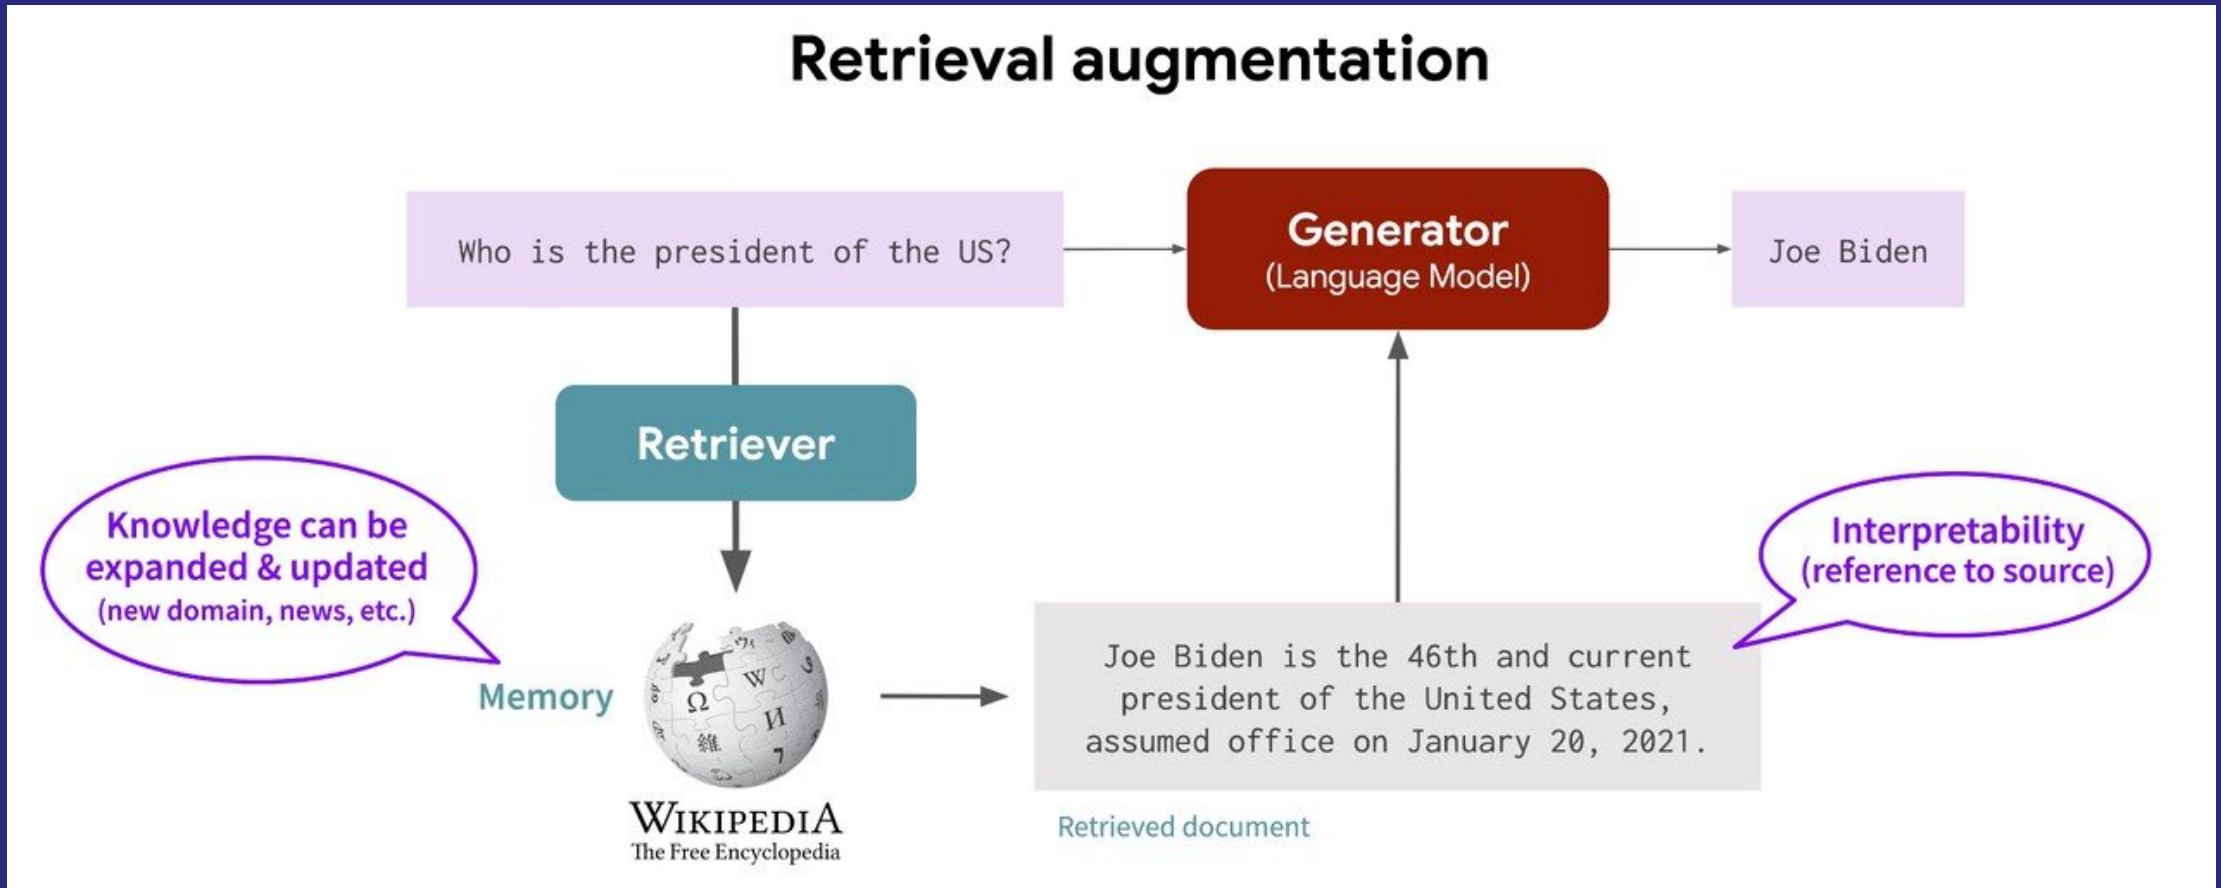
```
runtime.go course.rb time.js IsPrimeTest.java
1 package main
2
3 type Run struct {
4     Time int // in milliseconds
5     Results string
6     Failed bool
7 }
8
9 // Get |
10
11
12
13
14
15
16
17
18
19
20
21
22
23
```

More than
natural language

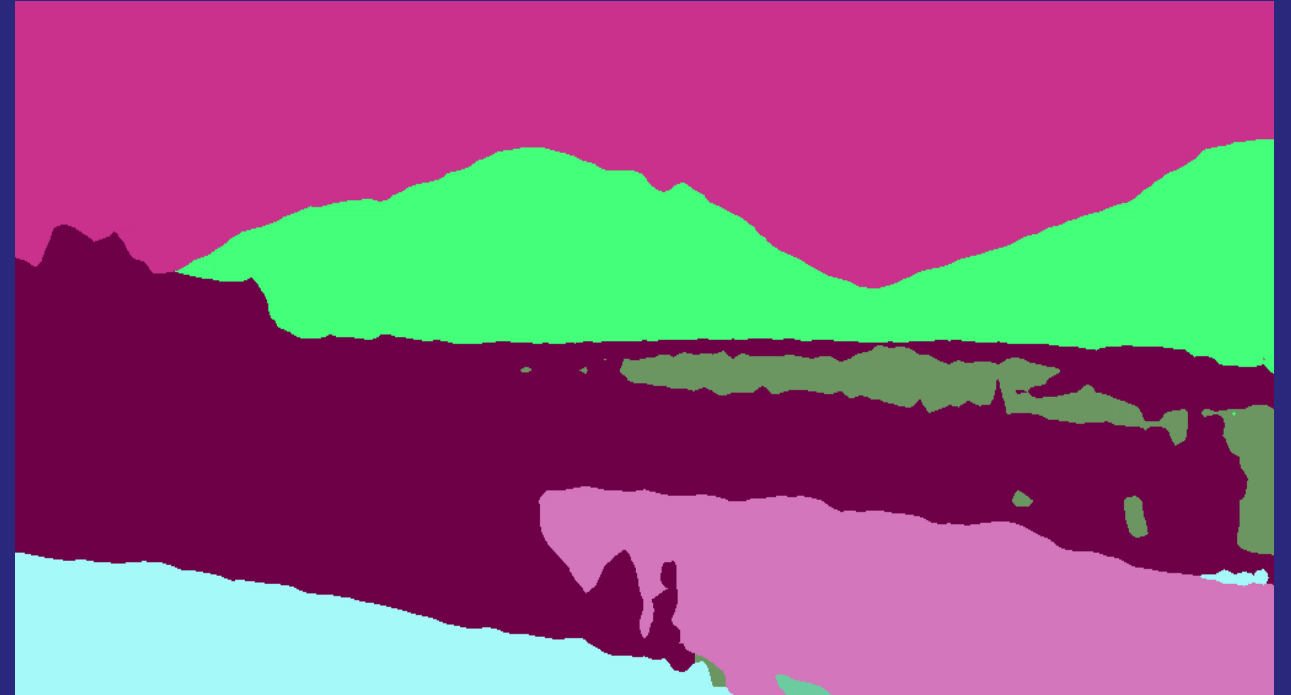
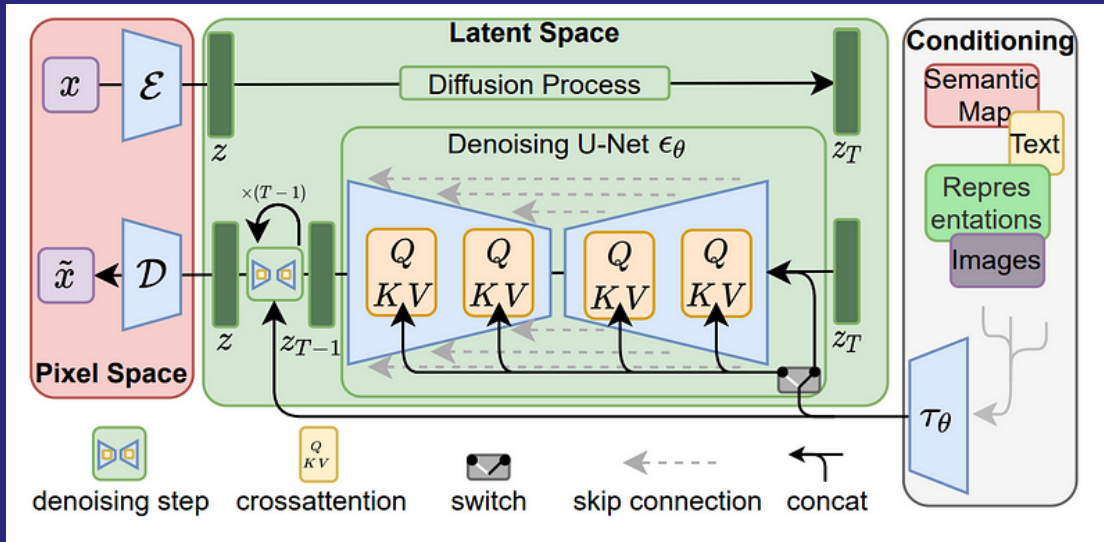


Going beyond LLMs: Generative AI Systems

RAG: Retrieval-augmented Generation



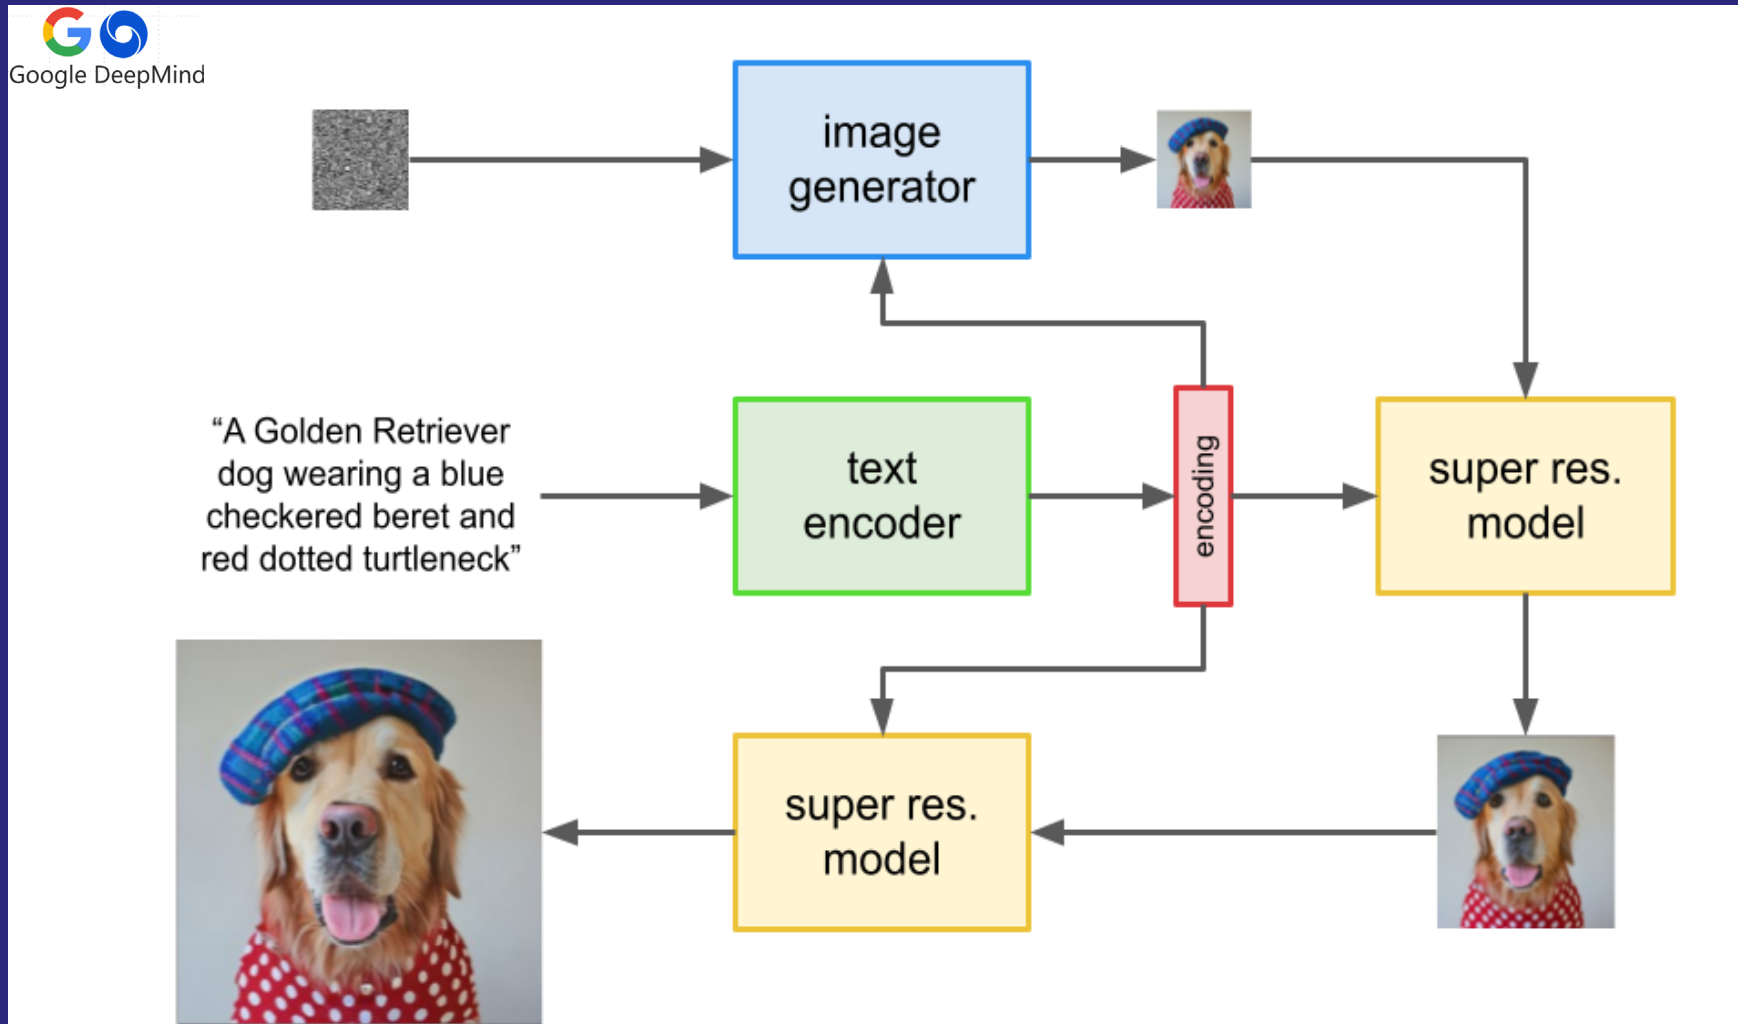
Stable Diffusion (a.k.a. Latent Diffusion)



Pre-trained Auto-Encoder
 Diffusion model in the latent space
 Sampling/Inference:
 Diffusion model -> Decoder

Segmentation-to-Image
 Inpainting
 Superresolution
 Text to Image

ImaGen: Diffusion + Superresolution




Multimodal Generative AI (various data modalities)



Prompt

Here's a solution to a Physics problem by a student.

23. A skier slides down the frictionless slope as shown. What is the skier's speed at the bottom?



Solution:

The conservation of energy means that the energy of the skier is same at the beginning and at the end.

Start: $E = mgL$

End: $E = \frac{1}{2}mv^2$

$\Rightarrow mgL = \frac{1}{2}mv^2$

$\Rightarrow v = \sqrt{2gL} = \sqrt{2 \times 9.81 \times 80} = 39.6 \text{ m.s}^{-1}$

Try to reason about the question step by step.

Did the student get the correct answer? If the solution is wrong, please explain what is wrong and solve the problem.

Make sure to use LaTeX for math and round off the final answer to two decimal places.

Response

The student did not get the correct answer. The student made a mistake in the calculation of the potential energy at the start of the slope. The potential energy at the start of the slope is given by mgh , where m is the mass of the skier, g is the acceleration due to gravity, and h is the height of the slope. The student used the length of the slope, L , instead of the height, H , in the calculation of the potential energy.

The correct solution is:

The conservation of energy means that the total energy of the skier is the same at the beginning and at the end of the slope.

At the start of the slope, the skier has only potential energy, which is given by mgh .

At the end of the slope, the skier has only kinetic energy, which is given by $\frac{1}{2}mv^2$.

Therefore, we can write:

$$mgh = \frac{1}{2}mv^2$$

Canceling out the mass m on both sides, we get:

$$gh = \frac{1}{2}v^2$$

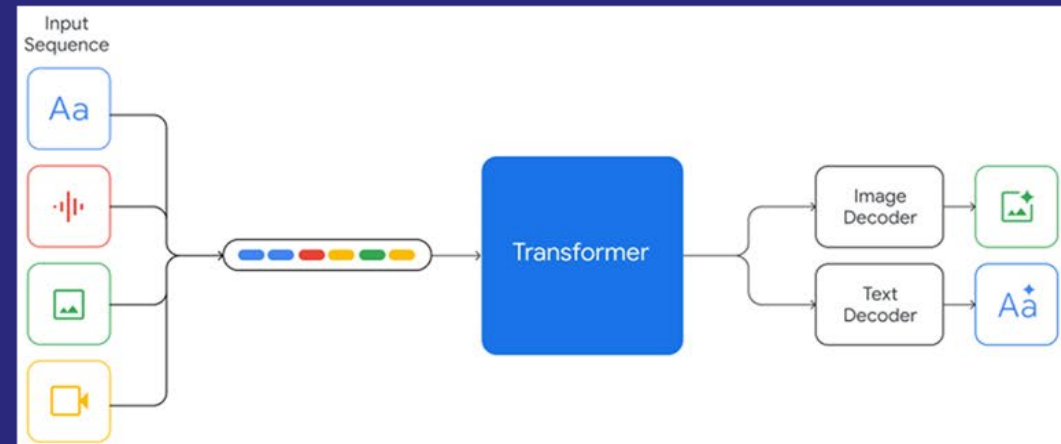
Solving for v , we get:

$$v = \sqrt{2gh}$$

Substituting the given values, we get:

$$v = \sqrt{2 \times 9.81 \times 40} = 28.01 \text{ m/s}$$

Therefore, the skier's speed at the bottom of the slope is 28.01 m/s.



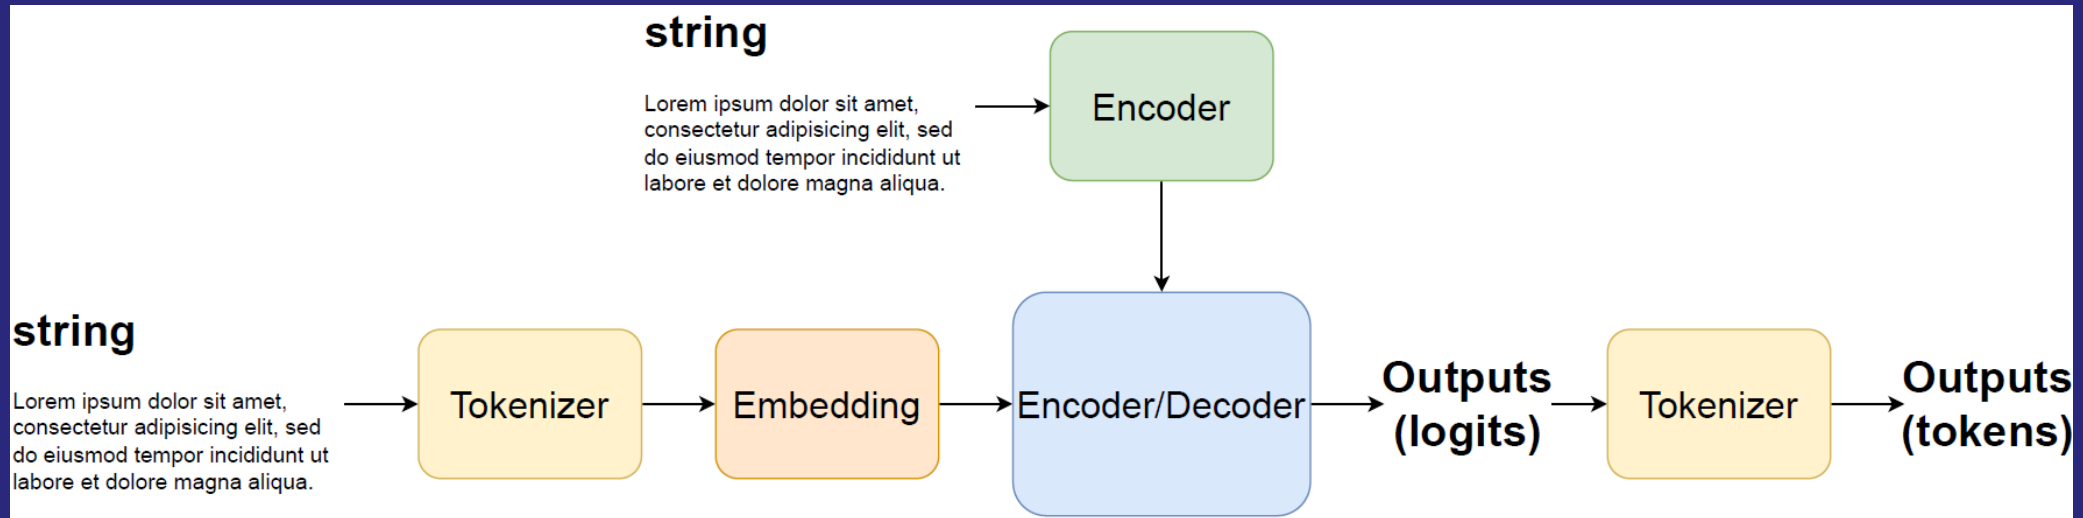
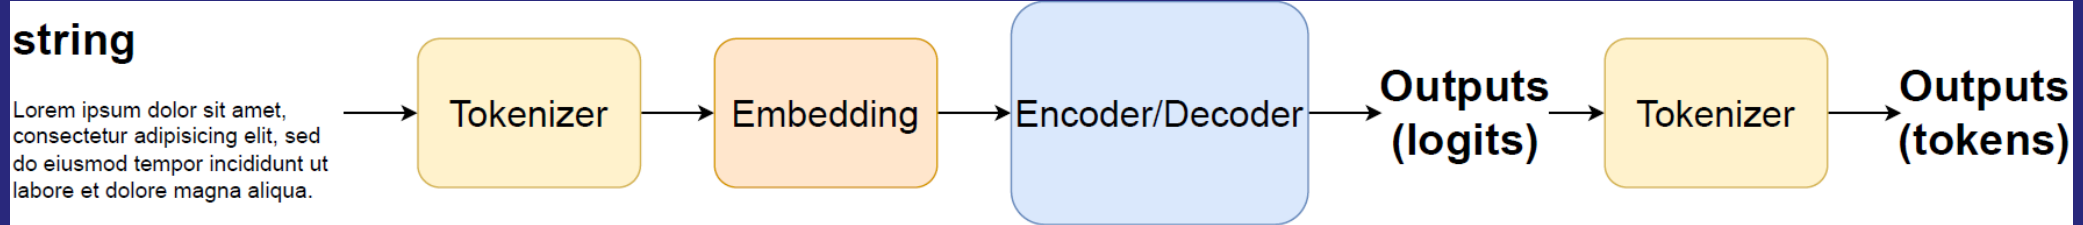
How to design such systems?

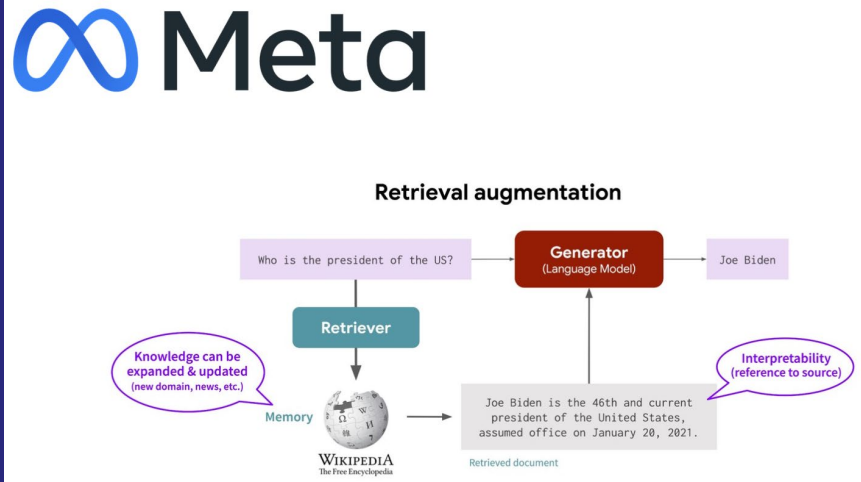
The following questions:

1. What are the common components?
2. What about compositionality?
3. Can we come up with a general structure?
4. Can we come up with general design patterns?

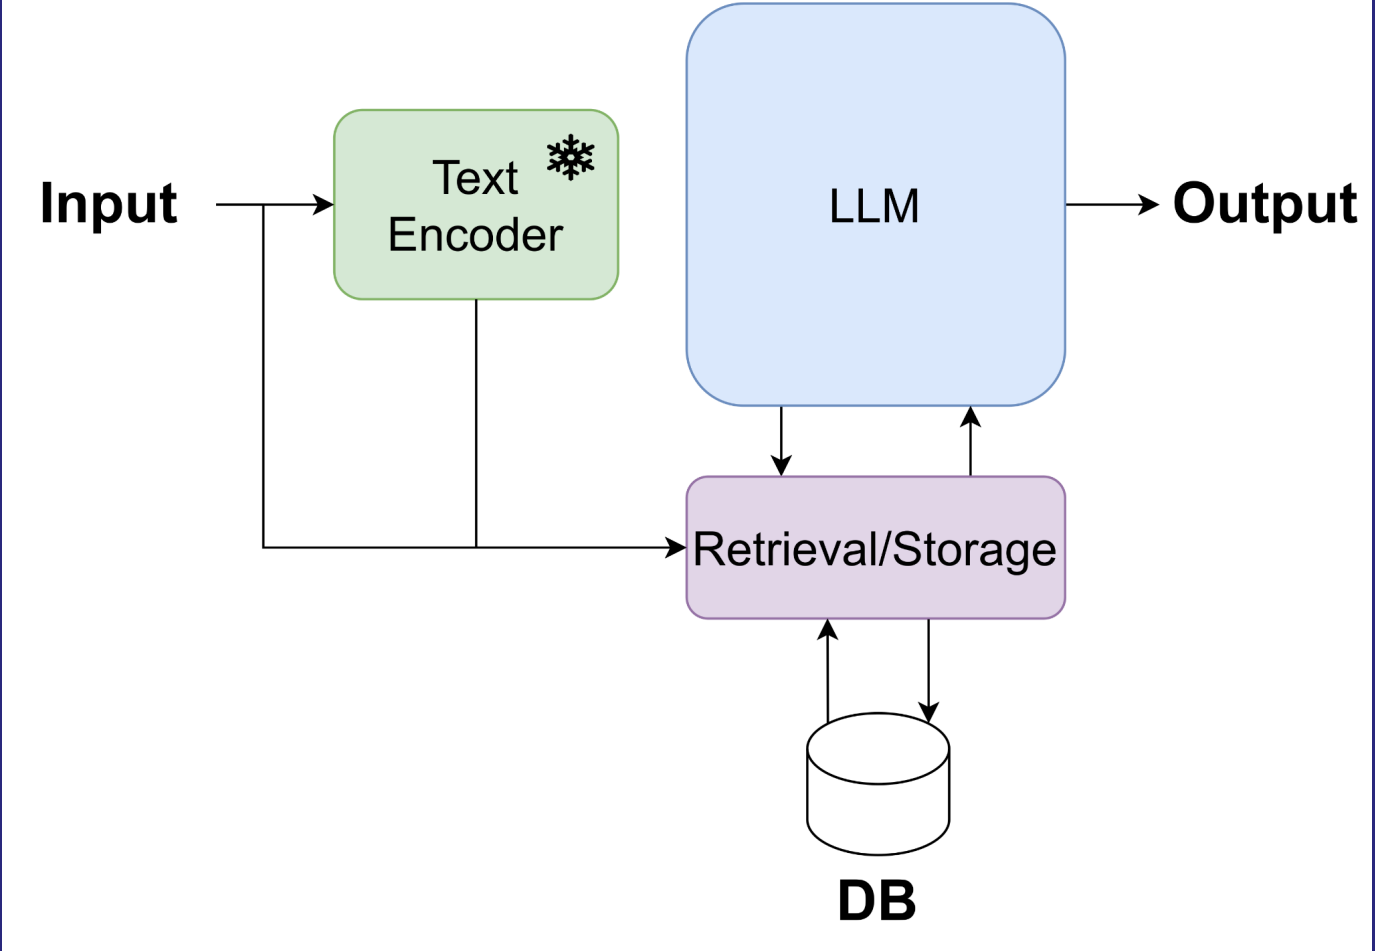


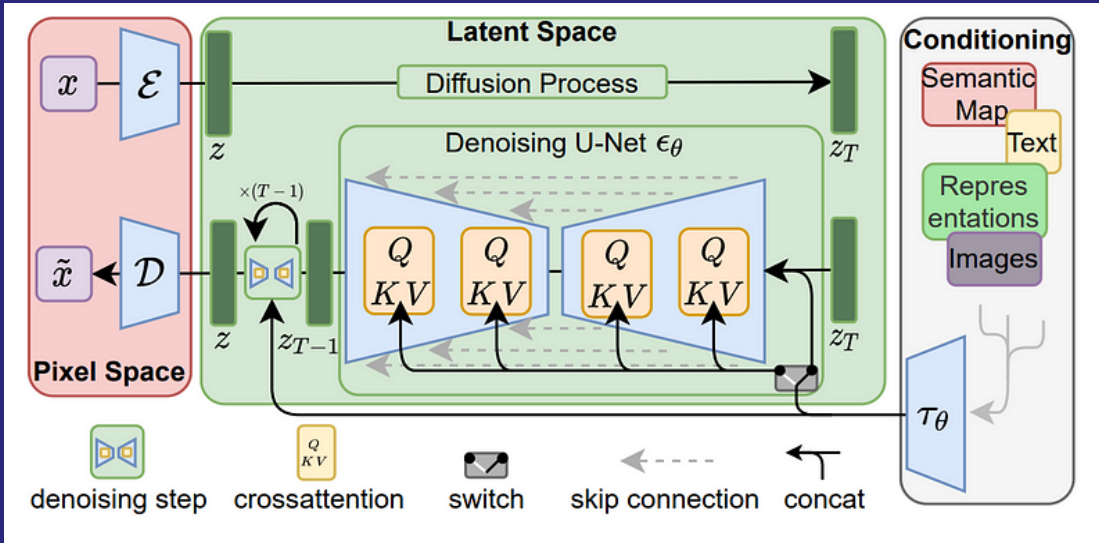
LLMs as systems





- Components:**
- Text Encoder
- Generative Model: LLM
- Retrieval/Storage unit



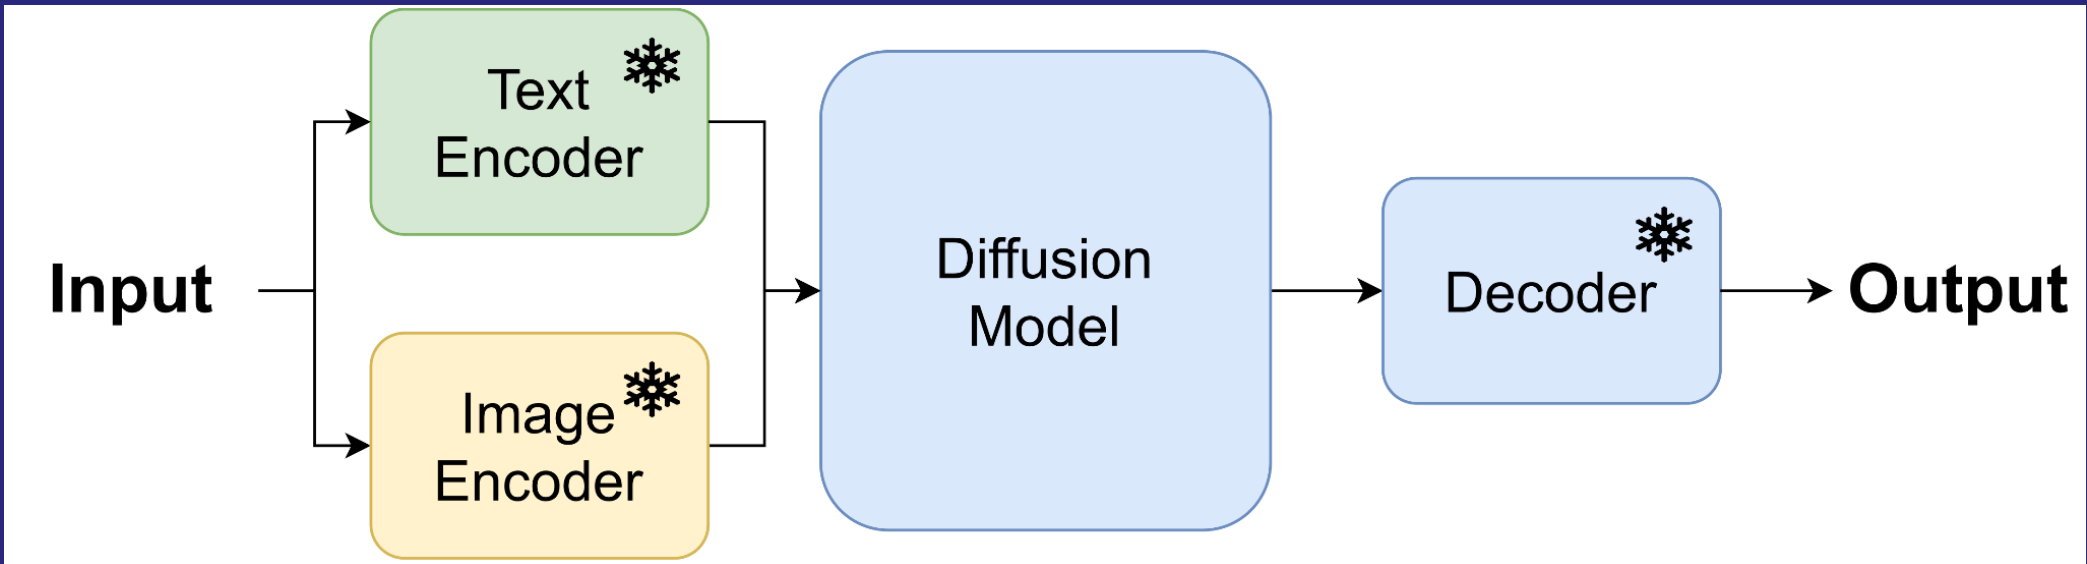


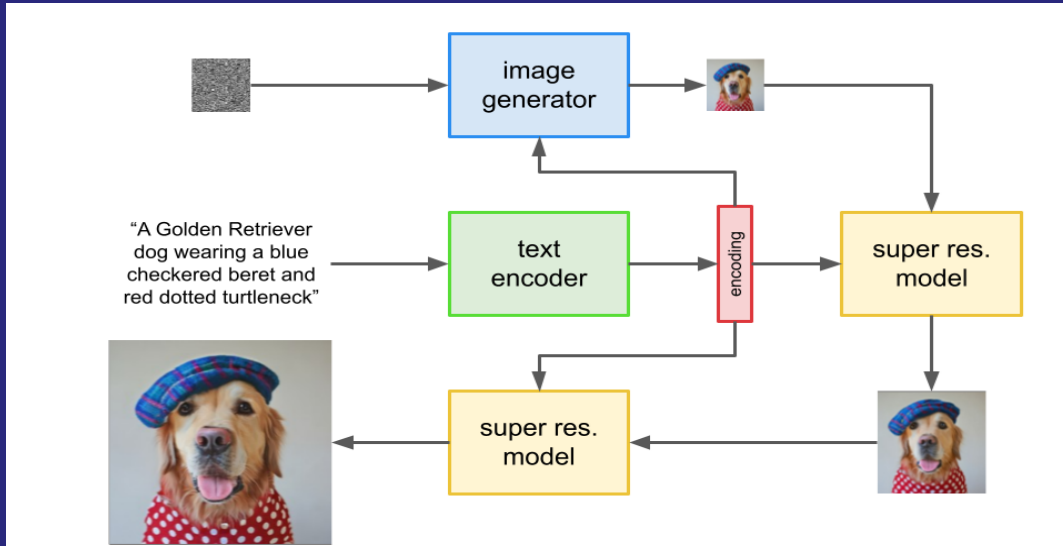
Components:

Text Encoder

Image encoder

Generative Model: Diffusion + Decoder

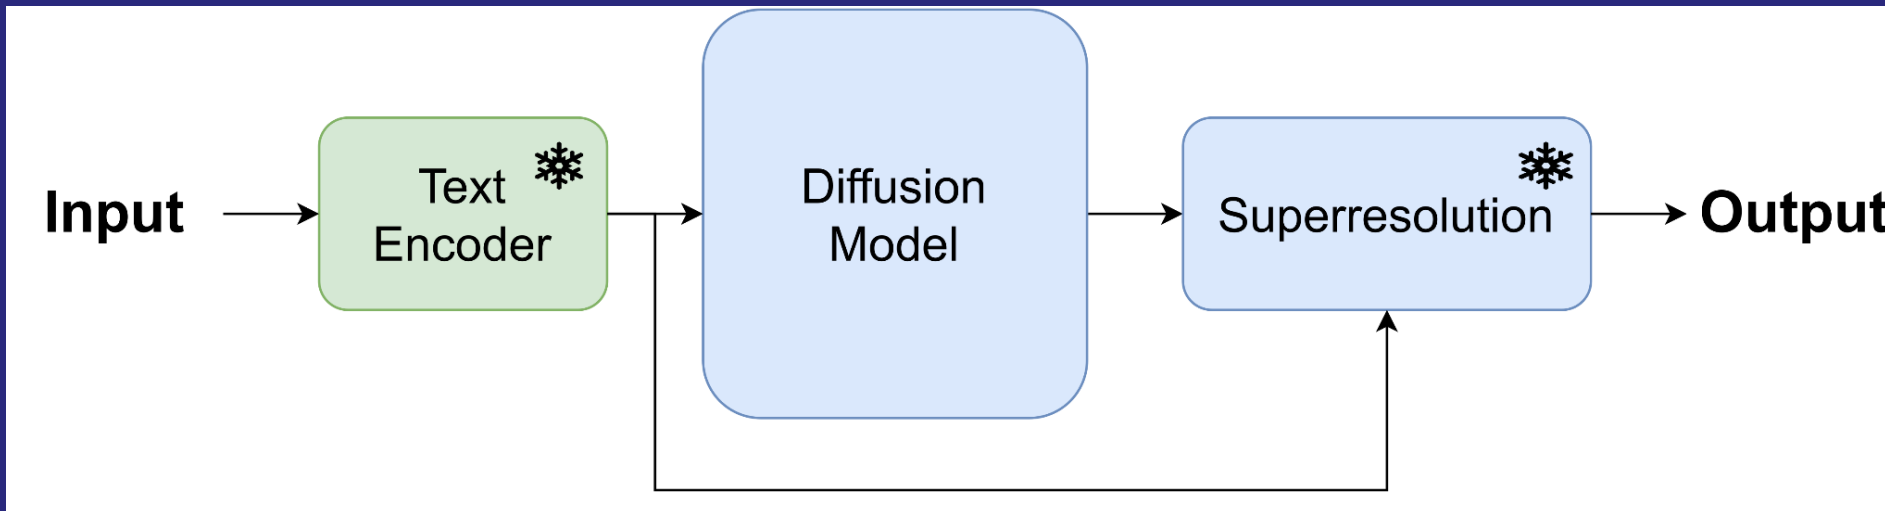




Components:

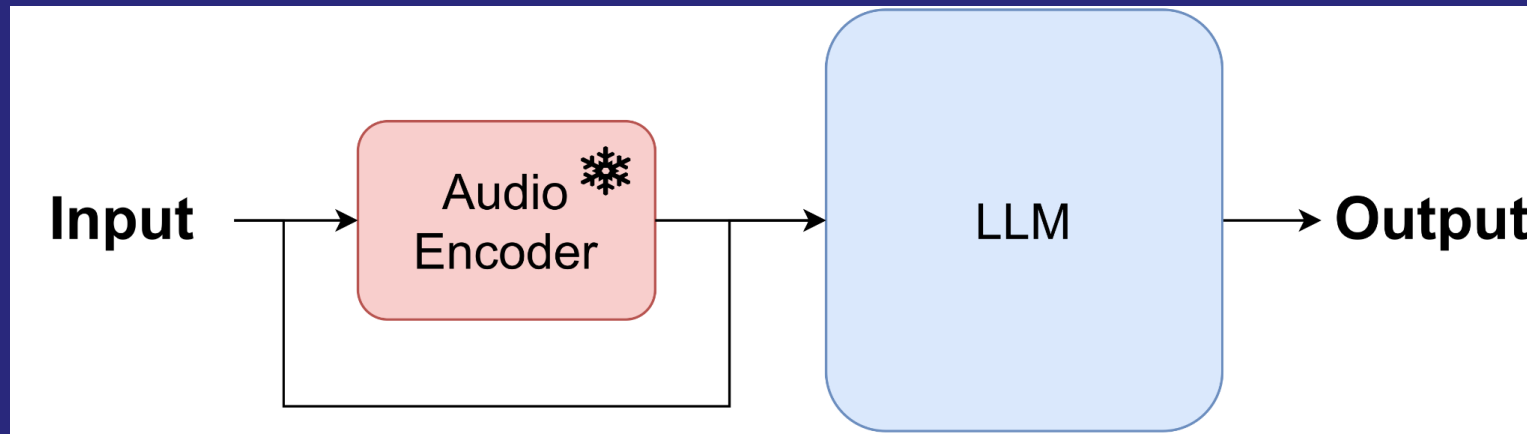
Text Encoder

Generative Model: Diffusion + Superresolution





Components:
Audio Encoder
Generative Model: LLM



Generative AI Systems (GenAISys)

Components:

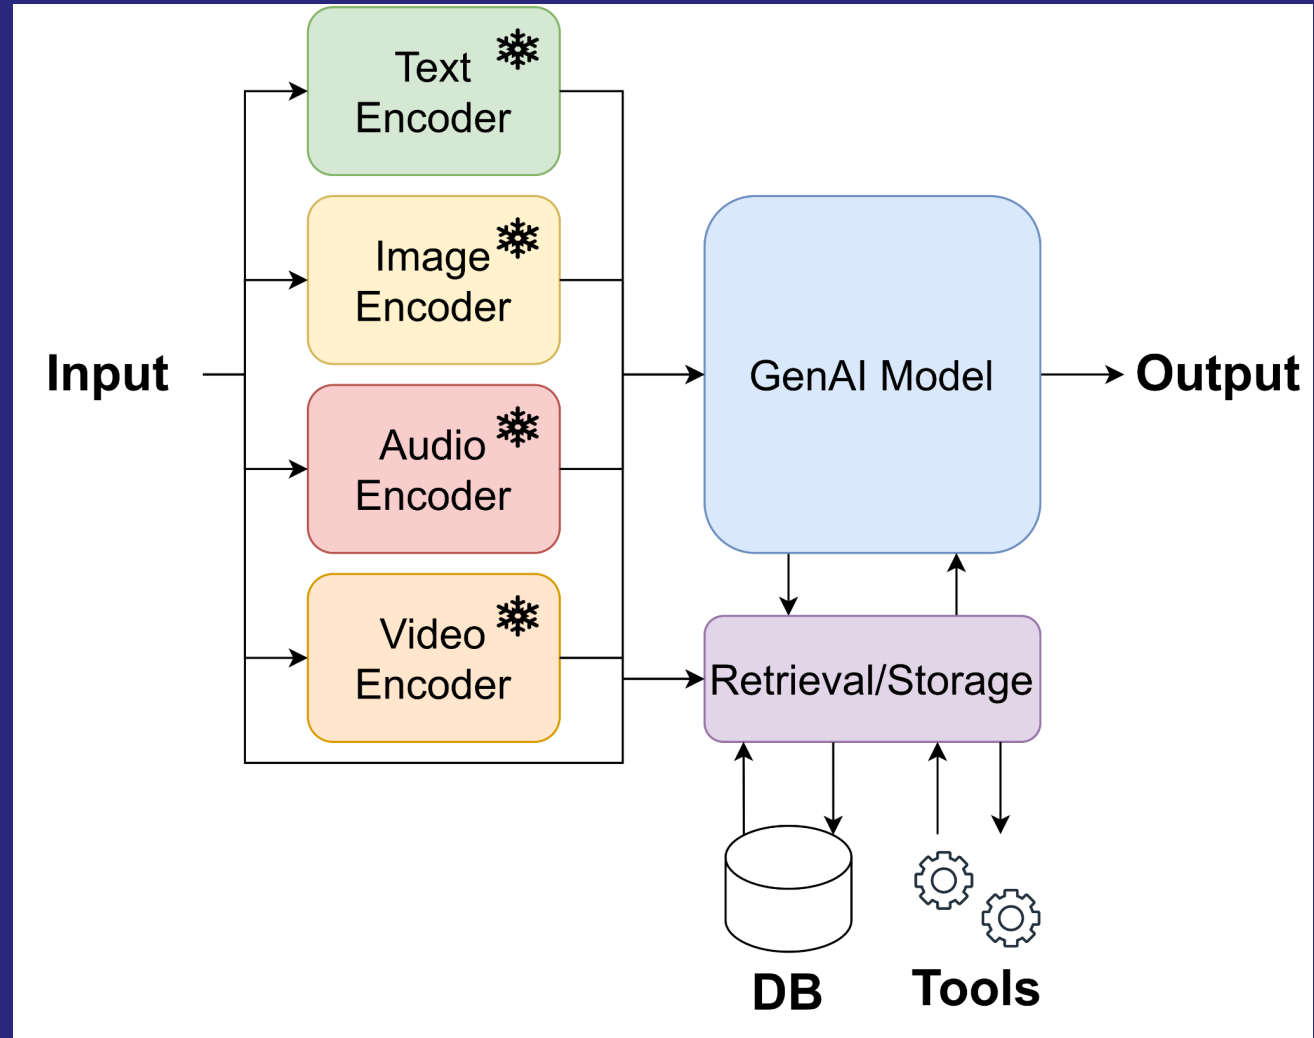
DEs: Data Encoders

GeM: Generative Model

R/S: Retrieval/Storage unit

Database (DB)

Tools



Generative AI Systems (GenAISys)

Compositionality:

atomic system (as) = state + dynamics

composite system (cs) = set of as

composition = how to combine as & cs

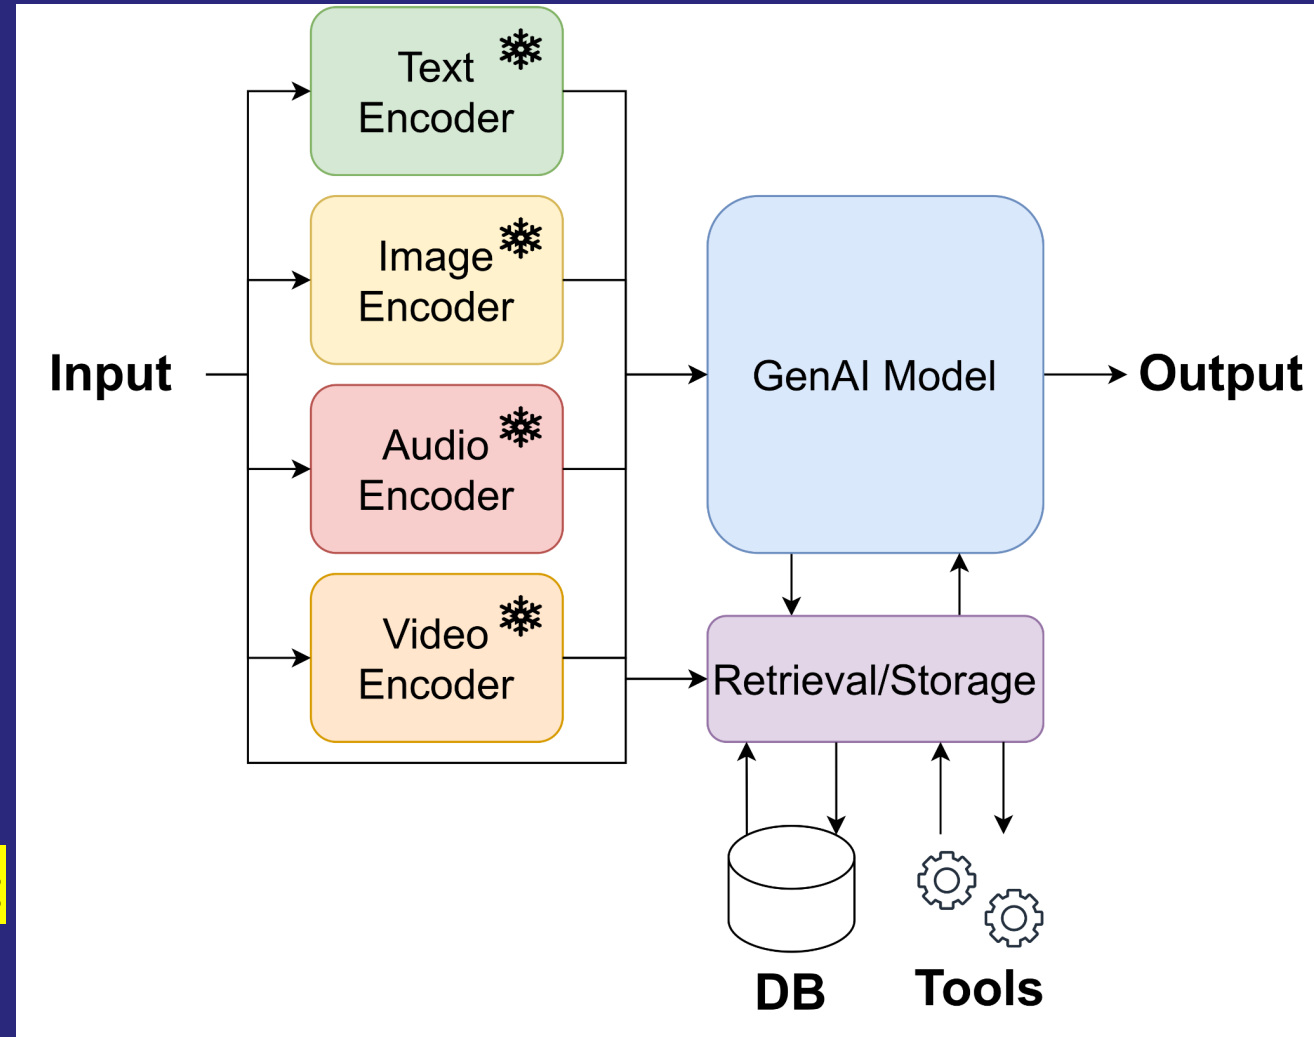
Compatibility:

Outputs of one system are legal as inputs of another system

System Verification & Validation (Reliable AI):

Verification is a set of actions used to check the correctness of a system and its elements.

Validation checks if a system solves a given problem



Conclusion

LLMs: not only a hype but an important landmark in GenAI

GenAISys: we have them, we use them but need to understand them

Future:

Generative AI + Responsive AI = **Agentic AI**

How to formally analyze/understand GenAISys? **Category Theory?**

Thank you!

Questions?

Contact: j.m.tomczak@tue.nl
jmk.tomczak@gmail.com

Generativ/e

Generative AI Group: <https://generativeai-tue.github.io/>