

# Generative AI

## The key concepts, Developments and Challenges

Jakub M. Tomczak  
Generative AI Group, TU/e (associate professor)  
Amsterdam AI Solutions (founder)

**Generativ/e**

**Amsterdam**  
**AI Solutions**

# Recent breakthroughs

# ChatGPT

User help me write a short note to introduce myself to my neighbor

ChatGPT Hi there,

My name is [Your Name] and I just moved in next door. I wanted to reach out and introduce myself. I'm looking forward to getting to know you and the neighborhood. If you ever need anything or just want to chat, please don't hesitate to stop by or give me a call.

Best,  
[Your Name]

User can you make it more formal?

ChatGPT Dear Neighbor,

I am writing to introduce myself. My name is [Your Name] and I have recently moved in next door. I wanted to take a moment to reach out and say hello. I am looking forward to getting to know you and the neighborhood.

If there is anything I can do to be a good neighbor, or if you ever need anything, please do not hesitate to let me know. I am happy to help in any way that I can.

Sincerely,  
[Your Name]

User can you send it for me?

ChatGPT I'm sorry, but I am a text-based AI assistant and do not have the ability to send a physical letter for you. I can only provide suggestions and advice on what to write. It is up to you to take that advice and send the letter yourself.

User this code is not working like i expect — how do i fix it?

```
resultWorkerErr := make(chan error)
defer close(resultWorkerErr)
go func() {
    defer cancel()
    resultWorkerErr <- b.resultWorker(ctx)
}()

err := b.worker(ctx)
cancel()
if err == nil {
    return <-resultWorkerErr
}
return multierror.Append(err, <-resultWorkerErr)
```

ChatGPT It's difficult to say what's wrong with the code without more context. Can you provide more information about what the code is supposed to do and what isn't working as expected? Also, is this the entire code or just a part of it?

User just part of it — the error never surfaces. I think there is something wrong with the channel

ChatGPT It's difficult to say without more information about what the code is supposed to do and what's happening when it's executed. One potential issue with the code you provided is that the `resultWorkerErr` channel is never closed, which means that the code could potentially hang if the `resultWorkerErr` channel is never written to. This could happen if `b.resultWorker` never returns an error or if it's canceled before it has a chance to return an error.

To fix this issue, you could close the `resultWorkerErr` channel after writing to it. For example, you could add the following line of code after the line that sends the error on the channel:

```
close(resultWorkerErr)
```

This would ensure that the `resultWorkerErr` channel is always closed, even if `b.resultWorker` never returns an error. Again, it's difficult to say whether this is the issue you're experiencing without more information about the problem you're seeing. Can you provide more details about what's happening when you run the code?

# GPT-4

User What is funny about this image? Describe it panel by panel.



Source: [hmmm \(Reddit\)](#)

GPT-4 The image shows a package for a "Lightning Cable" adapter with three panels.

Panel 1: A smartphone with a VGA connector (a large, blue, 15-pin connector typically used for computer monitors) plugged into its charging port.

Panel 2: The package for the "Lightning Cable" adapter with a picture of a VGA connector on it.

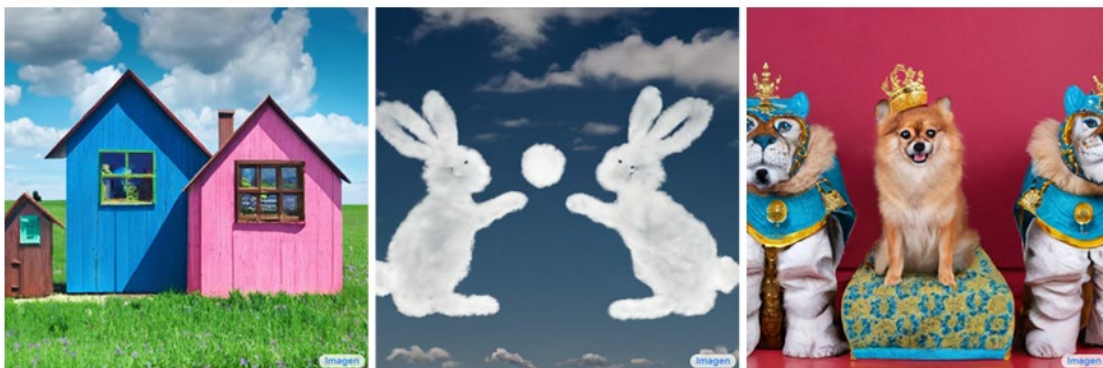
Panel 3: A close-up of the VGA connector with a small Lightning connector (used for charging iPhones and other Apple devices) at the end.

The humor in this image comes from the absurdity of plugging a large, outdated VGA connector into a small, modern smartphone charging port.

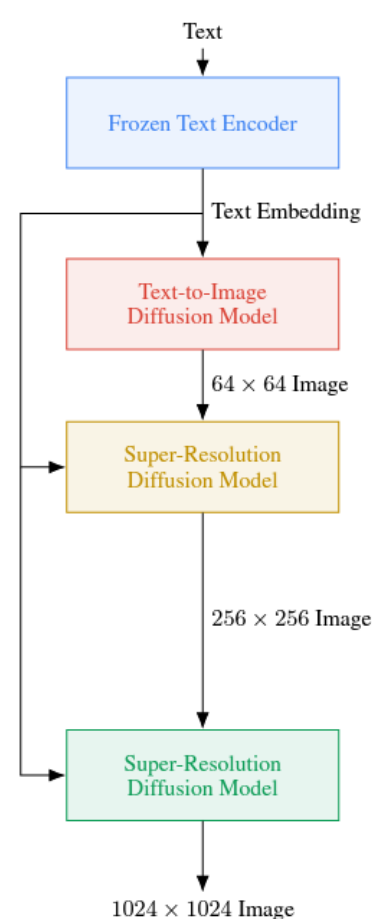
# Imagen: Transformer Encoder + Diffusion model



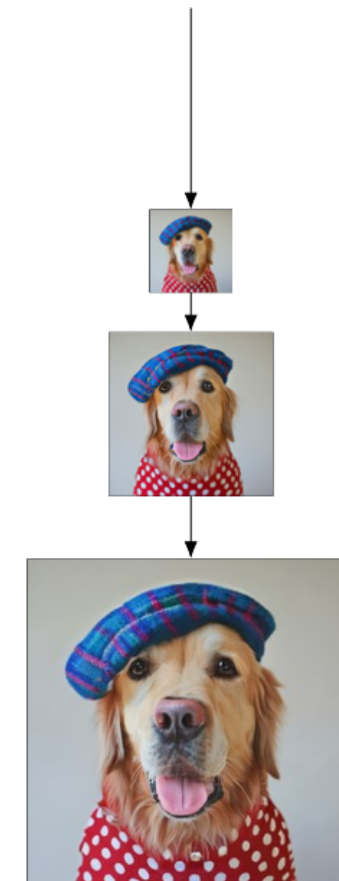
A wall in a royal castle. There are two paintings on the wall. The one on the left a detailed oil painting of celebrating the birthday of their friend. There is a pizza cake on the desk. A chrome-plated duck with a golden beak arguing with an angry turtle in a forest.



A family of three houses in a meadow. The Dad house is a large blue house. The Mom house is a large pink house. The Child house is a small wooden shed. A cloud in the shape of two bunnies playing with a ball. The ball is made of clouds too. A Pomeranian is sitting on the Kings throne wearing a crown. Two tiger soldiers are standing next to the throne.



"A Golden Retriever dog wearing a blue checkered beret and red dotted turtleneck."



# Stable Diffusion

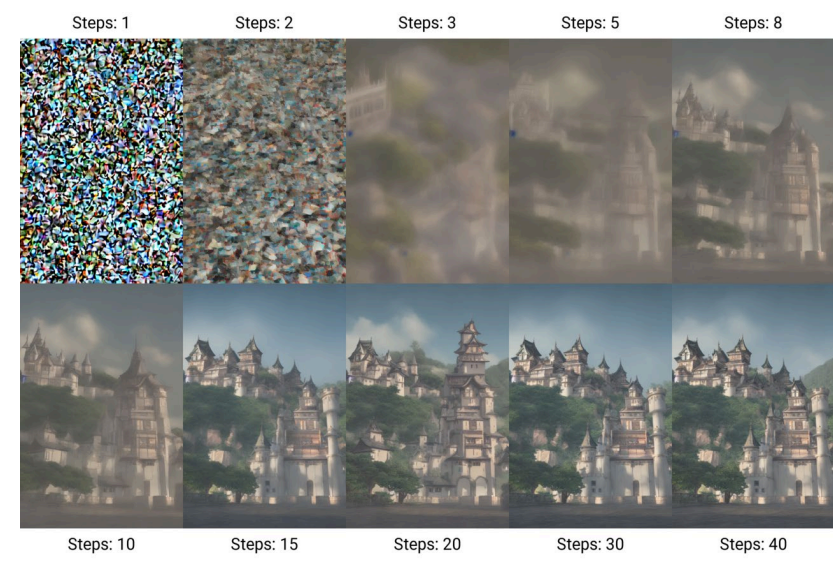
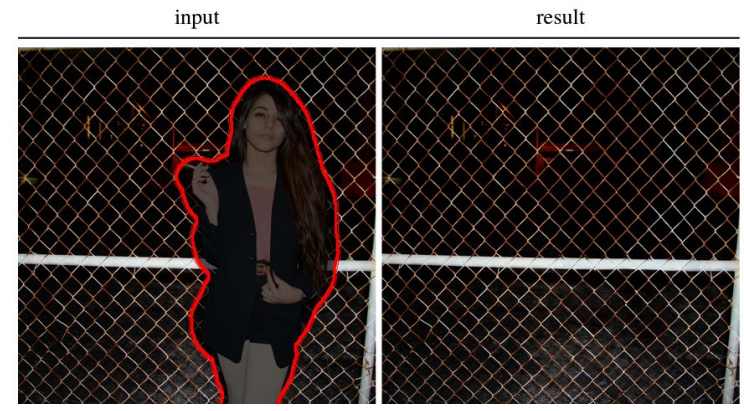
## Text2Image

'A street sign that reads "Latent Diffusion"'

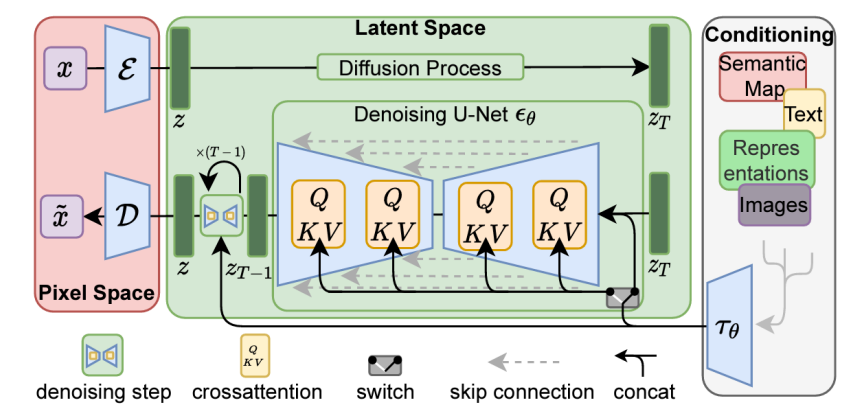
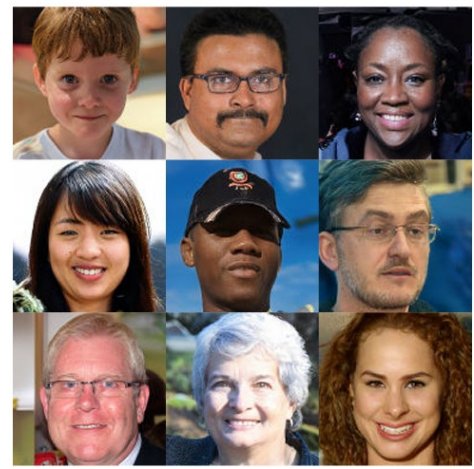
'A zombie in the style of Picasso'



## Inpainting



## Unconditional



# Generative AI

# Modeling: Discriminative vs. Generative

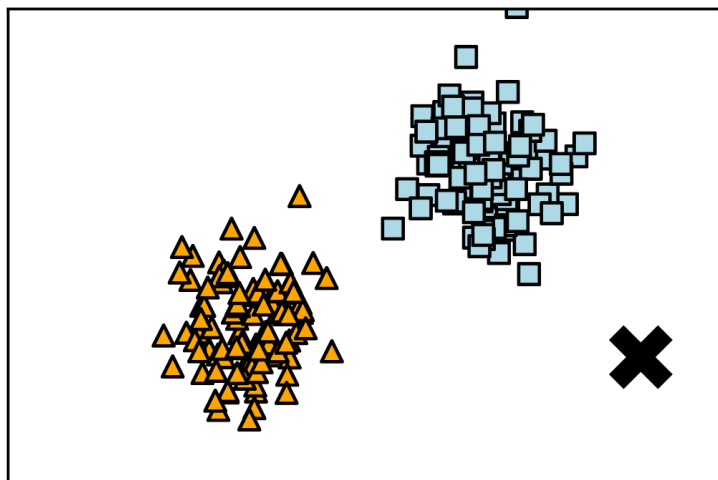
- **Discriminative models:** finding boundaries in the data space,  $p(y|x)$ .
- **Generative models:** explaining how the data was generated,  $p(x, y)$ .



# Modeling: Discriminative vs. Generative

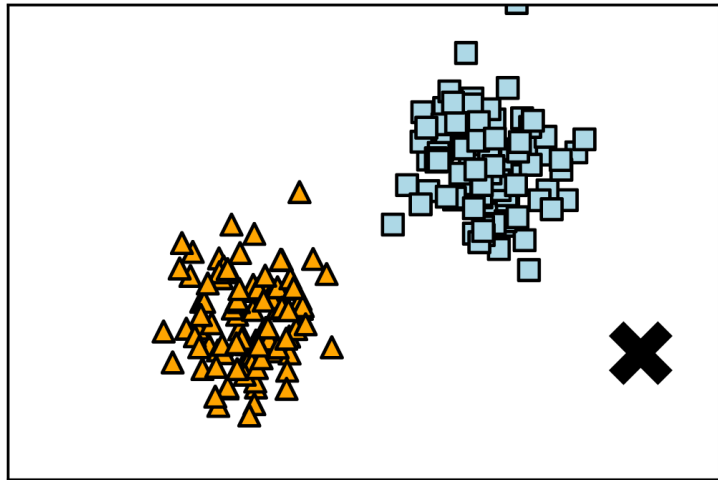
- **Discriminative models:** finding boundaries in the data space,  $p(y|x)$ .
- **Generative models:** explaining how the data was generated,  $p(x, y)$ .
- In ML, many models are generative:
  - Naive Bayes, Linear Discriminant Analysis
  - Bayesian networks & Markov random fields
  - Mixture Models, Latent Dirichlet Allocation, Factor Analysis, PCA
  - Chinese restaurant process, Indian buffet process

# Generative = Understanding how data is generated

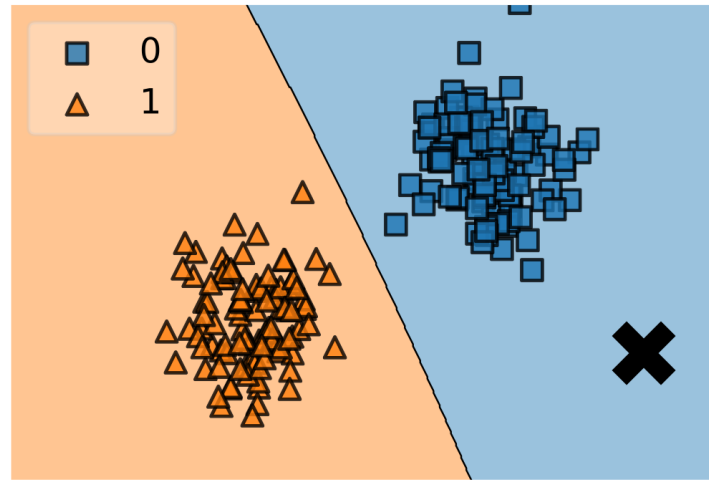


Data

# Generative = Understanding how data is generated



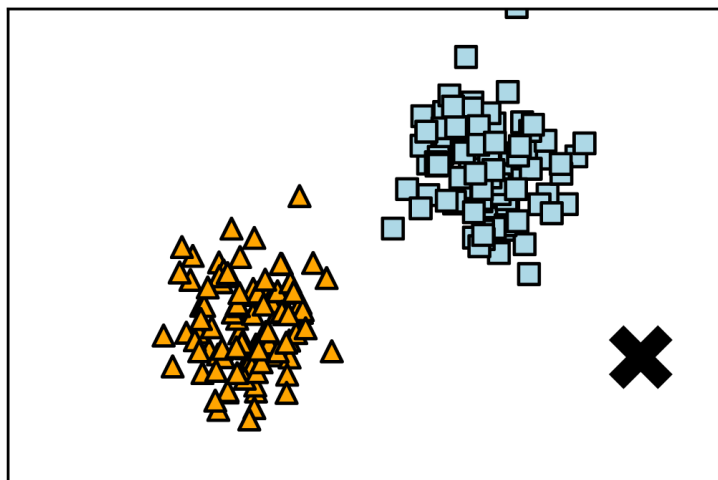
Data



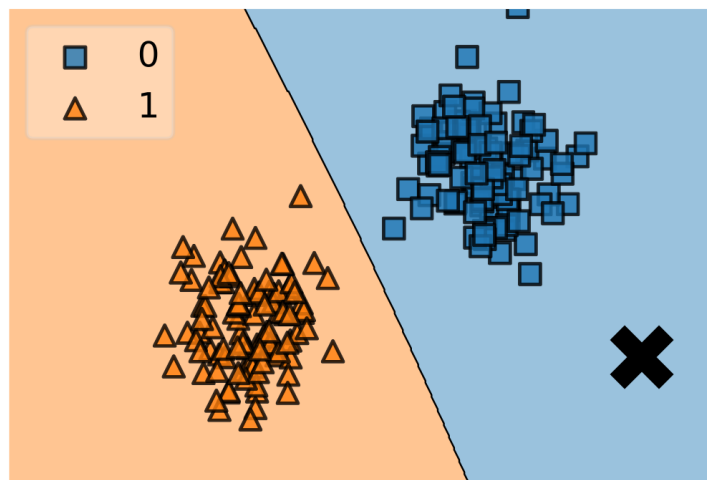
$p(y|\mathbf{x})$

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

# Generative = Understanding how data is generated

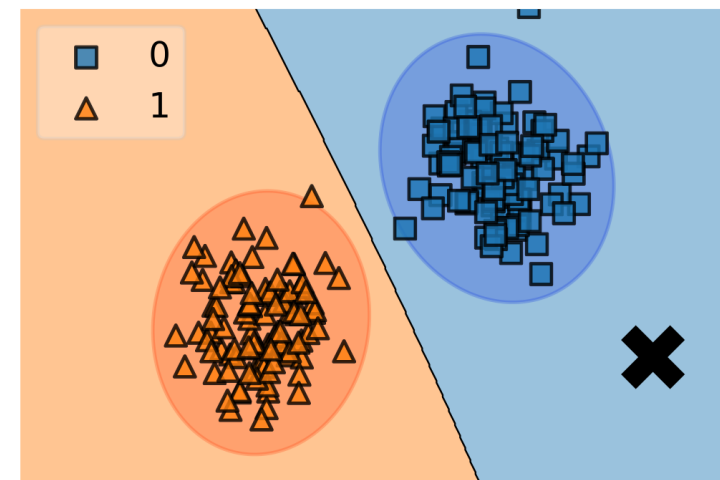


Data



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

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



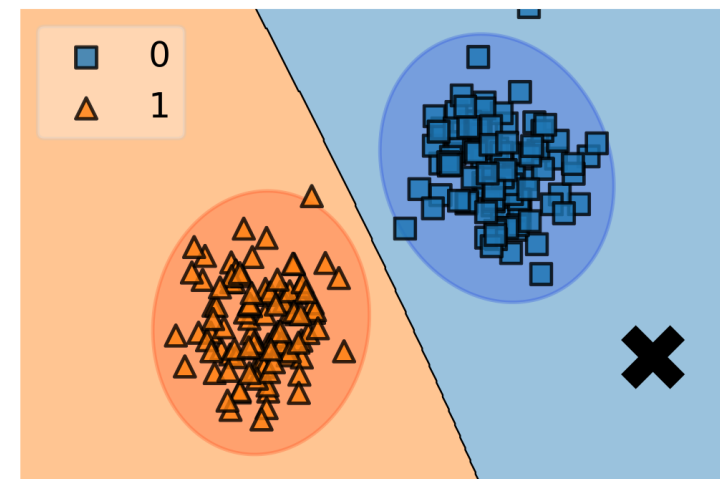
$p(\mathbf{x}, y) = p(y|\mathbf{x}) p(\mathbf{x})$

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

# Generative = Understanding how data is generated

## Pros and cons

- Uncertainty quantification
- Semi-supervised learning for *free*
- High flexibility
- Natural for multi-modal learning
- Flexibility for domain adaptation
- High complexity



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

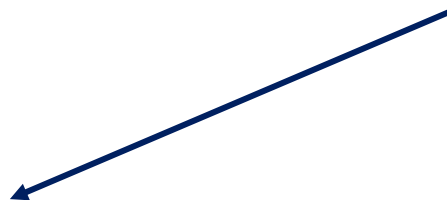
$p(\text{blue}|\mathbf{x})$  is high  
and  $p(\mathbf{x})$  is low  
= uncertain decision!

# Generative AI

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

# Generative AI

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



Any (deep learning) predictor

Relatively easy

# Generative AI

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

Any (deep learning) predictor

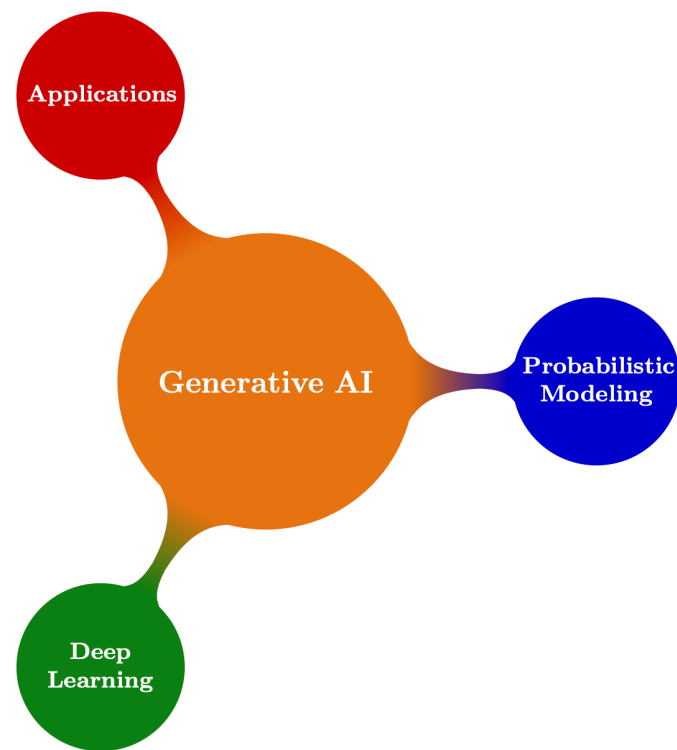
Relatively easy

A deep probabilistic model

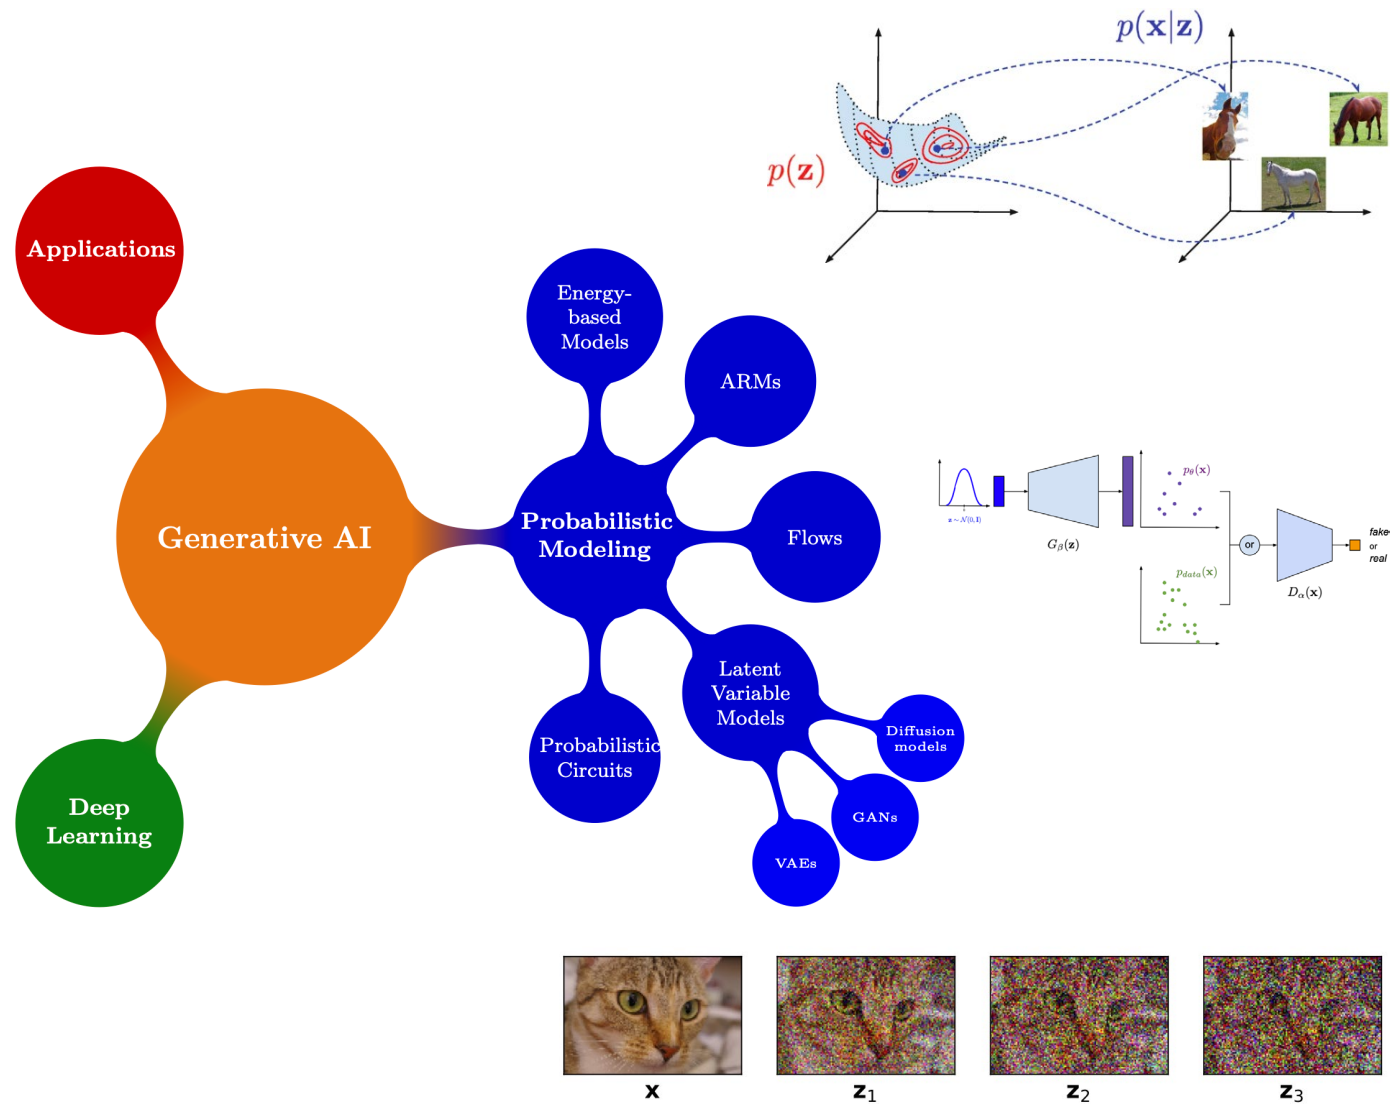
Challenging!



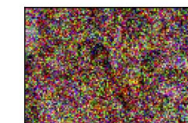
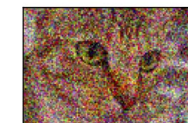
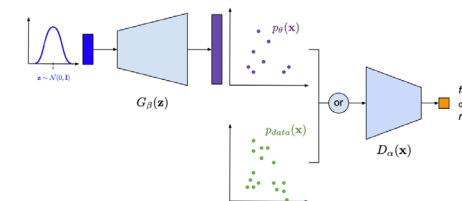
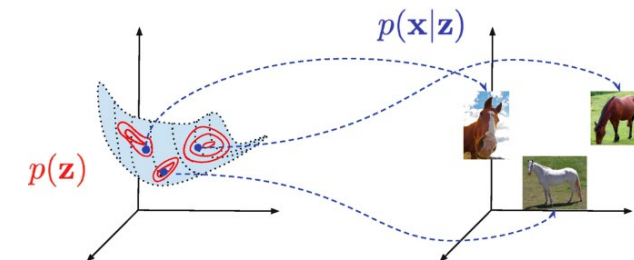
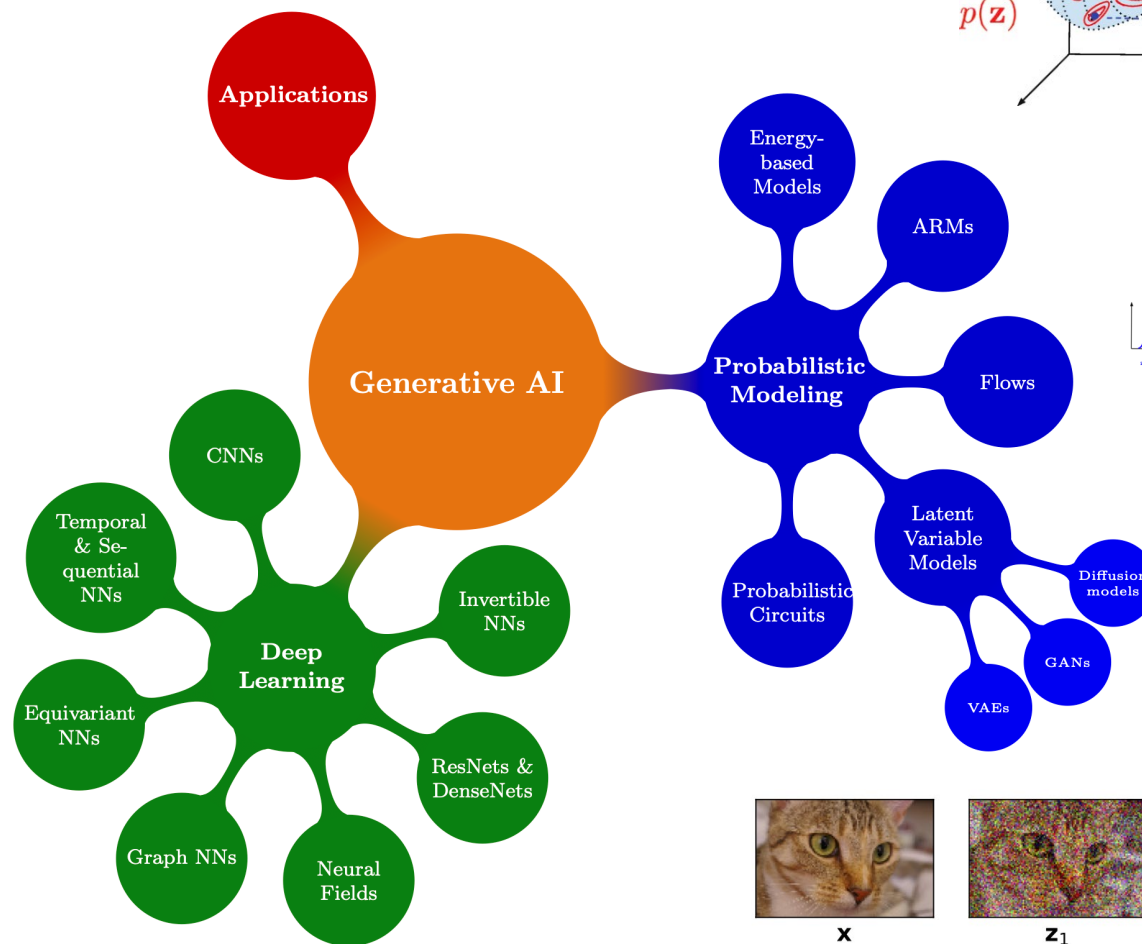
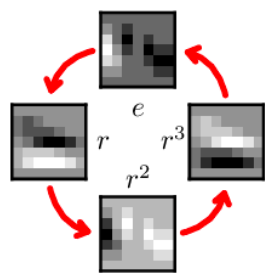
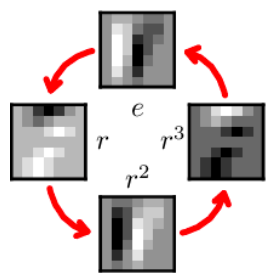
# Generative AI



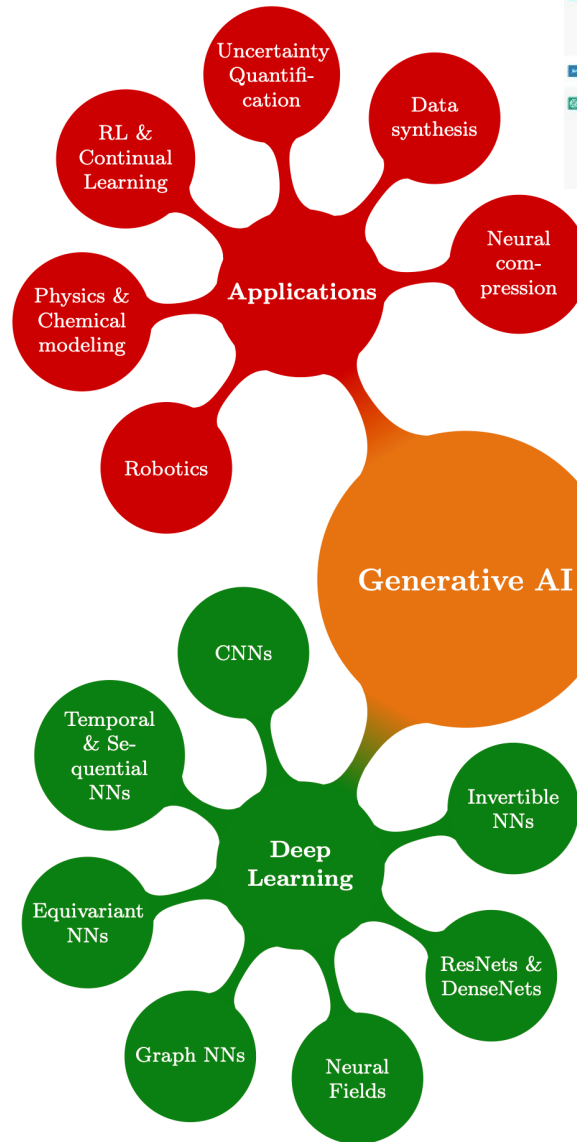
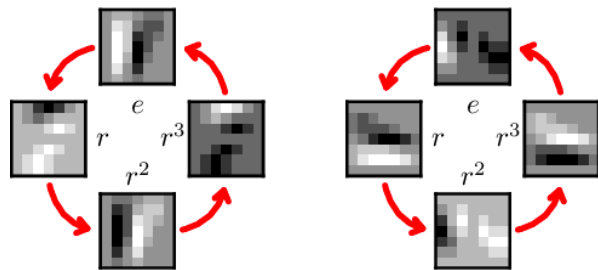
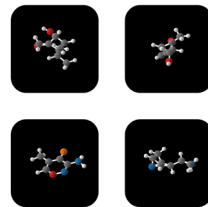
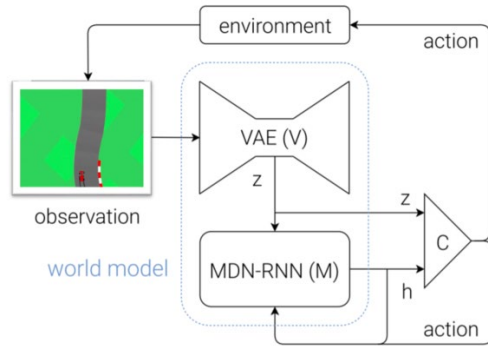
# Generative AI



# Generative AI



# Generative AI



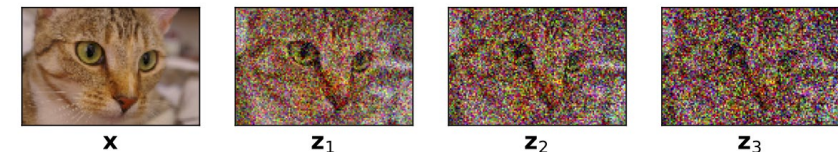
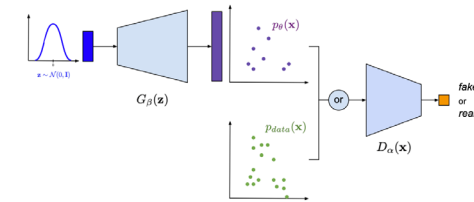
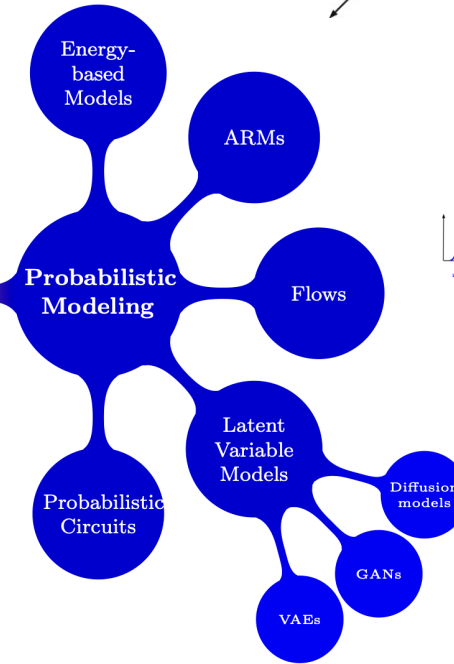
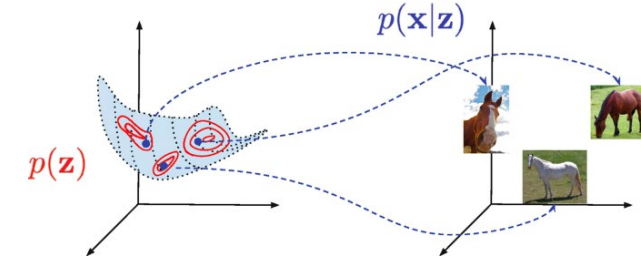
**Q** What is generative artificial intelligence?

**A** Generative artificial intelligence (AI) refers to a subset of AI that is focused on creating or generating new and original content or data. This can include tasks such as image generation, text generation, and audio synthesis. Generative AI models are typically trained on large datasets and use techniques such as deep learning and unsupervised learning to generate new content that is similar to the data they were trained on.

**Q** Do you know who works on that at the Technische University of Technology?

**A** I am not aware of specific researchers at Technische University of Technology who focus on generative artificial intelligence. However, the university does have a strong research program in the broader field of artificial intelligence, and it is likely that there are researchers there who work on various aspects of generative AI.

It would be best to check the university website or reach out to the university's Artificial Intelligence (AI) or Computer Science department, to find the researchers that are currently working on generative AI.



# Generative AI and (spherical) cows



High-dim data

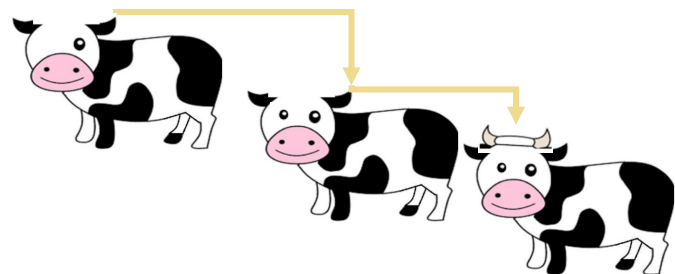


Latent variable models

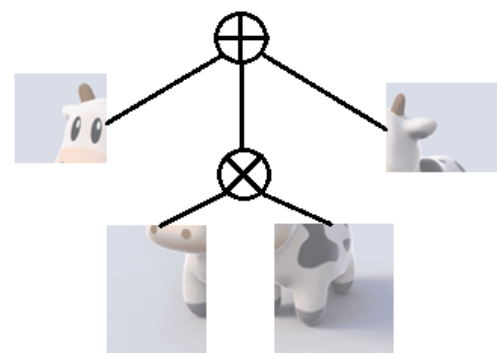


Flow-based models

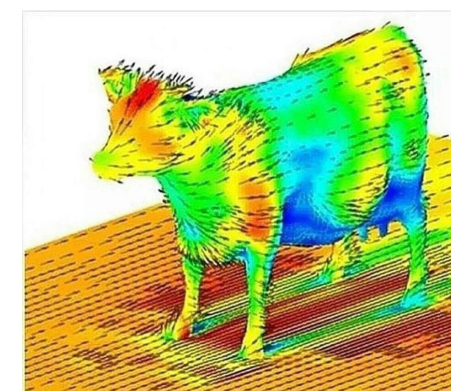
**GOAL:**  
 $p(\mathbf{x}, y)$  or  $p(\mathbf{x})$



Autoregressive models



Probabilistic circuits

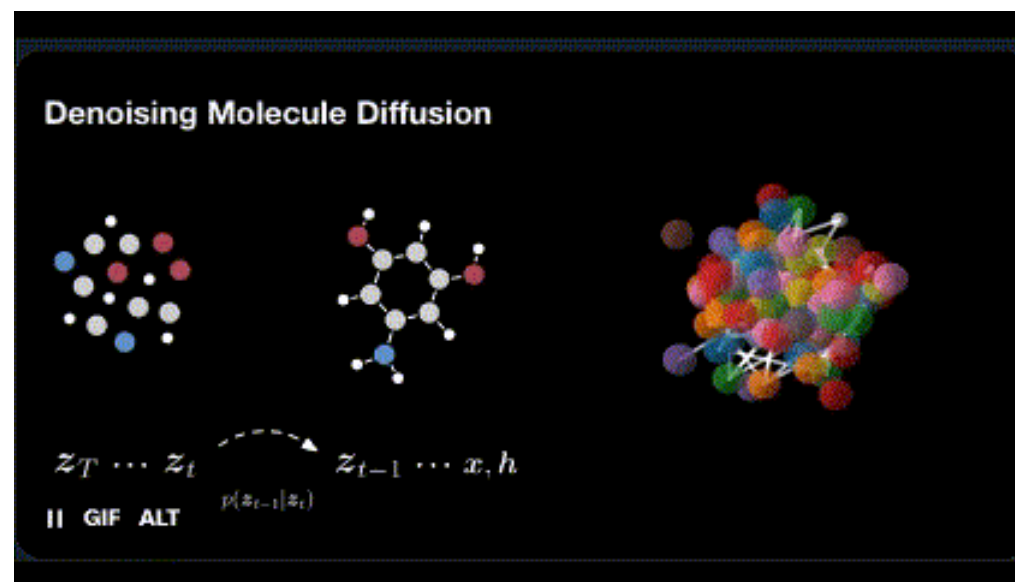
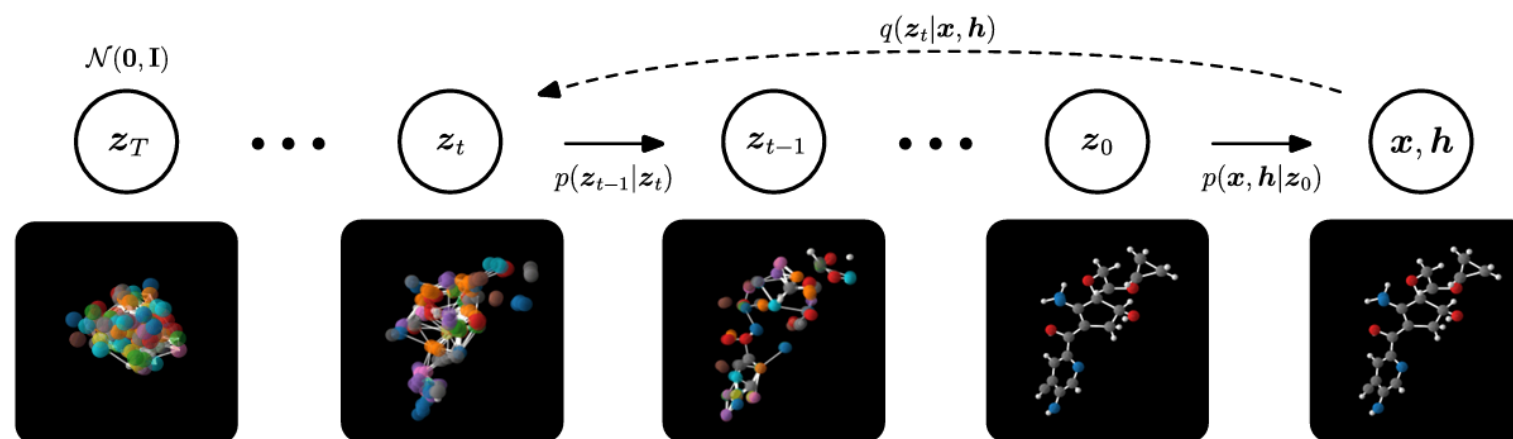


Diffusion models

Energy-based models

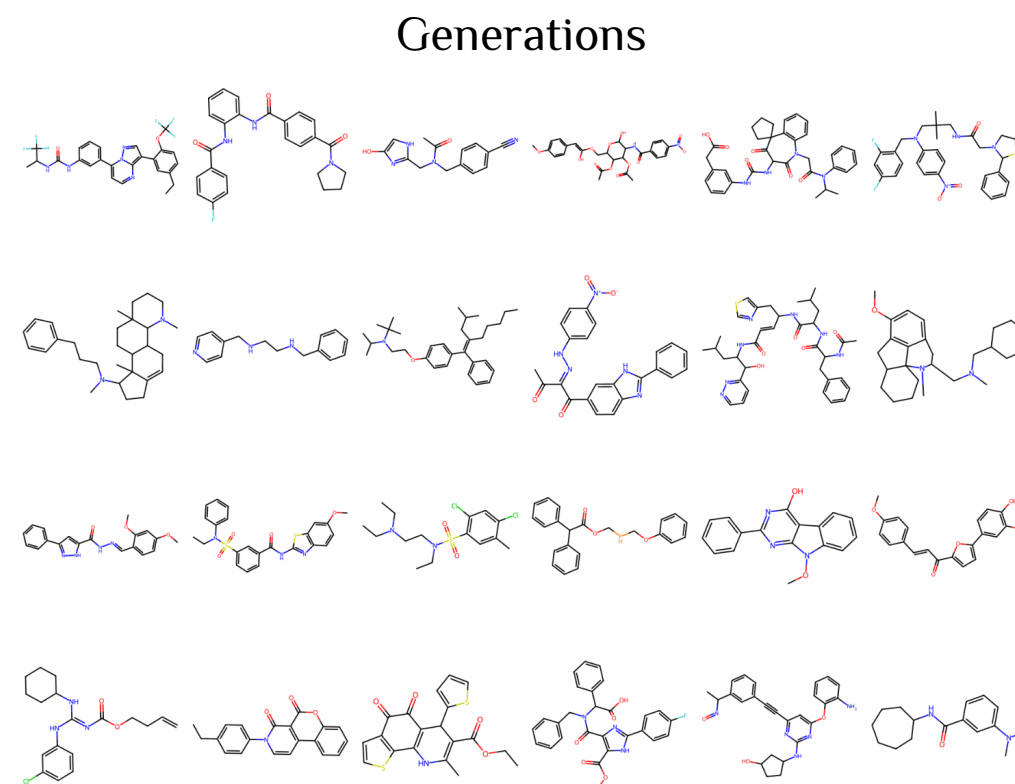
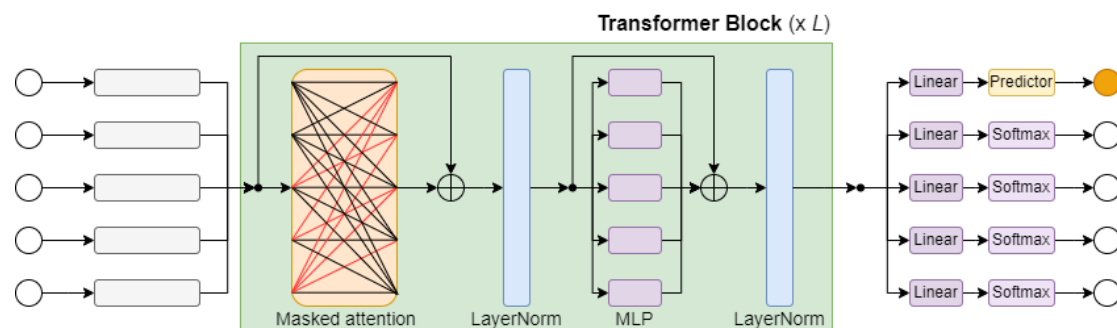
# A glance at *non-standard* applications of Generative AI

# Molecule generation with Diffusion models



# Molecule generation with Joint Transformers

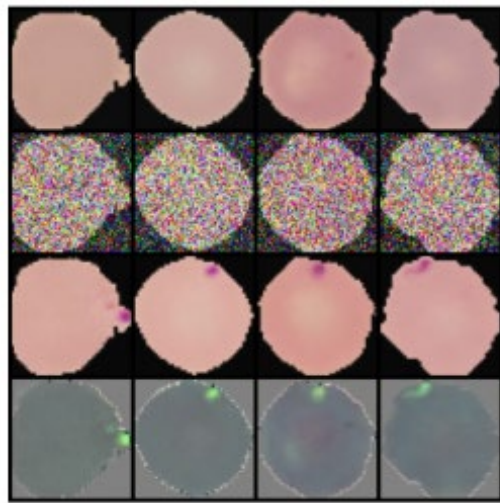
- Trained on 15M SMILES strings (1.5M molecules) with >36M updates.
- Tasks:
  - Molecule generation
  - Property prediction
  - *De novo* drug design



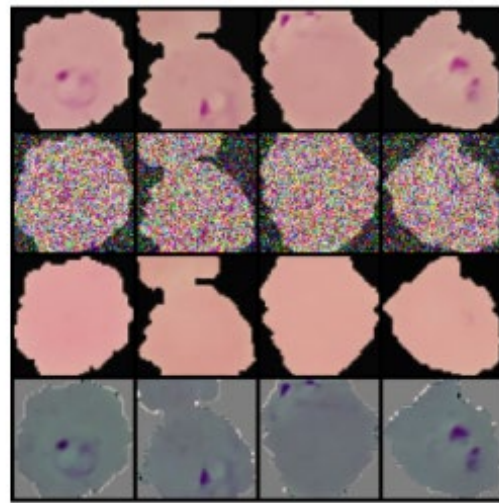


# Medical imaging with Joint Diffusion Models

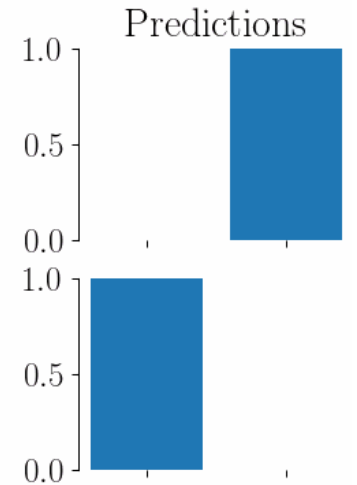
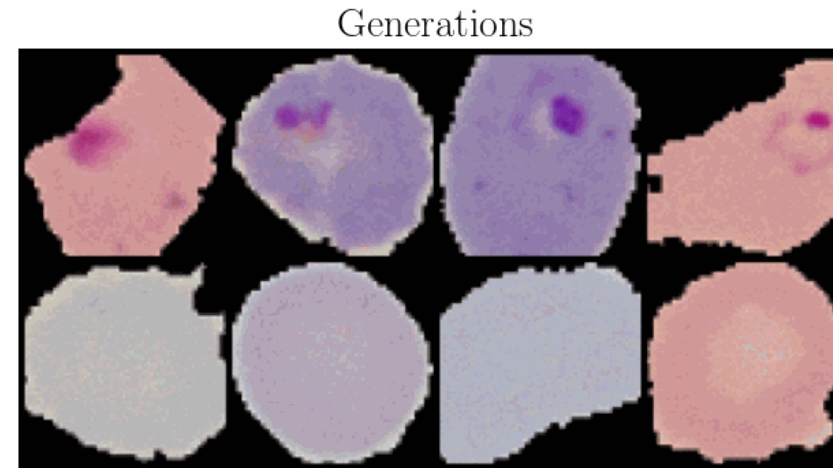
- Visual counterfactual explanations:



Negative examples



Positive examples



# Challenges

# Challenges

- Trustworthiness



what is heavier: kilo of water or kilo of air?



A kilogram of water is heavier than a kilogram of air. Water is a dense, liquid substance that has a mass of about 1 kilogram per liter at standard temperature and pressure. Air, on the other hand, is a gas that is much less dense than water and has a lower mass. A volume of air that weighs 1 kilogram would occupy a much larger space than a liter of water, which weighs the same amount.



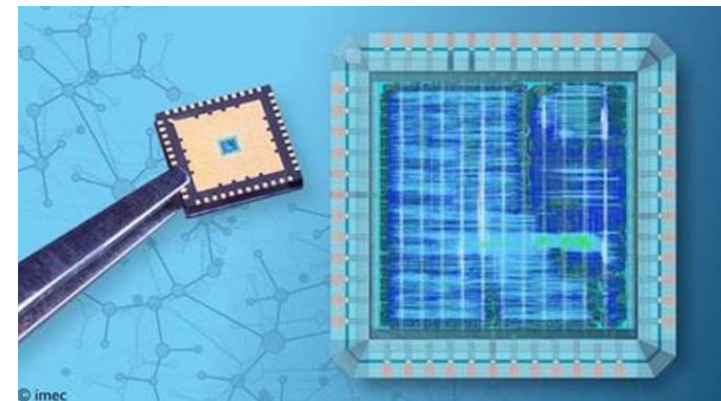
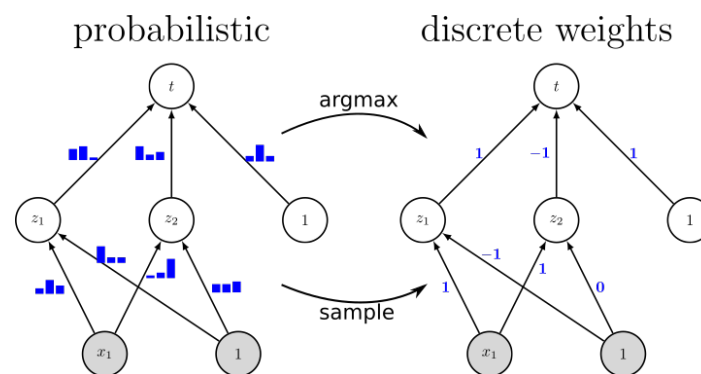
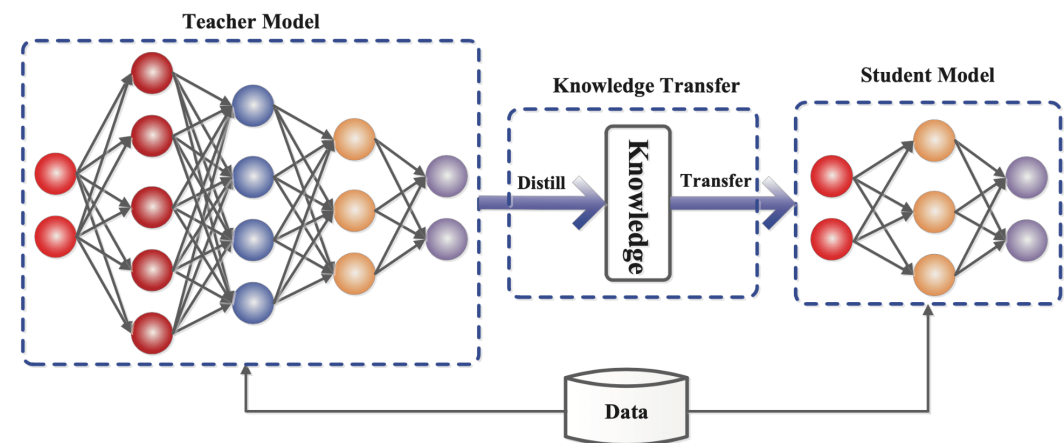
# Challenges

- Trustworthiness

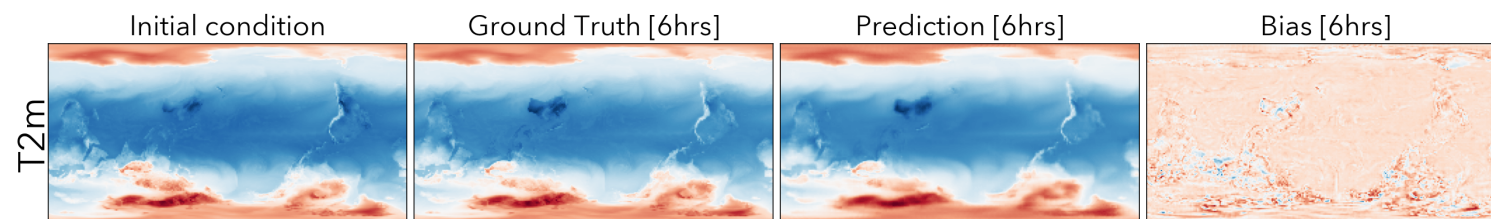


# Challenges

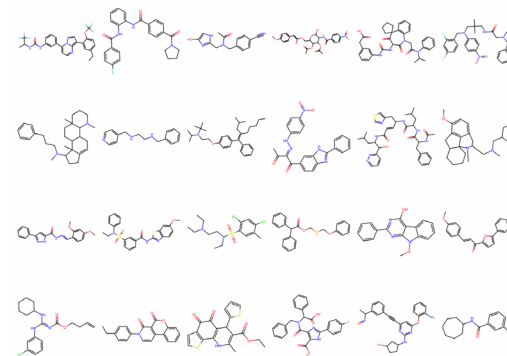
- Trustworthiness
- Sustainable Generative AI
  - *Lean* Language Models
  - Low-energy consumption (brain: ~0.5 kWh/day vs. 16xGPU: ~170 kWh/day)
  - Low-precision



# Challenges



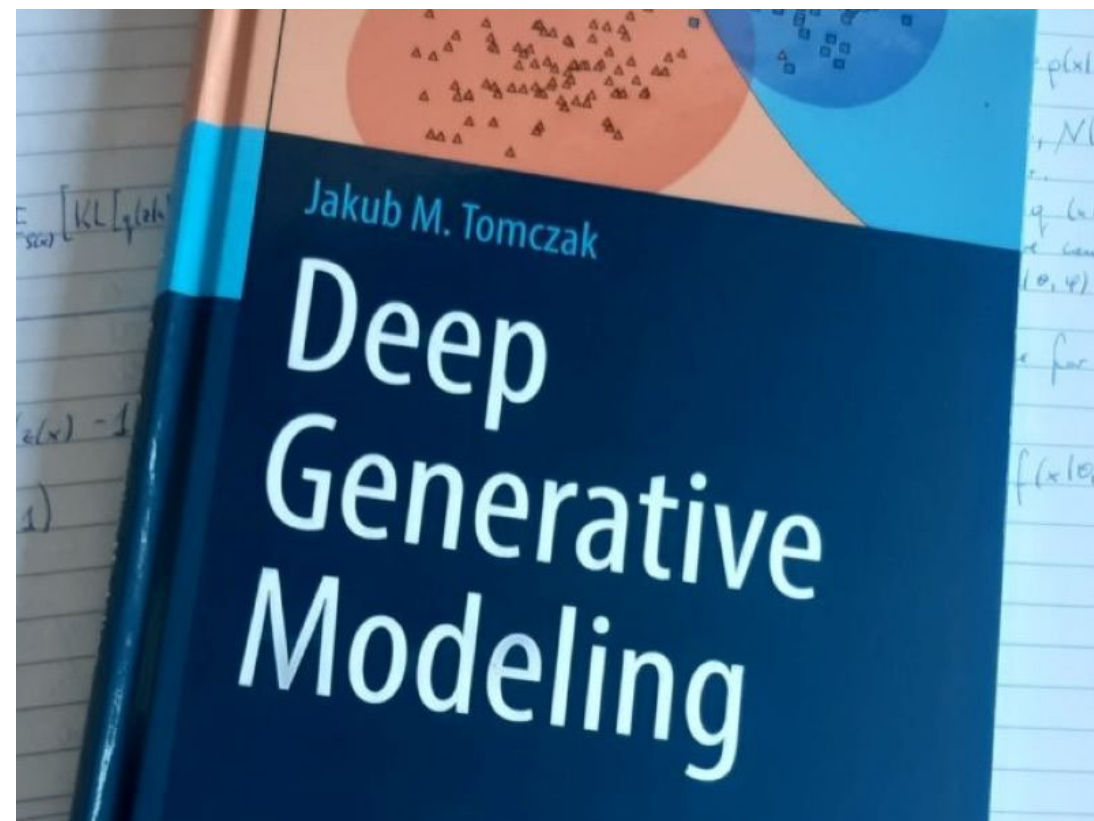
- Trustworthiness
- Sustainable Generative AI
  - *Lean* Language Models
  - Low-energy consumption (**brain: ~0.5 kWh/day vs. 16xGPU: ~170 kWh/day**)
  - Low-precision
- Useful applications:
  - From benchmarks to Life Sciences, Molecular Sciences, Ecology, ...



# Take-aways

# Take-aways

- **Generative AI has shown that almost everything is possible!**
- **Many open research questions**
- **Sustainability!**
- No killing robots (we are not there yet 😊) but think of **regulations!**
  - **Trustworthiness**



(Always remember about shameless self-promotion)



# Thank you!

Contact: [jmk.tomczak@gmail.com](mailto:jmk.tomczak@gmail.com)

## Generativ/e

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

## Amsterdam AI Solutions

Amsterdam AI Solutions: <https://amsterdamaisolutions.com/>