Welcome!

The program will begin soon.

You will not hear audio until we begin.

Unlocking the Potential of Generative Artificial Intelligence in the Oil and Gas Industry

Dr. Serkan Dursun
Leader
Artificial Intelligence Center of Excellence
July 7, 2023

Webinars

SPE Webinars are Sponsored by:

SPE

SPE Foundation

SPE International

Agenda

Webinars

- What is Generative AI
- Evolution of Generative AI
- Generative AI Landscape
- Transformer Architecture
- Unimodal and Multimodal Gen AI
- Responsible AI Practice
- Gen AI Applications in Oil and Gas Industry

SPE International

What is Generative Machine Learning

Supervised
- Data: (x,y) x:input, y:label
- Goal : Learn a function to map x to y
- Application: Classification, Regression

Unsupervised
- Data: x only
- Goal: learn underlying patterns/structure in the data
- Application: Clustering, Dimension Reduction

	Discriminative model	Generative model
Goal	Directly estimate $P(y\|x)$	Estimate $P(x\|y)$ to then deduce $P(y\|x)$
What's learned	Decision boundary	Probability distributions of the data
Illustration	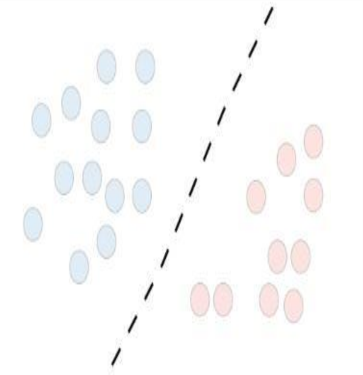	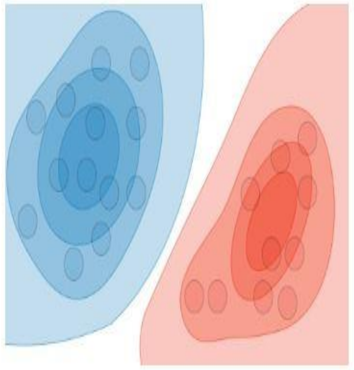
Examples	Regressions, SVMs	GDA, Naive Bayes

Source: CS 229 - Supervised Learning Cheatsheet (stanford.edu)

- A discriminative machine learning model learns parameters that maximize the conditional probability P(Y|X),
- A generative machine learning model learns parameters that maximize the joint probability P(X,Y).

Evolution of Generative AI

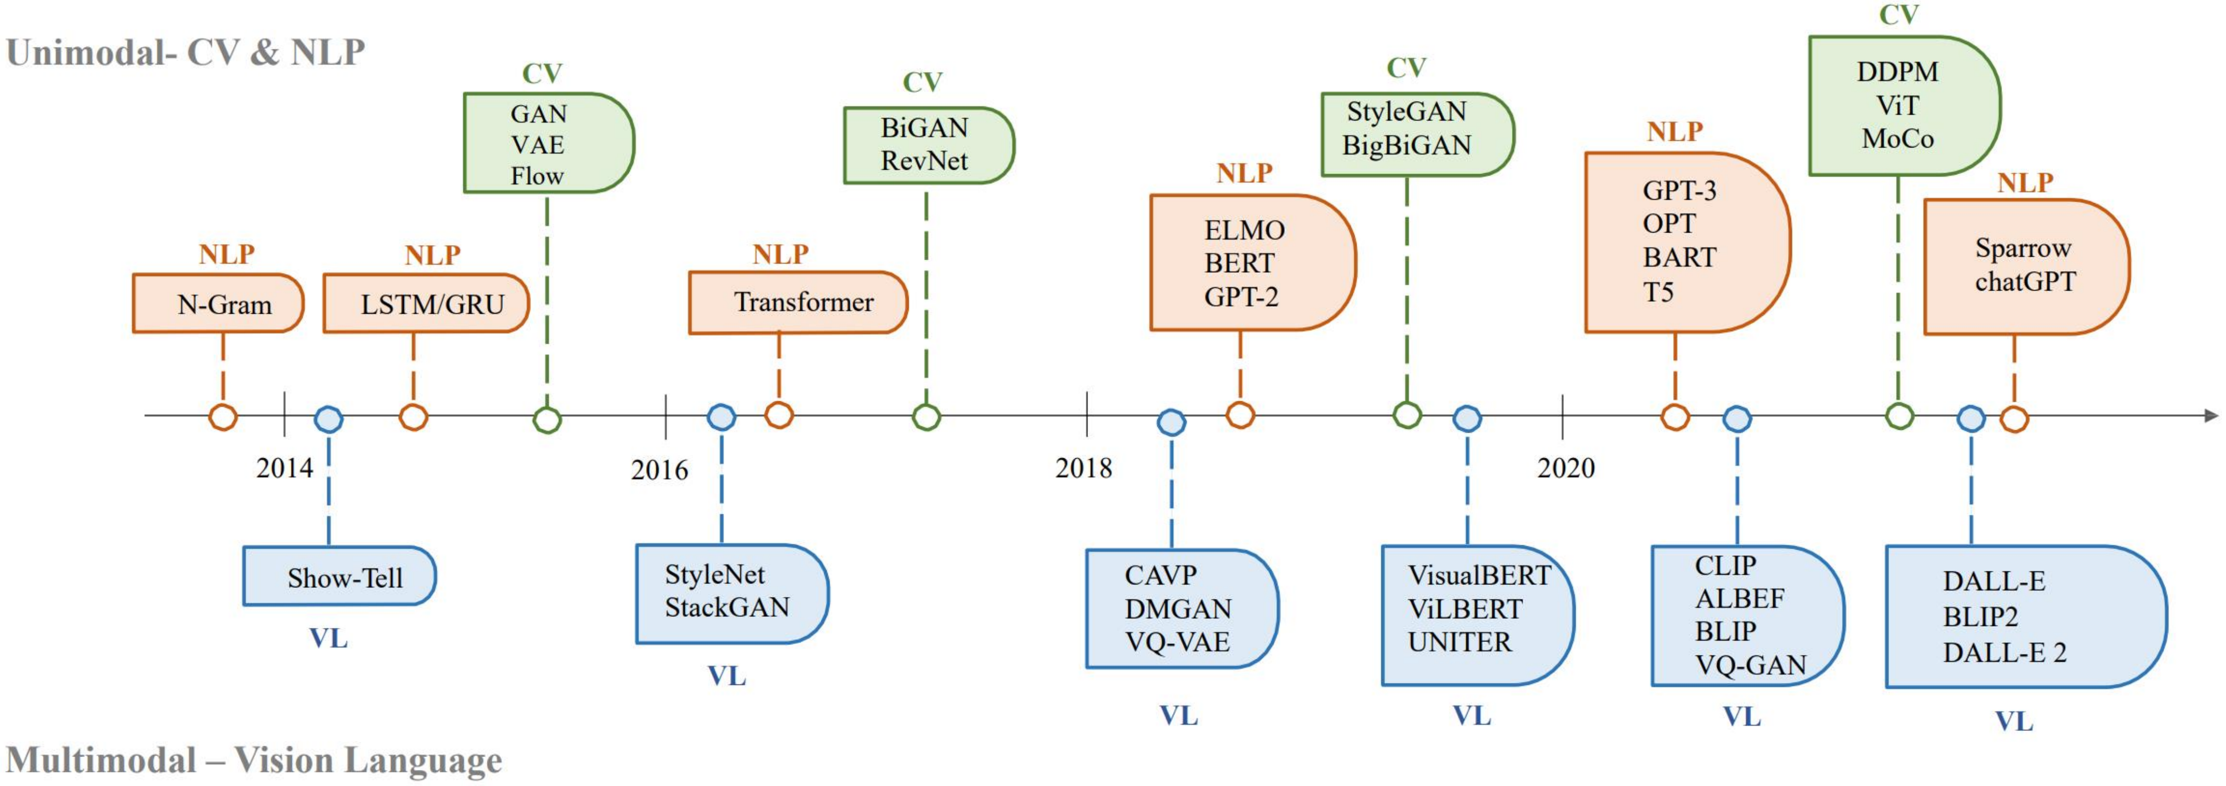

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Generative AI Landscape

Unimodal

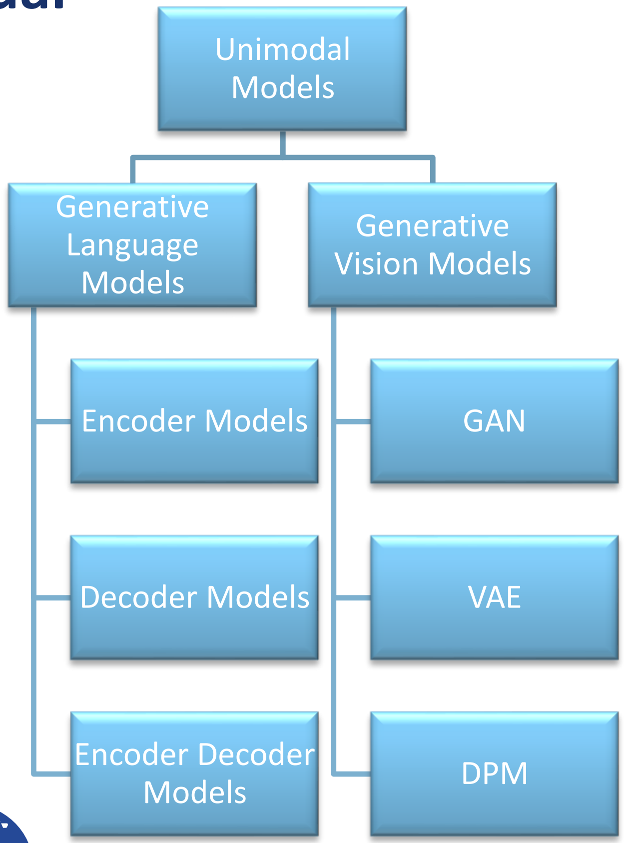

Multimodal

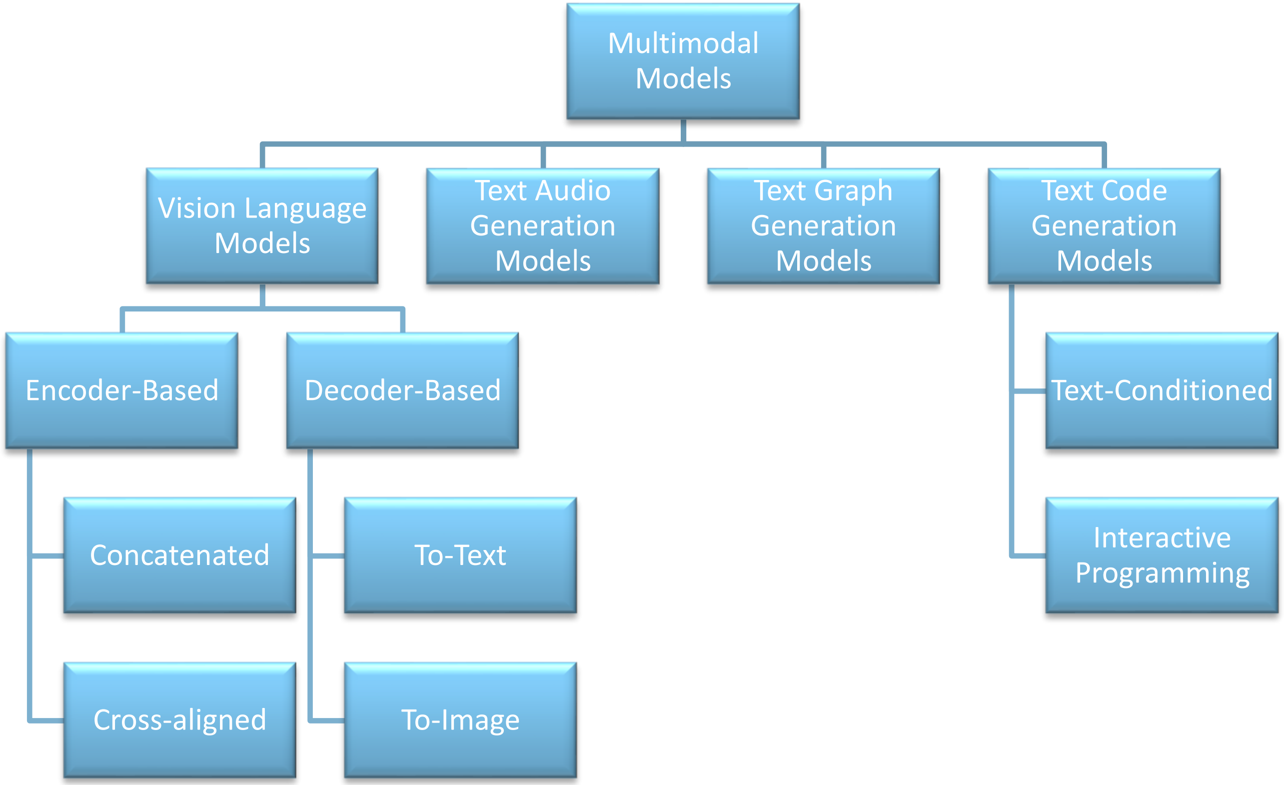

Generative AI Scenarios

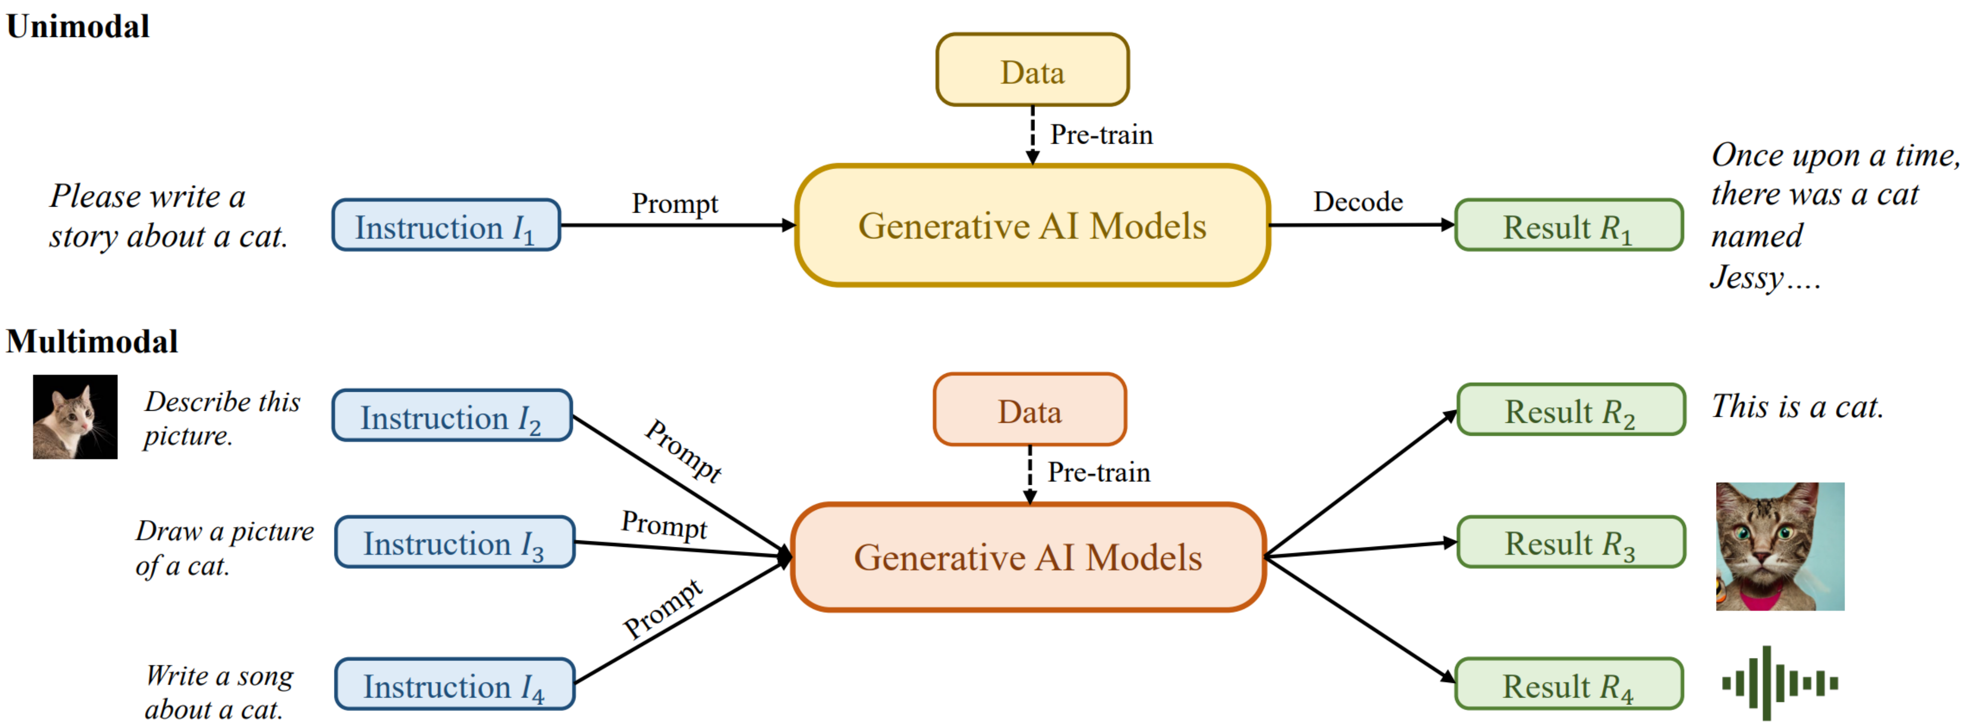

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Transformer Model Architecture

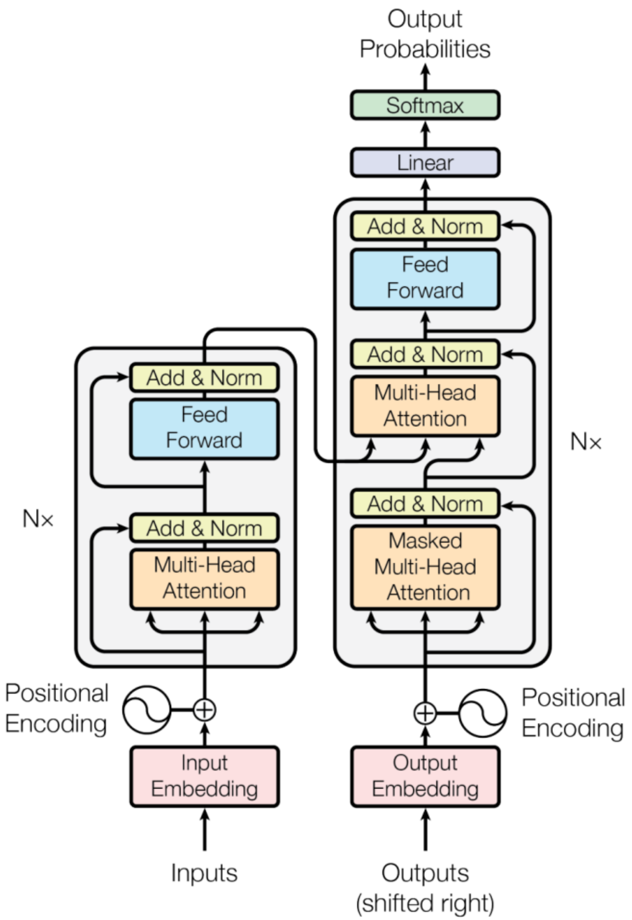

- Embedding
- Positional Encoding
- Encoder
- Decoder
- Multi-head Attention

Source: [1706.03762] Attention Is All You Need (arxiv.org)

Multimodal
Vision Language Generation

- Vision Language Encoders
  - Concatenated Encoders
    - by concatenating the embeddings from single encoders.
    - VisualBERT
  - Cross-aligned Encoders
    - Two tower for each modality and learn joint representation space by using a cross-modality encoder.
    - LXMERT
    - ViLBERT
- Vision Language Decoders
  - To-Text Decoders
  - To-Image Decoders

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Vision Language Encoders

- Concatenated Encoder

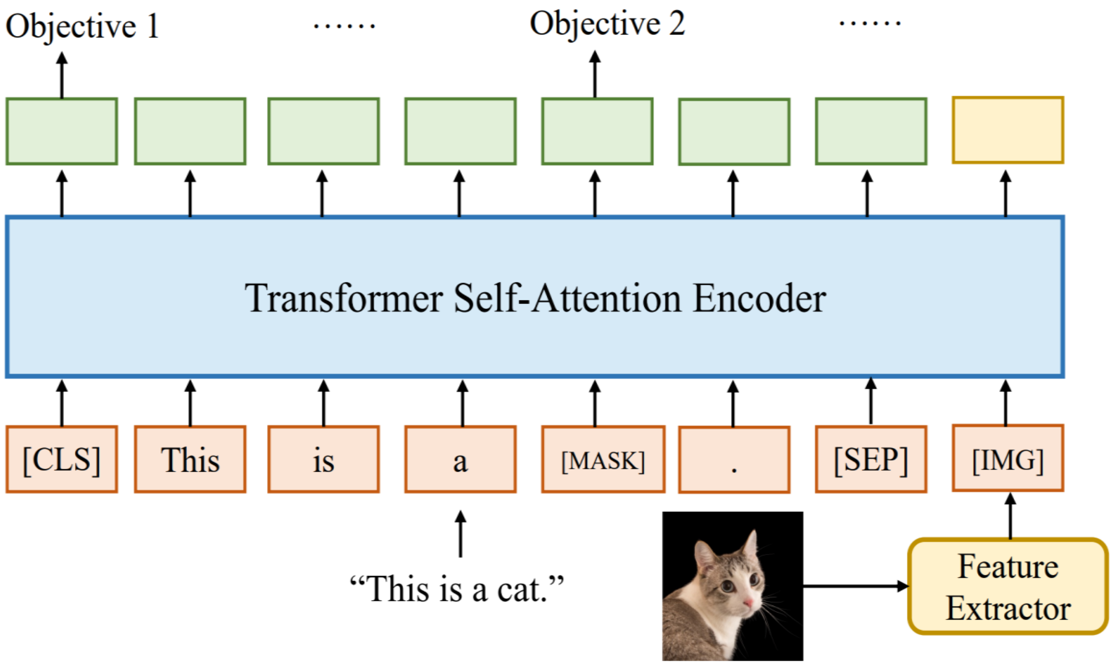

- Cross-aligned Encoder

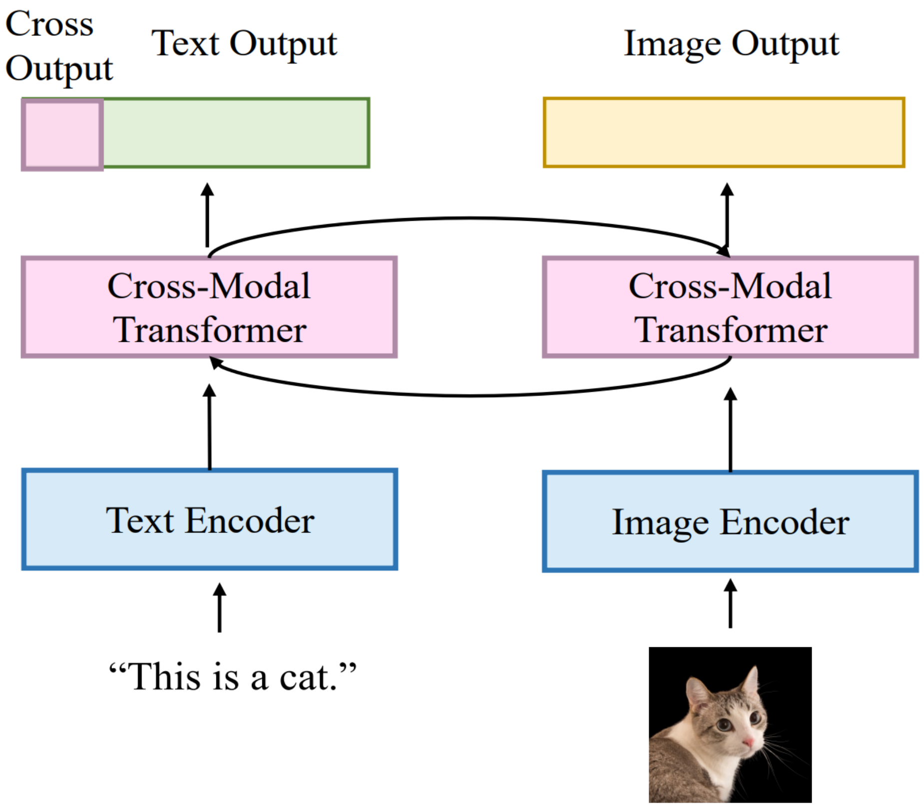

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Vision Language Decoders

- To Text Decoders
  - Jointly-trained Decoders
  - Frozen Decoders
- To Image Decoders
  - GAN-based Decoders
    - StackGAN:two stage
    - AttnGAN: Attn+StackGAN
  - Diffusion-based Decoders
    - GLIDE

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Vision Language Decoders

- Jointly Trained Decoders
- Frozen Decoders

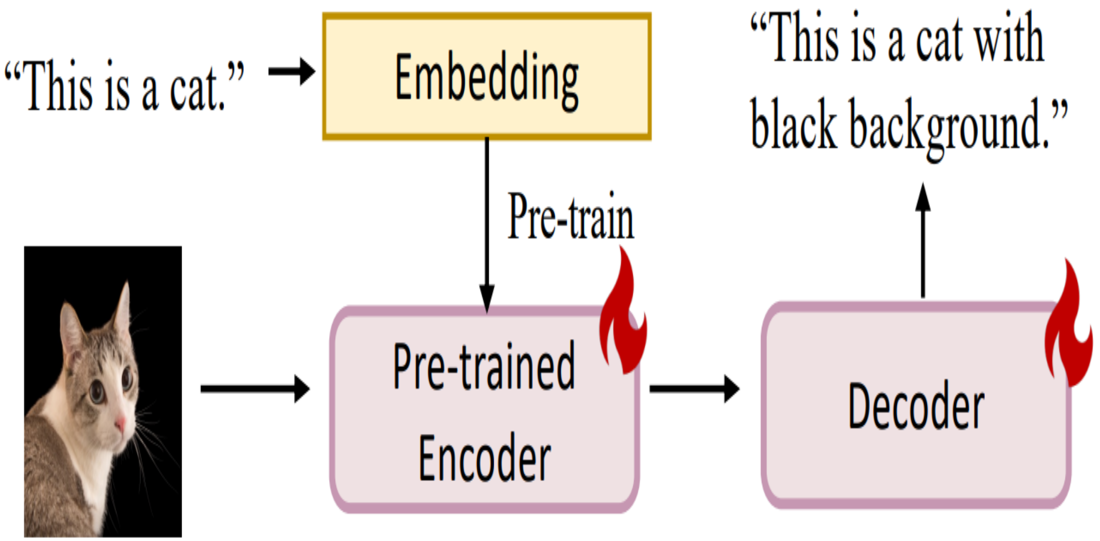

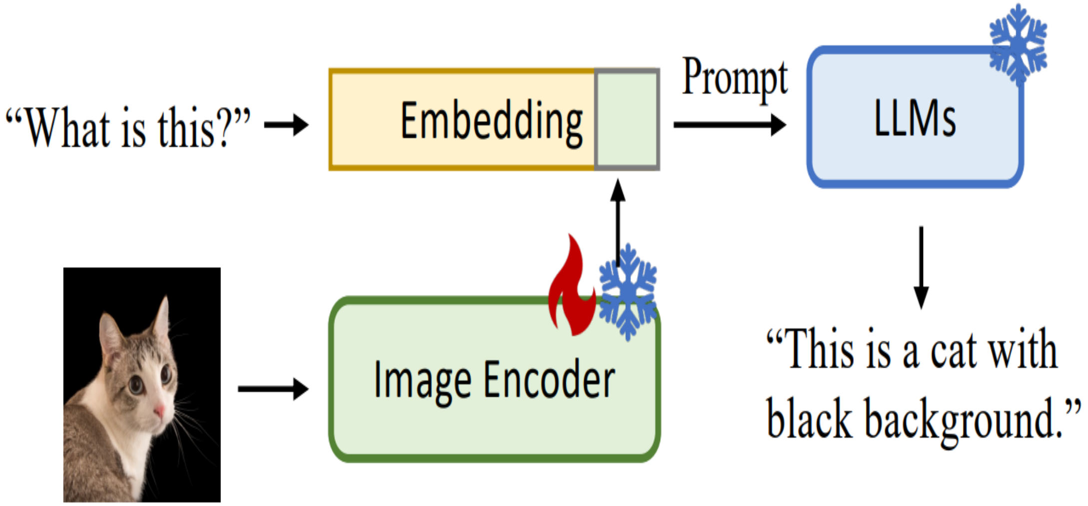

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Multimodal
Text Code Generation

- Text-conditioned Programming
  Code Generation
  - CodeBERT
  - CuBERT
  - CodeT5
  - PLBART

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Selected Generative AI Models

- Large Language Models (LLMs)
- Generative Adversarial Networks (GANs)
- Variational Autoencoders (VAEs)
- Diffusion Probabilistic Models (DPMs)

What is a Language Model ?

Webinars

A Language Model is a probabilistic model that defines a probability distribution over the sequence of words to predict the next ______ in a sentence.

SPE International ®

Evolution of Large Language Models

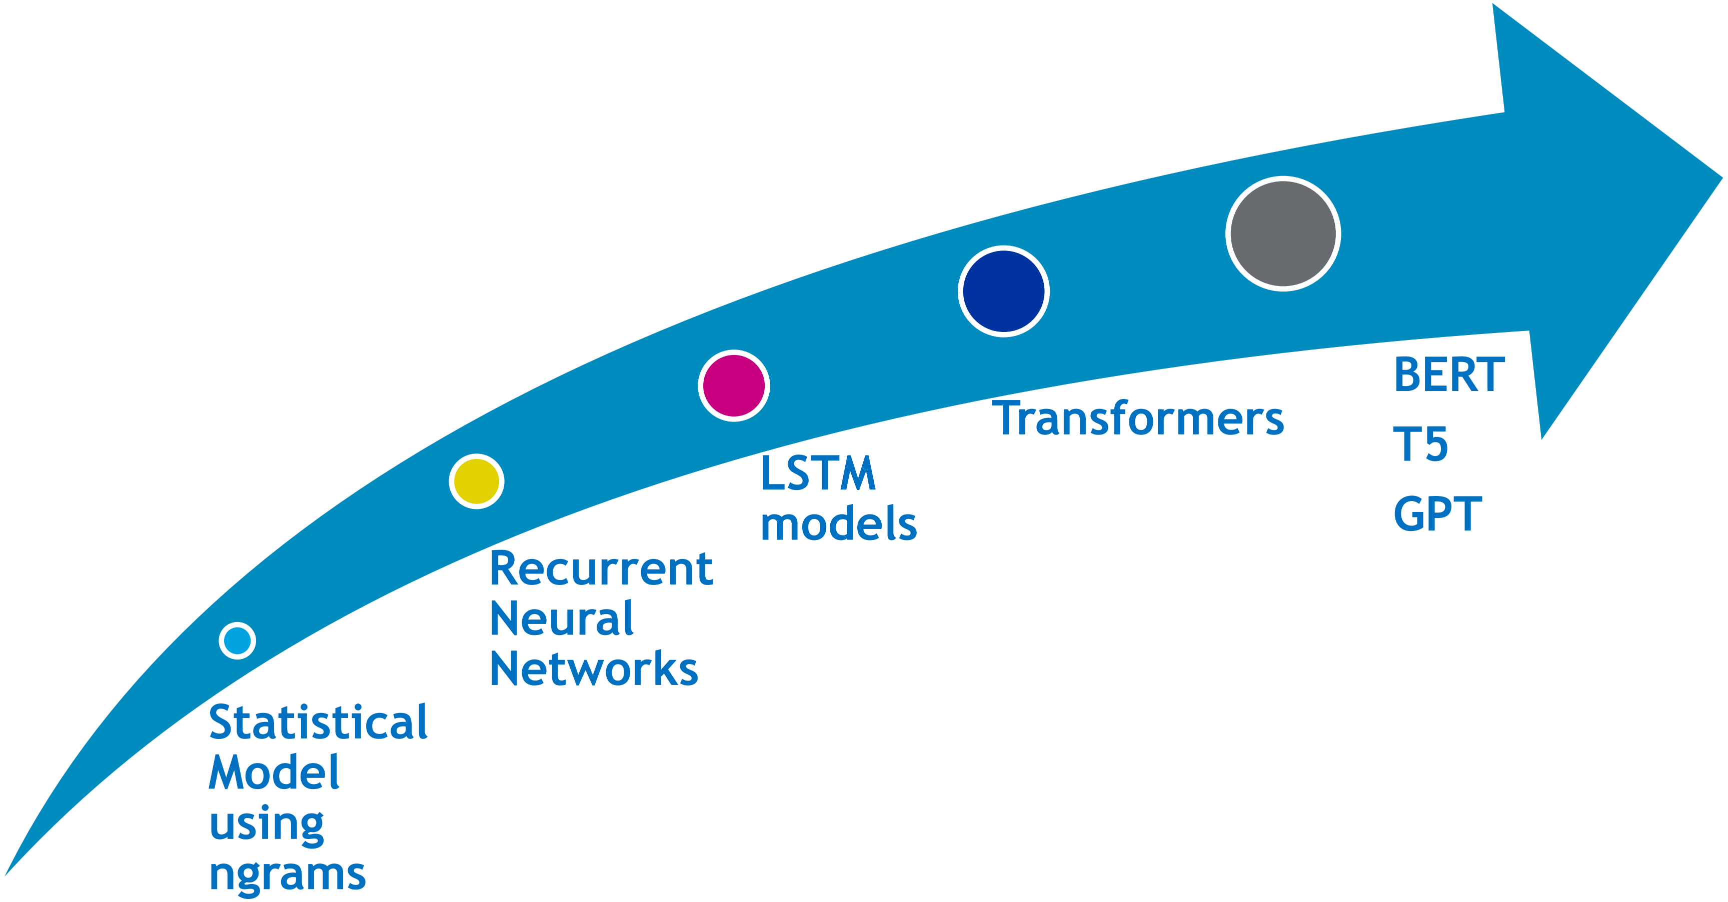

Large Language Models Landscape

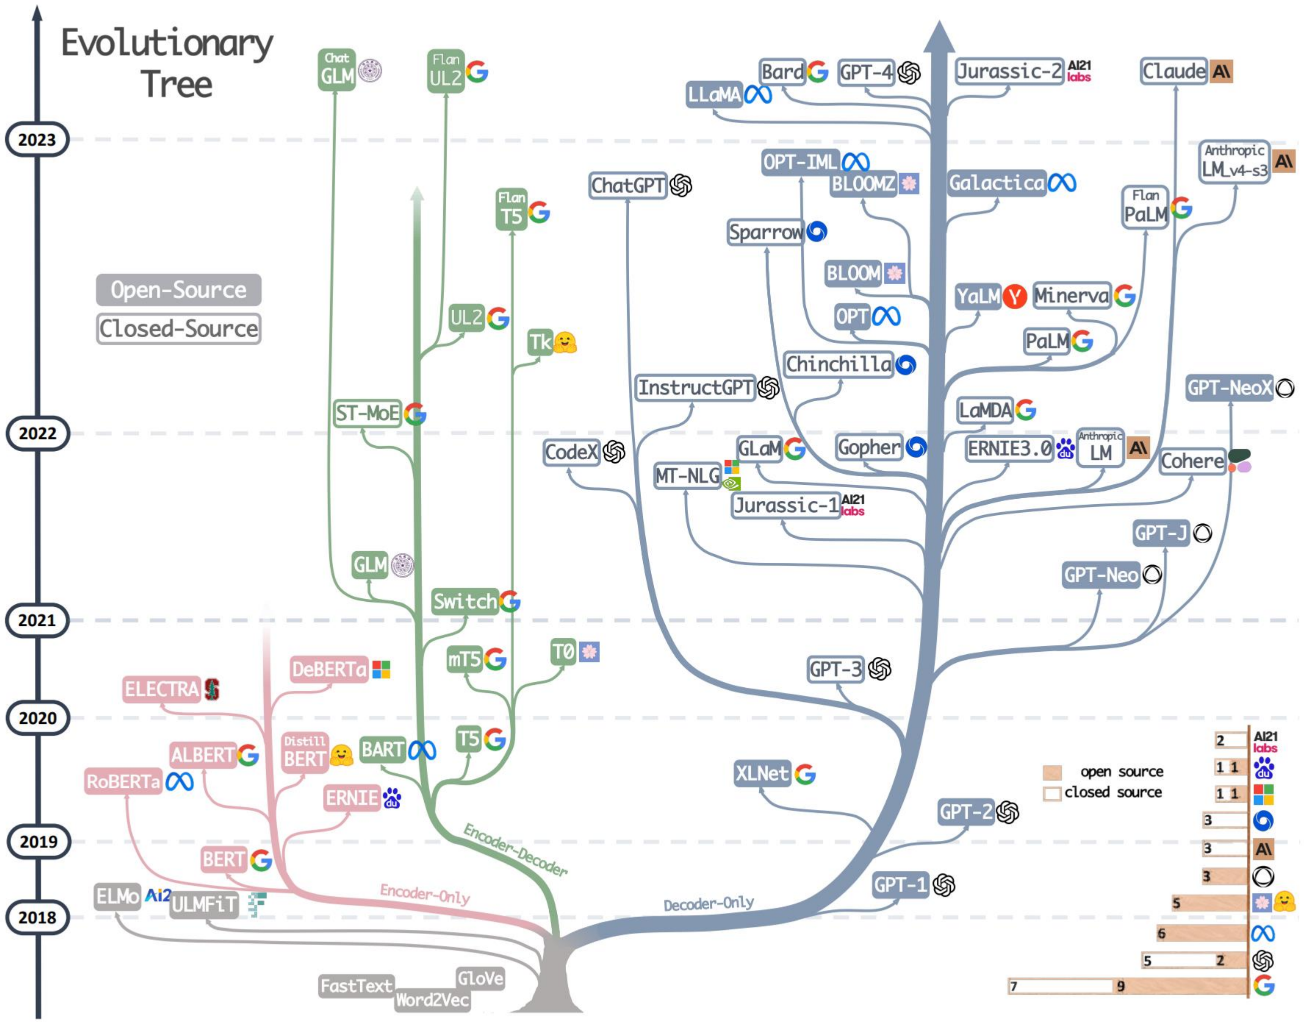

Source: [2304.13712] Harnessing the Power of LLMs in Practice: A Survey on ChatGPT and Beyond (arxiv.org)

Training Lifecycle of LLM

- **Pre-training**
  - Data Collection
  - Data Preprocessing
  - Architecture
  - Model Training
- **Adaptation Tuning**
  - Instruction Tuning
  - Alignment Tuning
  - Efficient Tuning
- **Utilization**
  - In-context Learning
  - Zero-shot prompt
  - Few-shot prompt
  - Chain of Thought prompt
- **Performance Evaluation**
  - Language Generation
  - Knowledge Utilization
  - Complex Reasoning

Webinars

SPE International

What is ChatGPT?

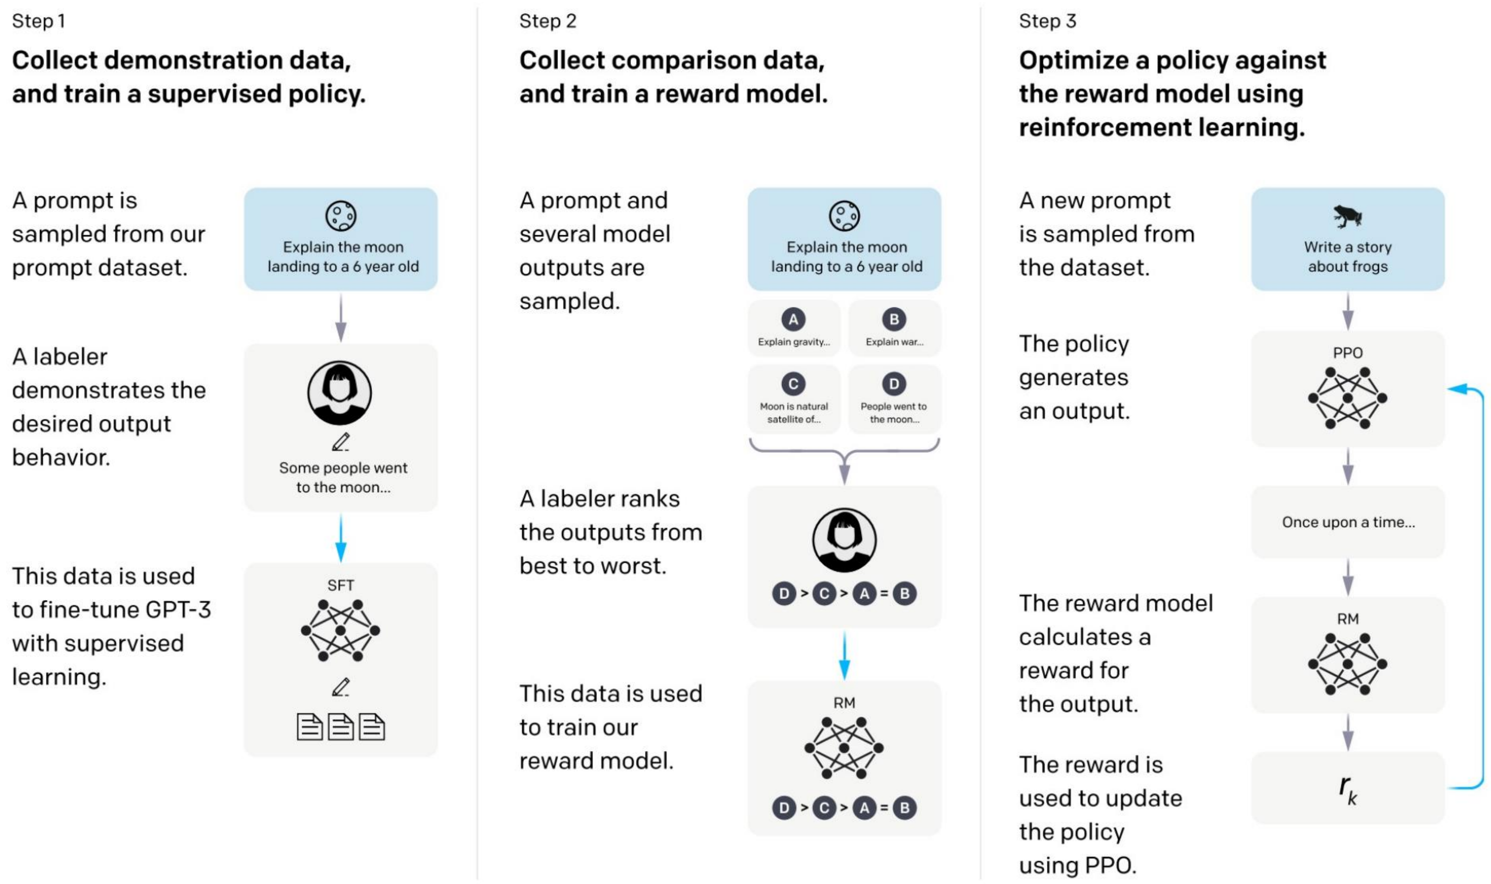

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Prompt Design and Engineering

- Prompt Engineering is a technique for improving the performance of LLMs by providing them with additional input in the form of prompts or instructions.
- The idea behind Prompt Engineering is to guide the model towards generating more relevant and coherent text by providing it with context-specific information.
- Prompt Engineering can be used to improve the performance of LLMs on specific tasks such as question answering or summarization.

- Zero-shot Prompting
  - No examples
- Few-shot Prompting
  - One or more examples
- Chain of Thought Prompting
  - Steps for guidance

Unimodal Image Generative AI Modeling

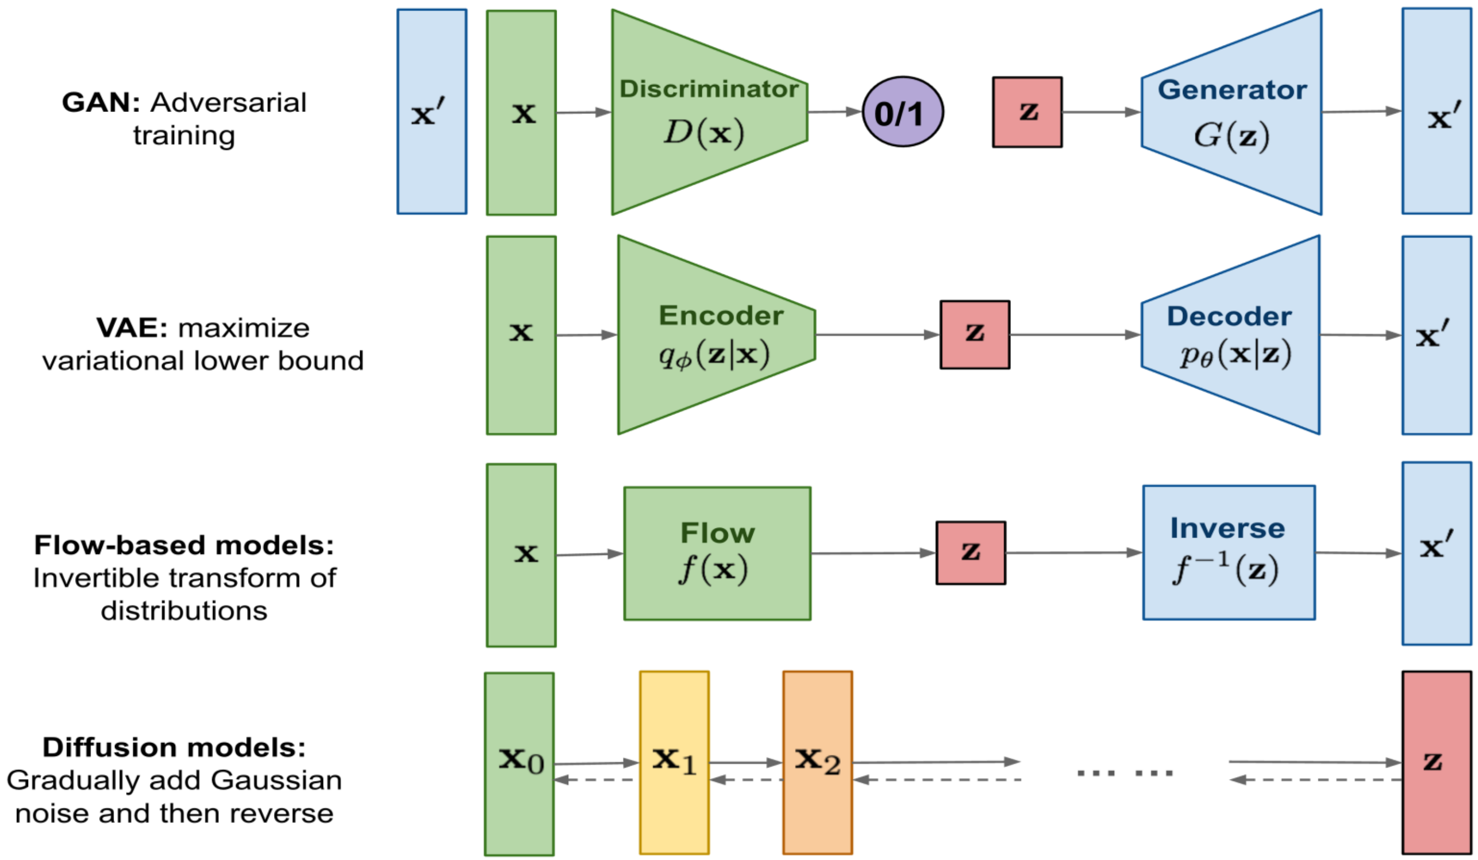

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT
What are Diffusion Models? | Lil'Log (lilianweng.github.io)

Generative Adversarial Networks

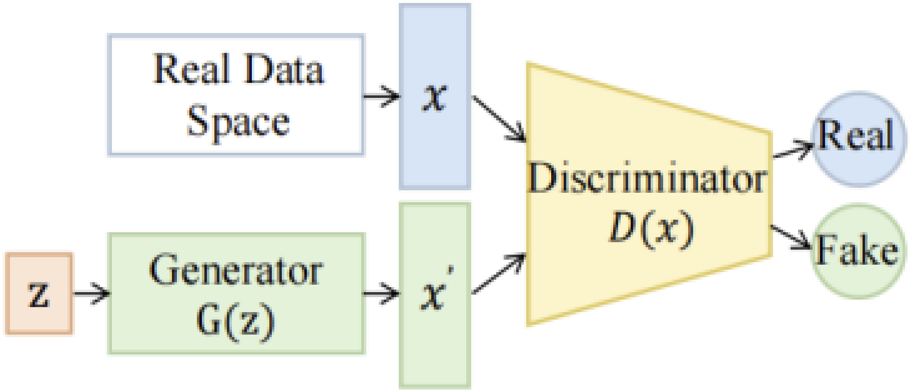

- Generator
- Discriminator

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Variational Autoencoders

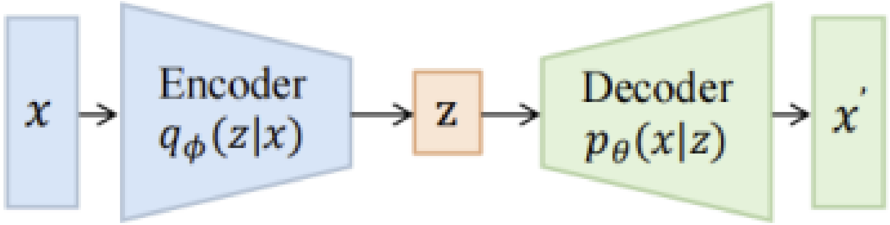

- Variational Inference
- Mean and Variance
- Latent Space

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Diffusion Probabilistic Models

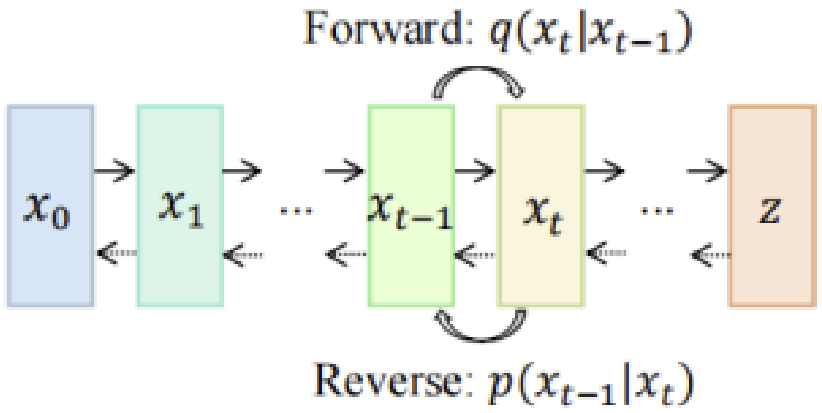

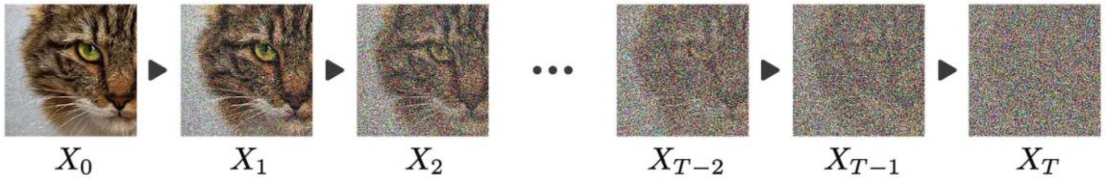

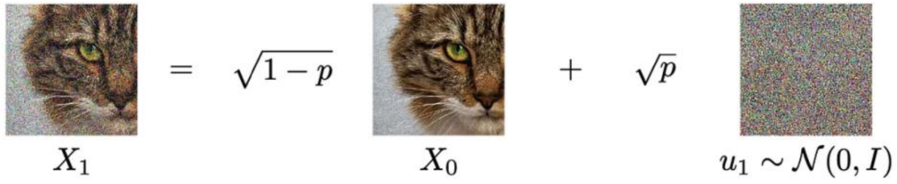

- Diffusion Mechanism
- Parametrized as Markov Chain
- Gaussian Noise
- Forward Process
- Reverse Process

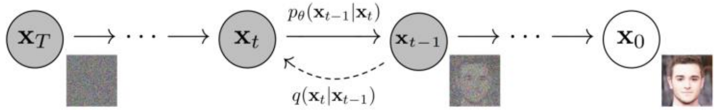

Source: Understanding Diffusion Probabilistic Models (DPMs) | by Joseph Rocca | Towards Data Science

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

CLIP: Contrastive Language-Image Pre-training

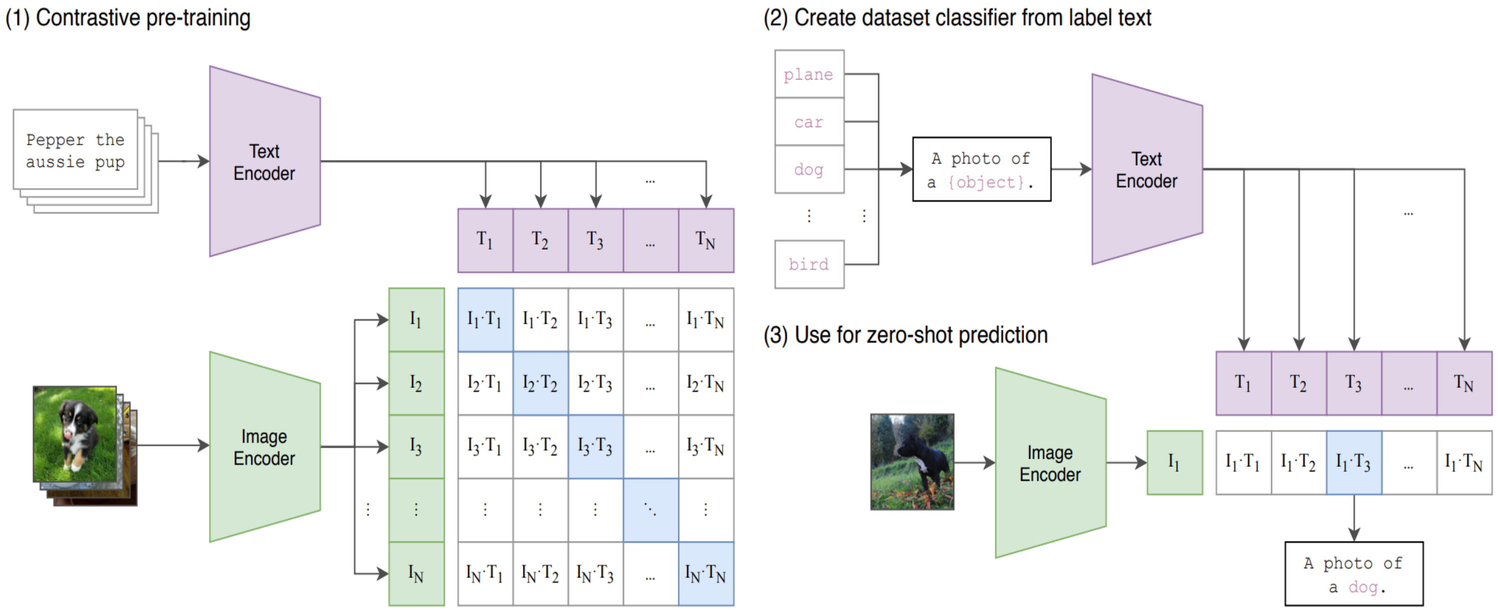

CLIP jointly trains an image encoder and a text encoder to predict the correct pairings of a batch of (image, text) training examples. At test time the learned text encoder synthesizes a zero-shot linear classifier by embedding the names or descriptions of the target dataset's classes.

Source: Learning Transferable Visual Models From Natural Language Supervision (arxiv.org)

Imagen

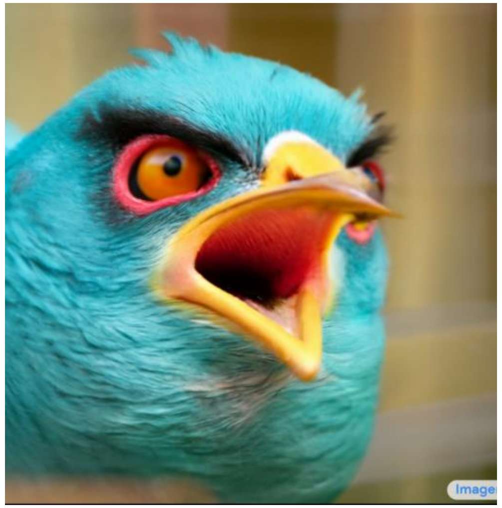

extremely angry bird

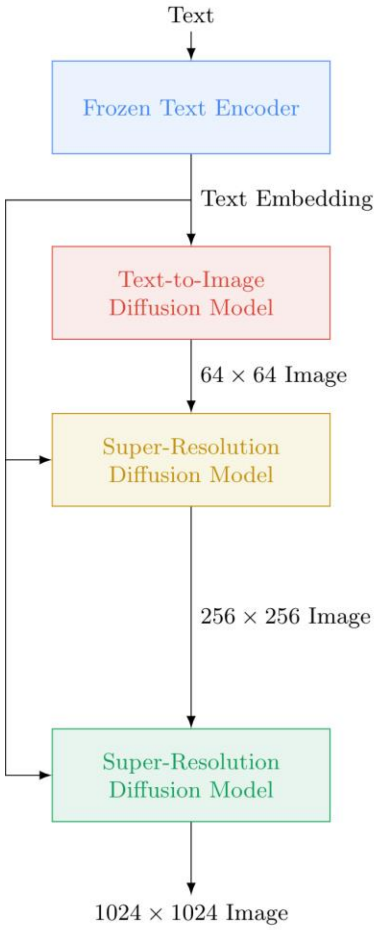

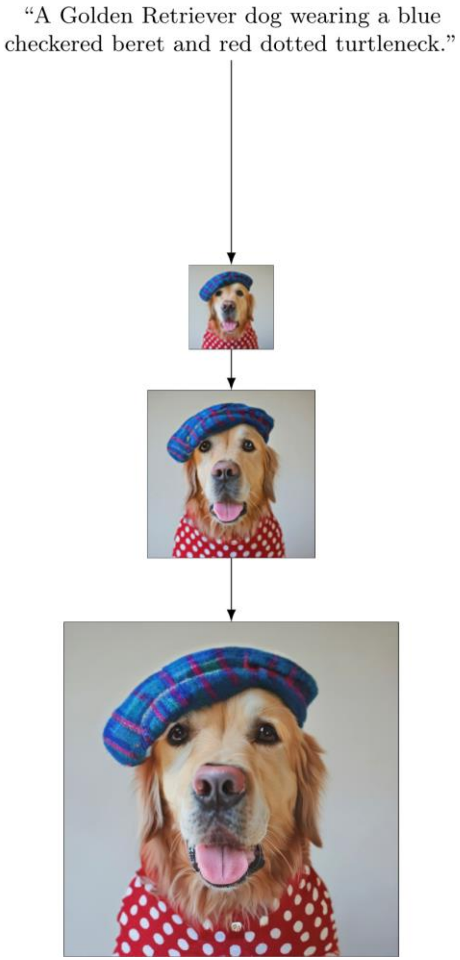

- Imagen uses a large frozen T5-XXL encoder to encode the input text into embeddings.
- A conditional diffusion model maps the text embedding into a 64×64 image.
- Imagen further utilizes text-conditional super-resolution diffusion models to upsample the image
- 64×64→256×256 and 256×256→1024×1024.

Source: https://imagen.research.google/

DALL-E 2

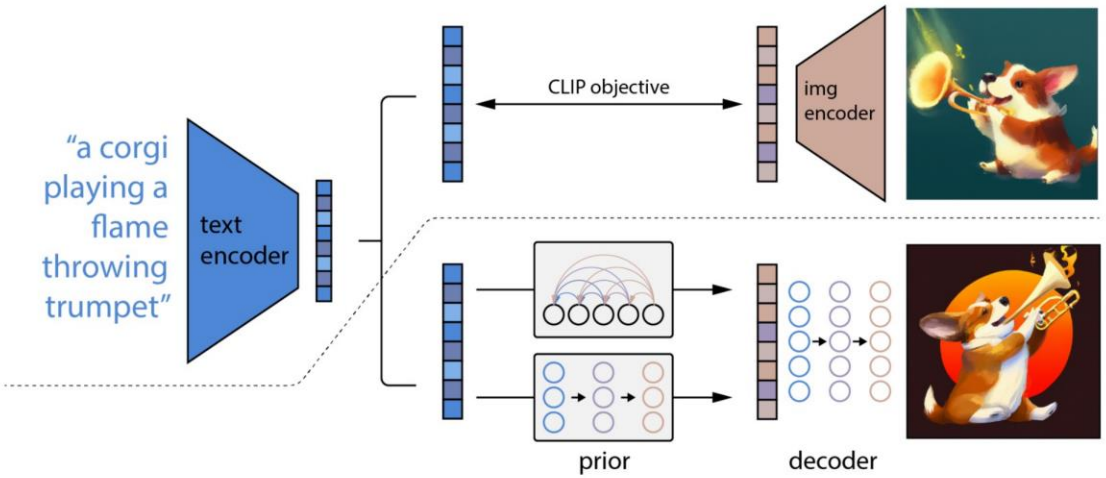

- The model structure of DALL-E-2. Above the dotted line is the CLIP pre-training process, which aims to align the vision and language modalities. And below the dotted line is the image generation process. The text encoder accepts an instruction and encodes it into a representation, then the prior network and diffusion model decodes this representation to generate the final output.

Source: A Comprehensive Survey of AI-Generated Content (AIGC): A History of Generative AI from GAN to ChatGPT

Responsible Artificial Intelligence

Webinars

Transparency:

The ability to explain how an AI system works and what data it uses to make decisions. This includes making the decision-making process clear and understandable to users and stakeholders.

Fairness:

The principle that AI systems should not discriminate against individuals or groups based on their race, gender, age, or other characteristics. This requires careful consideration of the data used to train AI systems and the algorithms used to make decisions.

Accountability:

The principle that AI systems should be designed to take responsibility for their actions and decisions. This includes providing clear lines of responsibility for the development and deployment of AI systems.

Reliability:

The principle that AI systems should be designed to work consistently and accurately over time. This requires careful testing and validation of AI systems before they are deployed.

Representation and Equitable:

The principle that AI systems should be designed to represent all individuals and groups fairly and equitably. This requires careful consideration of the data used to train AI systems and the algorithms used to make decisions.

Explainable and Interpretable:

The principle that AI systems should be designed to provide clear explanations of their decisions and actions. This includes providing users with information about how decisions are made and what data is used to make them.

SPE International ®

Responsible Generative Artificial Intelligence

Webinars

Transparency:	Fairness:	Accountability:	Reliability:	Representation and Equitable:	Explainable and Interpretable:
GAI systems should be designed to provide clear explanations of how they generate content and what data they use to do so. This includes making the decision-making process clear and understandable to users and stakeholders.	GAI systems should be designed to generate content that does not discriminate against individuals or groups based on their race, gender, age, or other characteristics. This requires careful consideration of the data used to train GAI systems and the algorithms used to generate content.	GAI systems should be designed to take responsibility for the content they generate. This includes providing clear lines of responsibility for the development and deployment of GAI systems.	GAI systems should be designed to generate content consistently and accurately over time. This requires careful testing and validation of GAI systems before they are deployed.	GAI systems should be designed to represent all individuals and groups fairly and equitably in the content they generate. This requires careful consideration of the data used to train GAI systems and the algorithms used to generate content.	GAI systems should be designed to provide clear explanations of how they generate content. This includes providing users with information about how decisions are made and what data is used to generate content.

SPE International

LLM Applications in Oil & Gas Industry

- NLP applications in the oil and natural gas industry | Second International Meeting for Applied Geoscience & Energy (seg.org)
  - Attribute Extraction
  - Semantic Information Retrieval
  - Query Intent Translation
  - User-Defined Classification of O&G Documents

- Geoscience Language Processing for Exploration | Abu Dhabi International Petroleum Exhibition and Conference | OnePetro
  - GeoBERT
  - Geoscience Q&A
  - Query based Summarization

GAN Applications in Oil & Gas Industry

- Geophysical model generation with generative adversarial networks | Geoscience Letters | Full Text (springeropen.com)
  - 2D density and stratigraphic models generated using StyleGAN2 and StyleGAN2 ADA

VAE Applications in Oil & Gas Industry

- Machine learning inversion of geophysical data by a conditional variational autoencoder | SEG International Exposition and Annual Meeting | OnePetro
  - Prior Geological knowledge is incorporated into geophysical inversions using conditional VAE.

Diffusion Applications in Oil & Gas Industry

- [2207.10451] Deep Diffusion Models for Seismic Processing (arxiv.org)
  - – DPM used to process the seismic data for demuiltiple, denoising and interpolation.

Questions