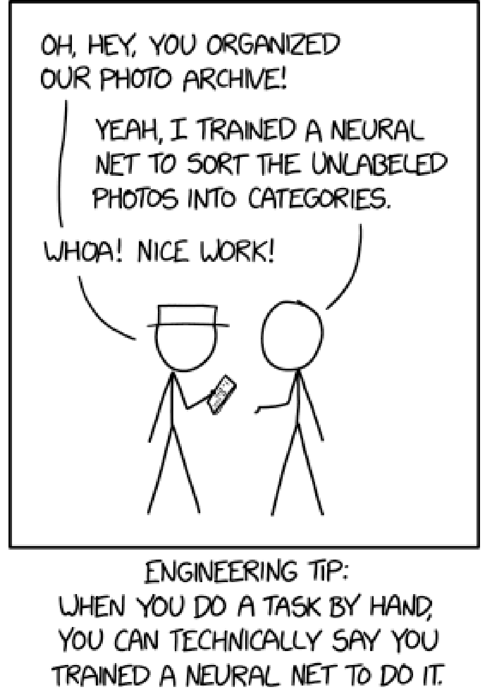

Diffusion models for generating structures

Sai Gautam Gopalakrishnan,[1] Keith Tobias Butler[2], Reshma Devi[1], and Pritam Ghosh[1]

[1]Materials Engineering, Indian Institute of Science

[2]Chemistry, University College London

saigautamg@iisc.ac.in[1]; k.t.butler@ucl.ac.uk[2]; https://sai-mat-group.github.io

Jan 8, 2025

Inverse materials design

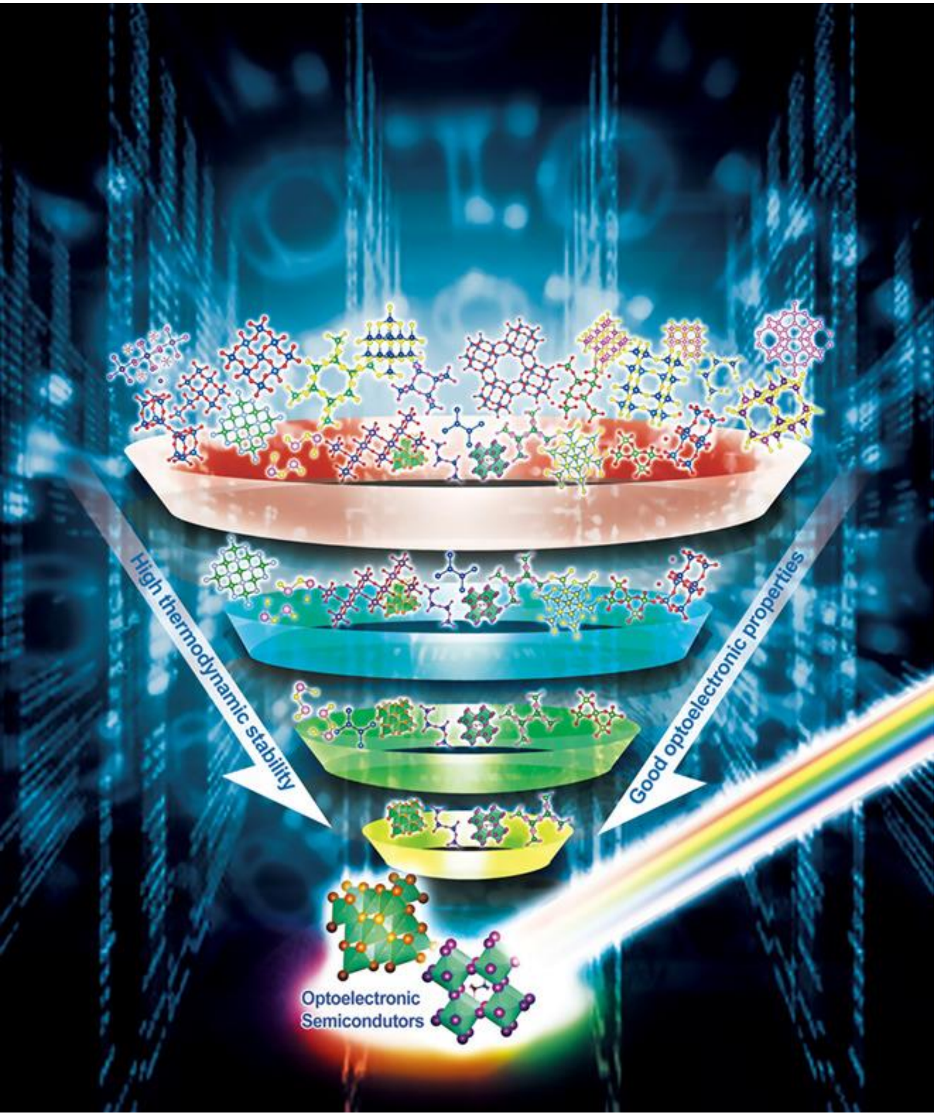

https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcms.1489

Property → Structure

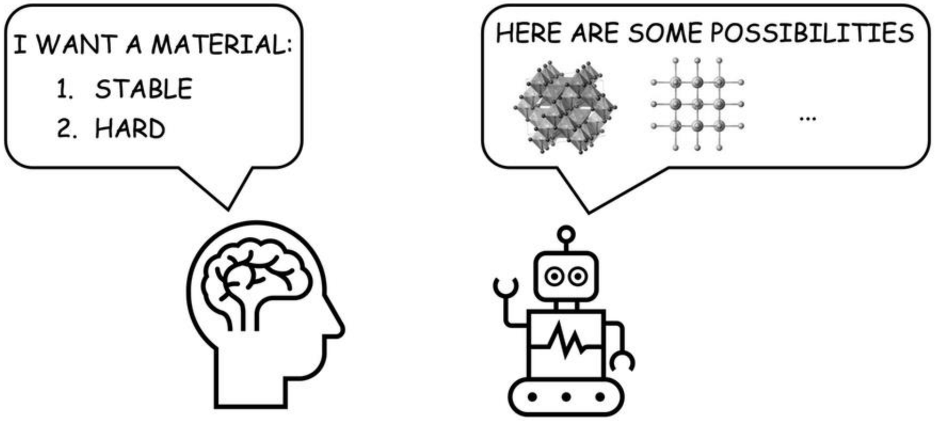

https://news.mit.edu/2022/new-way-perform-general-inverse-design-high-accuracy-0118

Generative models (classic)

- Step 1: encode a configuration (σ) into a latent/feature space (Z)
    - $Z = f(\sigma)$

- Critical info of any structure
    - Composition
    - Lattice parameters
    - Atomic positions
    - Use graph neural networks to obtain Z
    - Z can be mapped to labelled properties

- Step 2: decode configuration from latent space using a learnable function
    - $\sigma' = f'(Z)$
    - Introduces noise
    - Provides a probability distribution (compositions, lattice parameters, and positions)

- Step 3: generate configuration by sampling probabilities
    - $\sigma_{sampled} = p(Z)$
    - Given constraints on target properties, composition, and/or lattice geometry

Park et al., Matter 7, 2355-2367 (2024)

Variational autoencoder (VAE) models:~classic

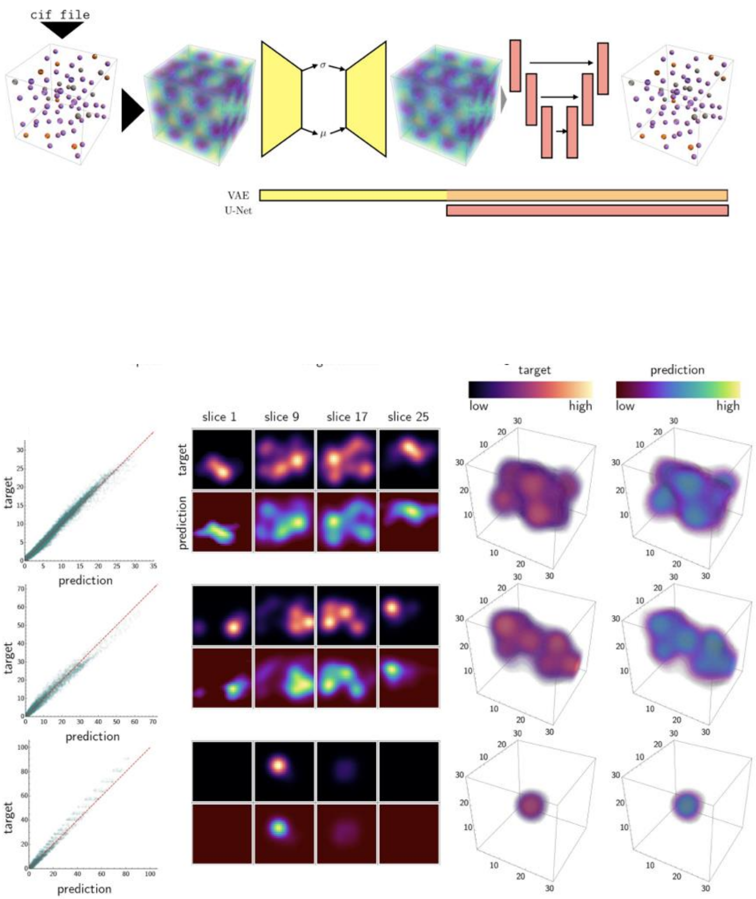

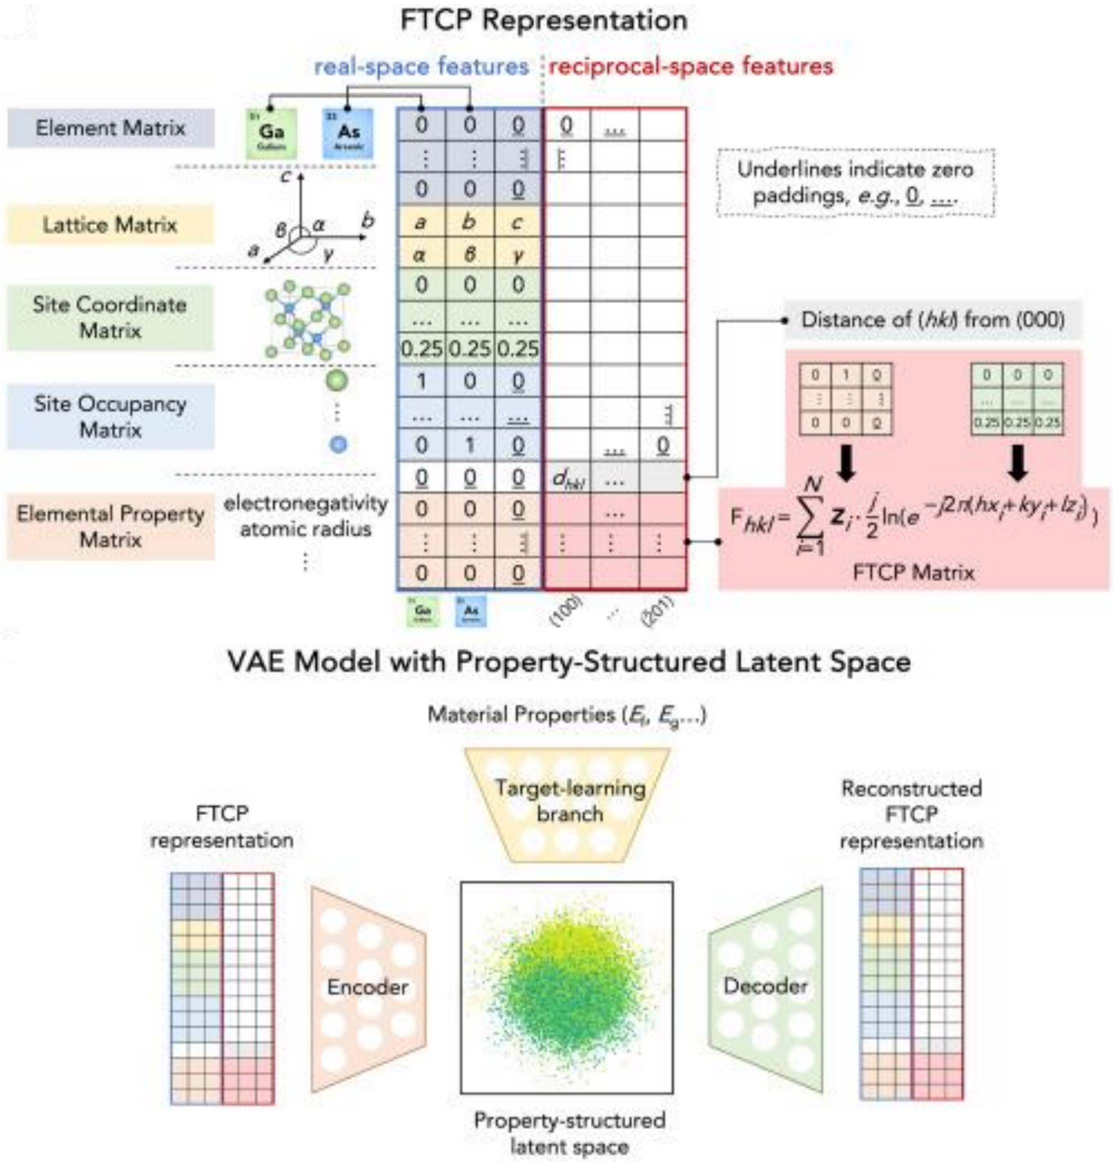

Hoffman et al., arXiv, 1909.00949v1 (2019)

Ren et al., Matter 5, 314-335 (2022)

Advancements in generative models

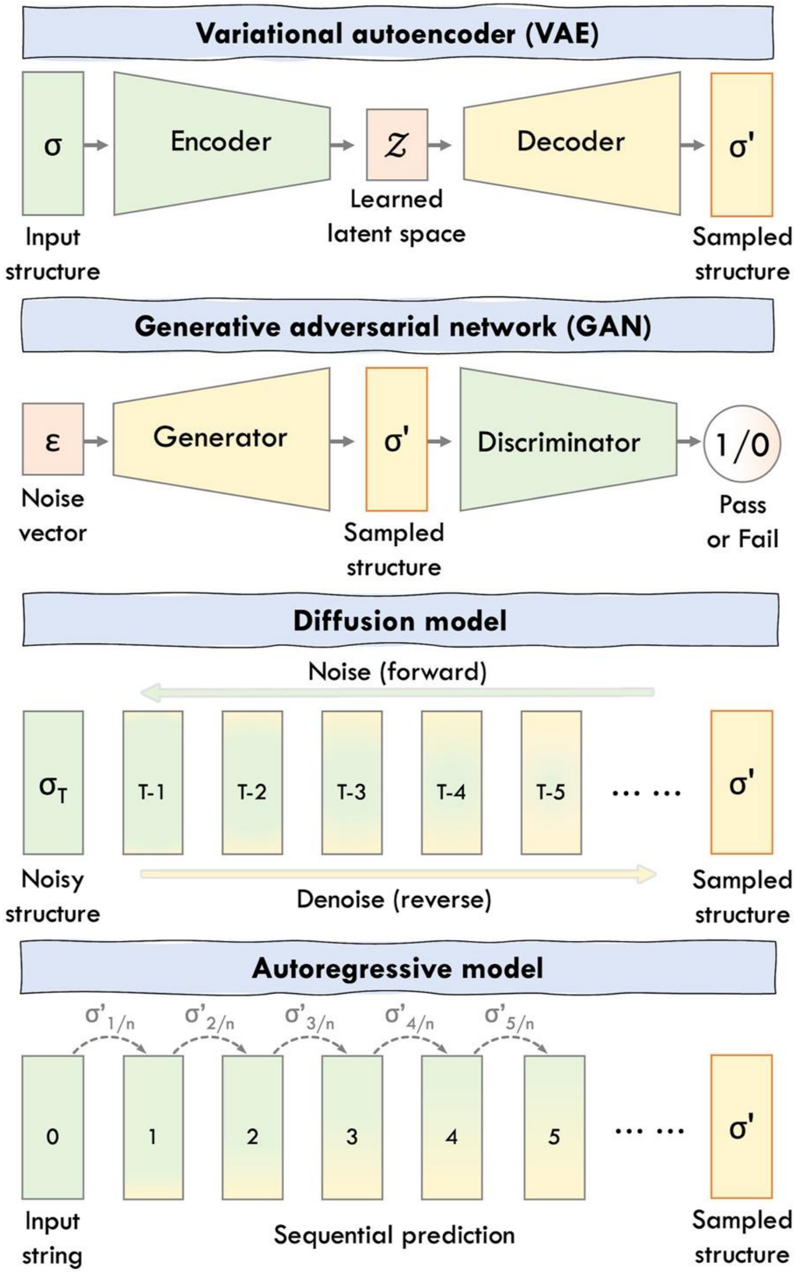

Classic

Generator confuses discriminator with synthetic data

- Beaten by diffusion models ☹

Progressive noise addition/removal

Sequential probability (language models)

Park et al., Matter 7, 2355-2367 (2024)

What is diffusion?

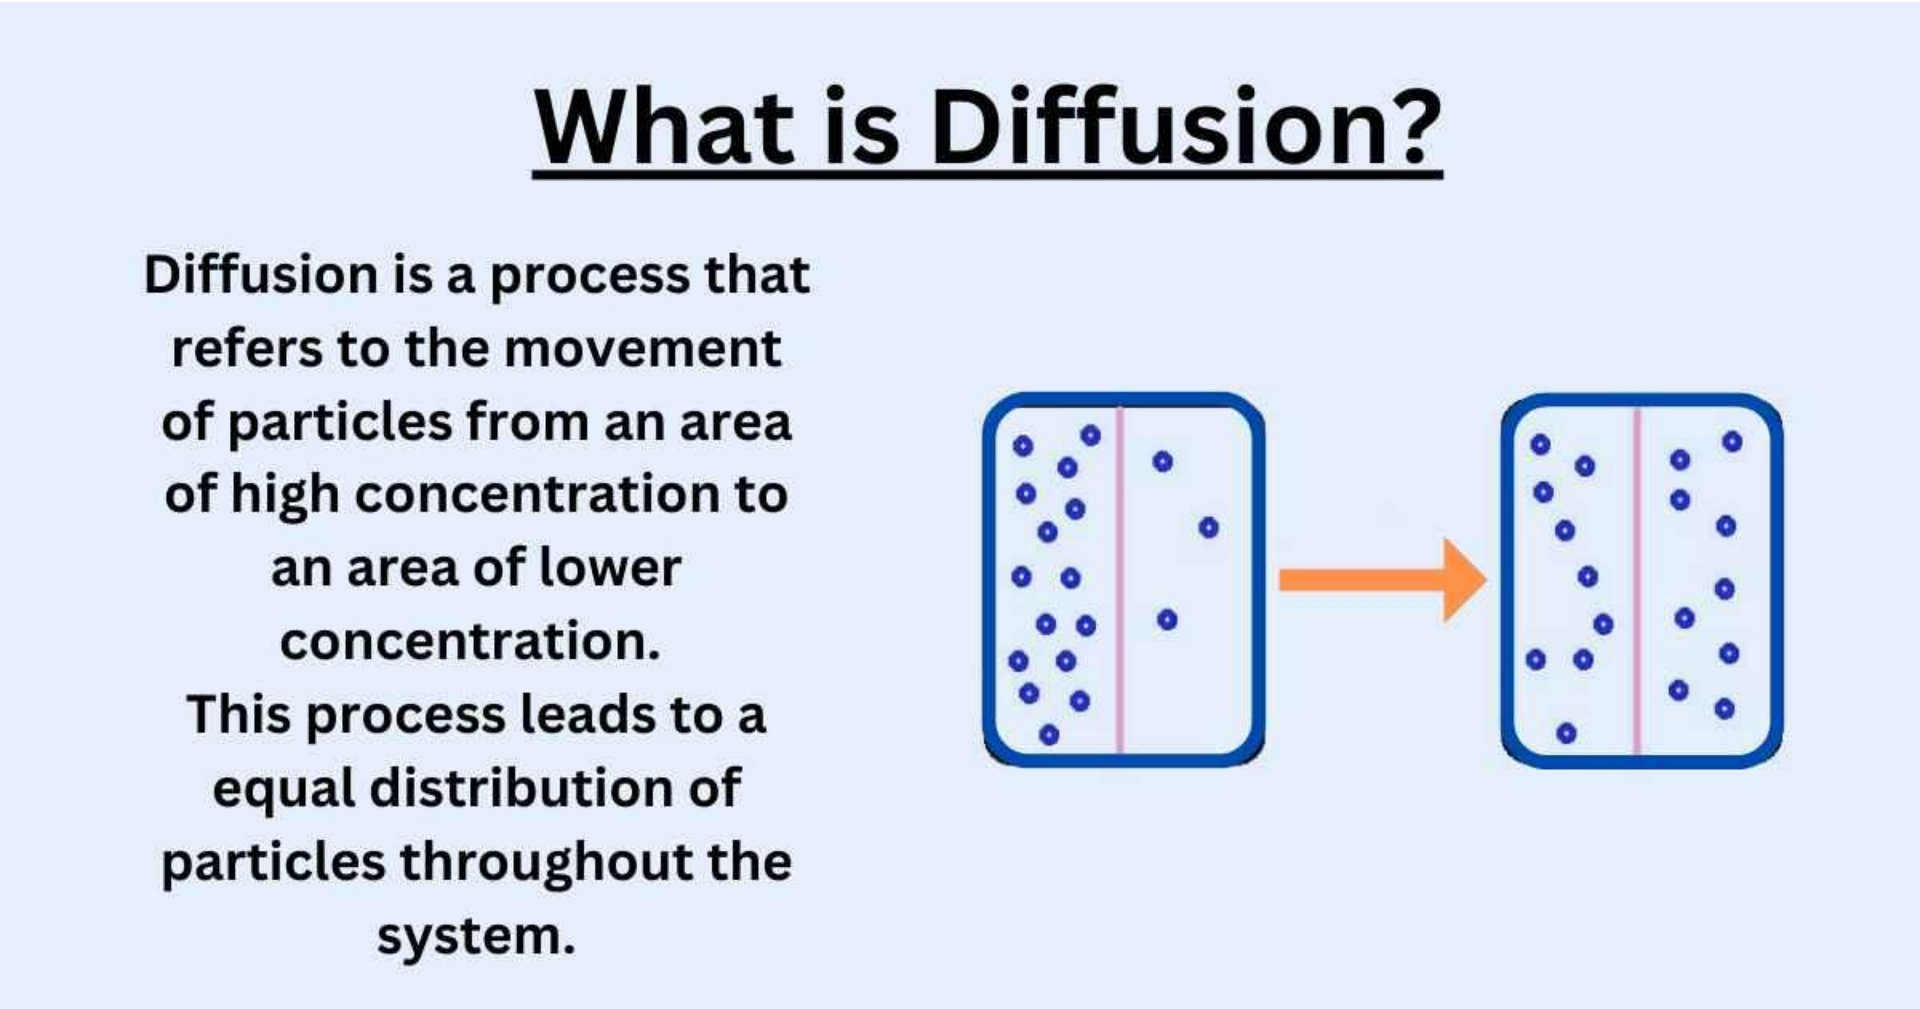

https://eduinput.com/what-is-diffusion/

It’s actually chemical potential gradients that drive…

What is diffusion?

Forward diffusion

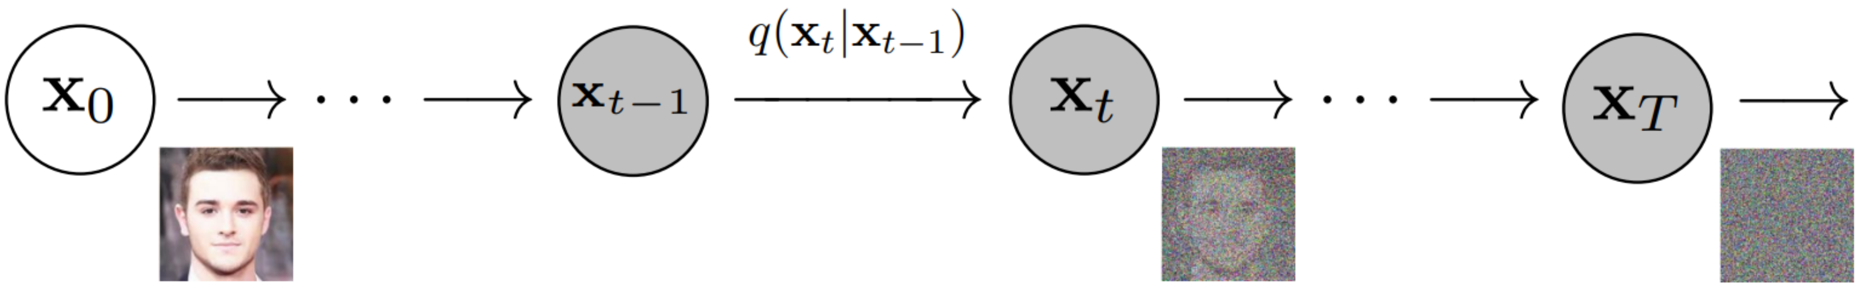

Gaussian noise added in a Markov chain

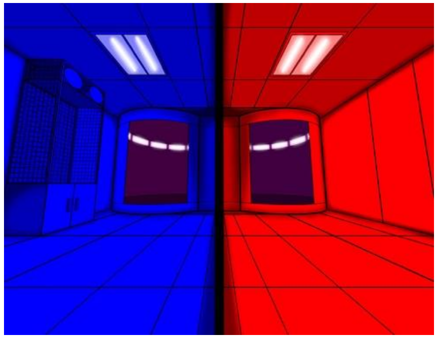

Learn conditional probabilities

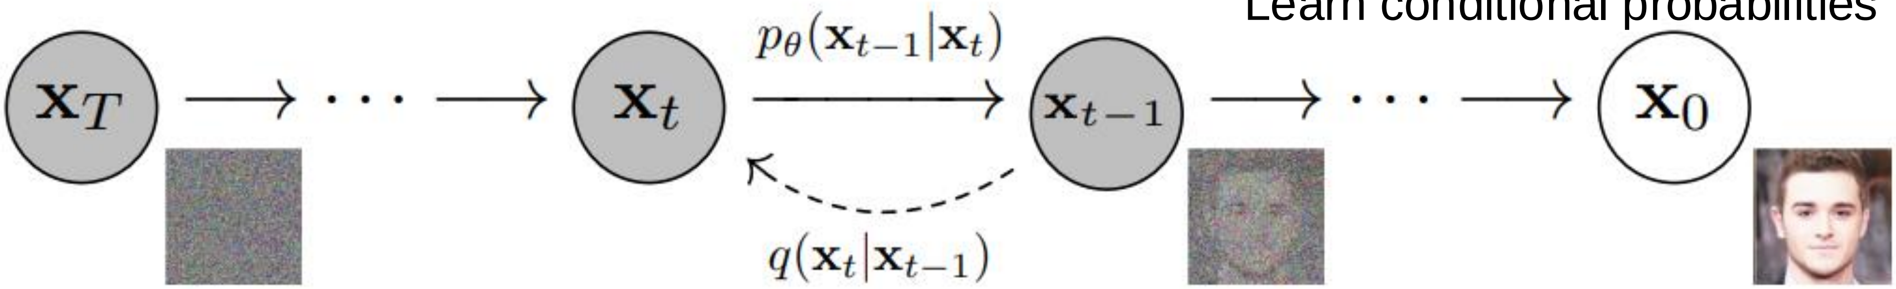

Reverse diffusion

Another example of diffusion in images

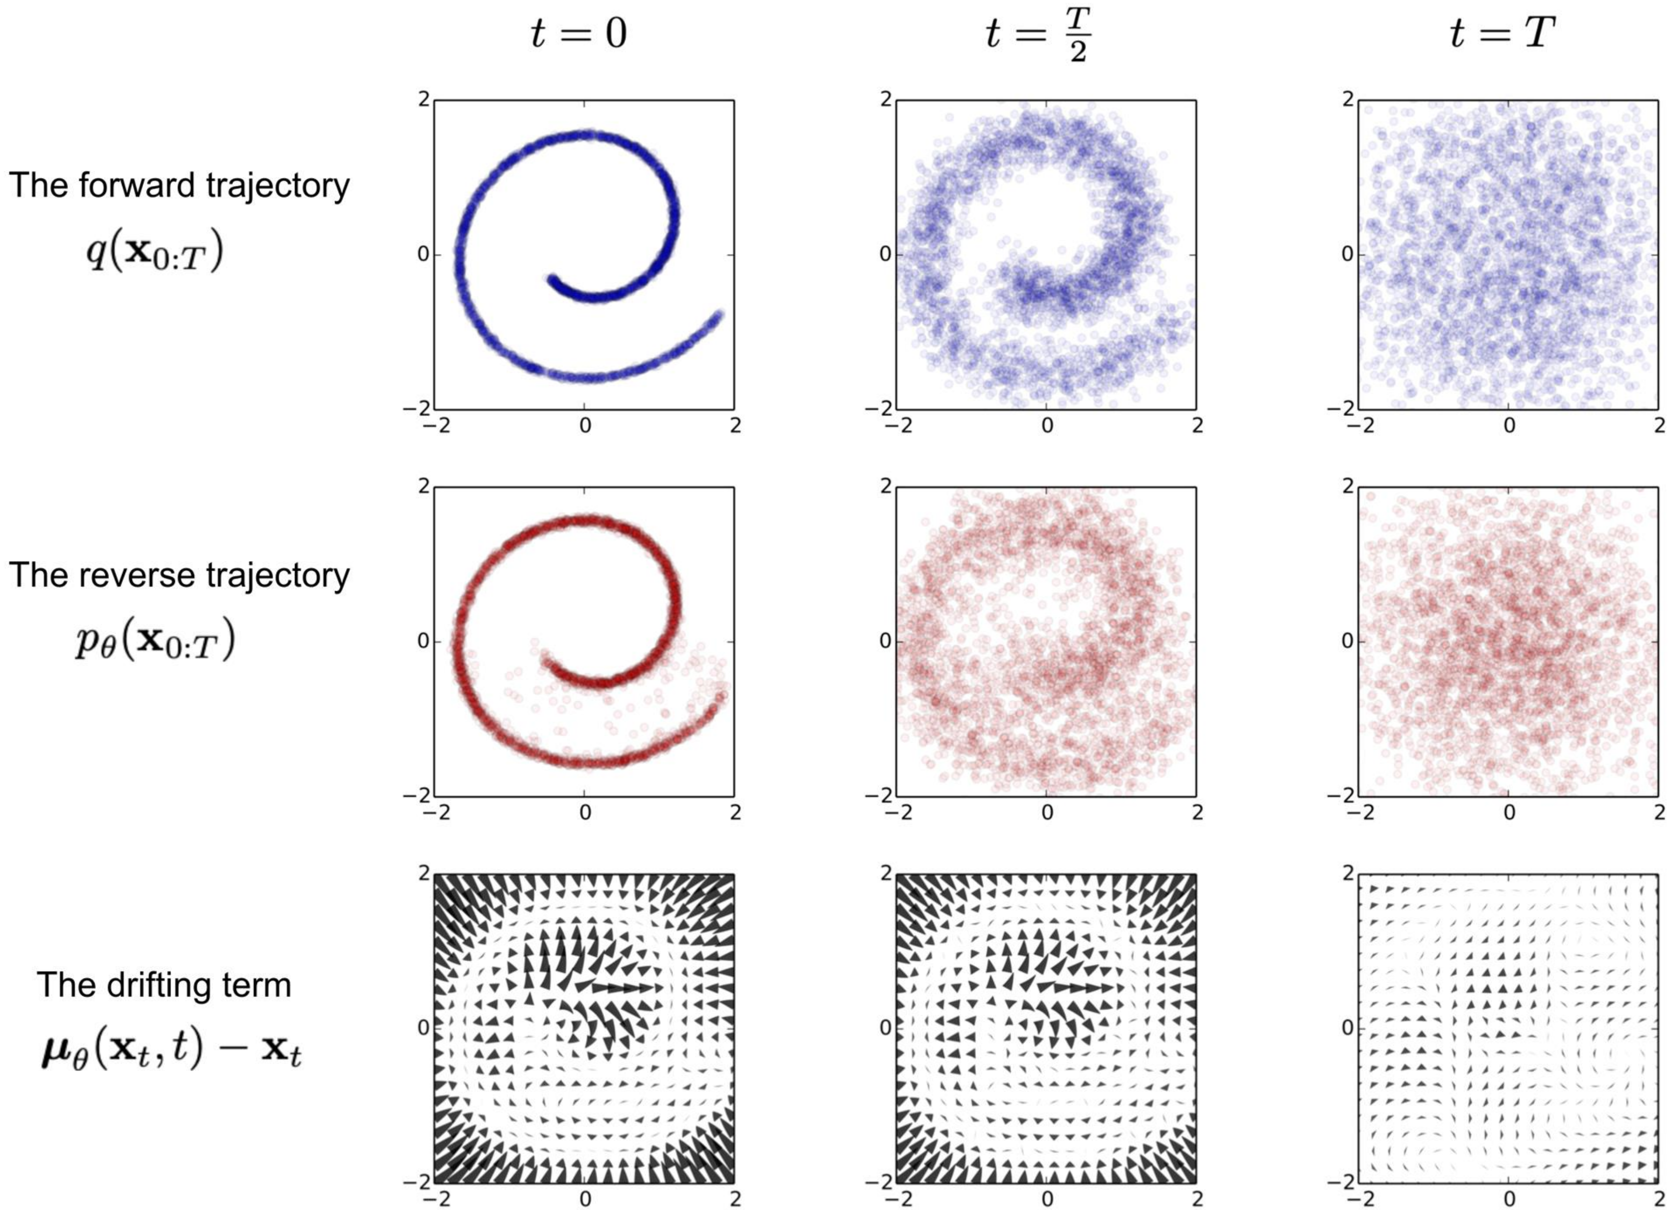

https://lilianweng.github.io/posts/2021-07-11-diffusion-models/

Sample workflow: diffusion model in image latent space

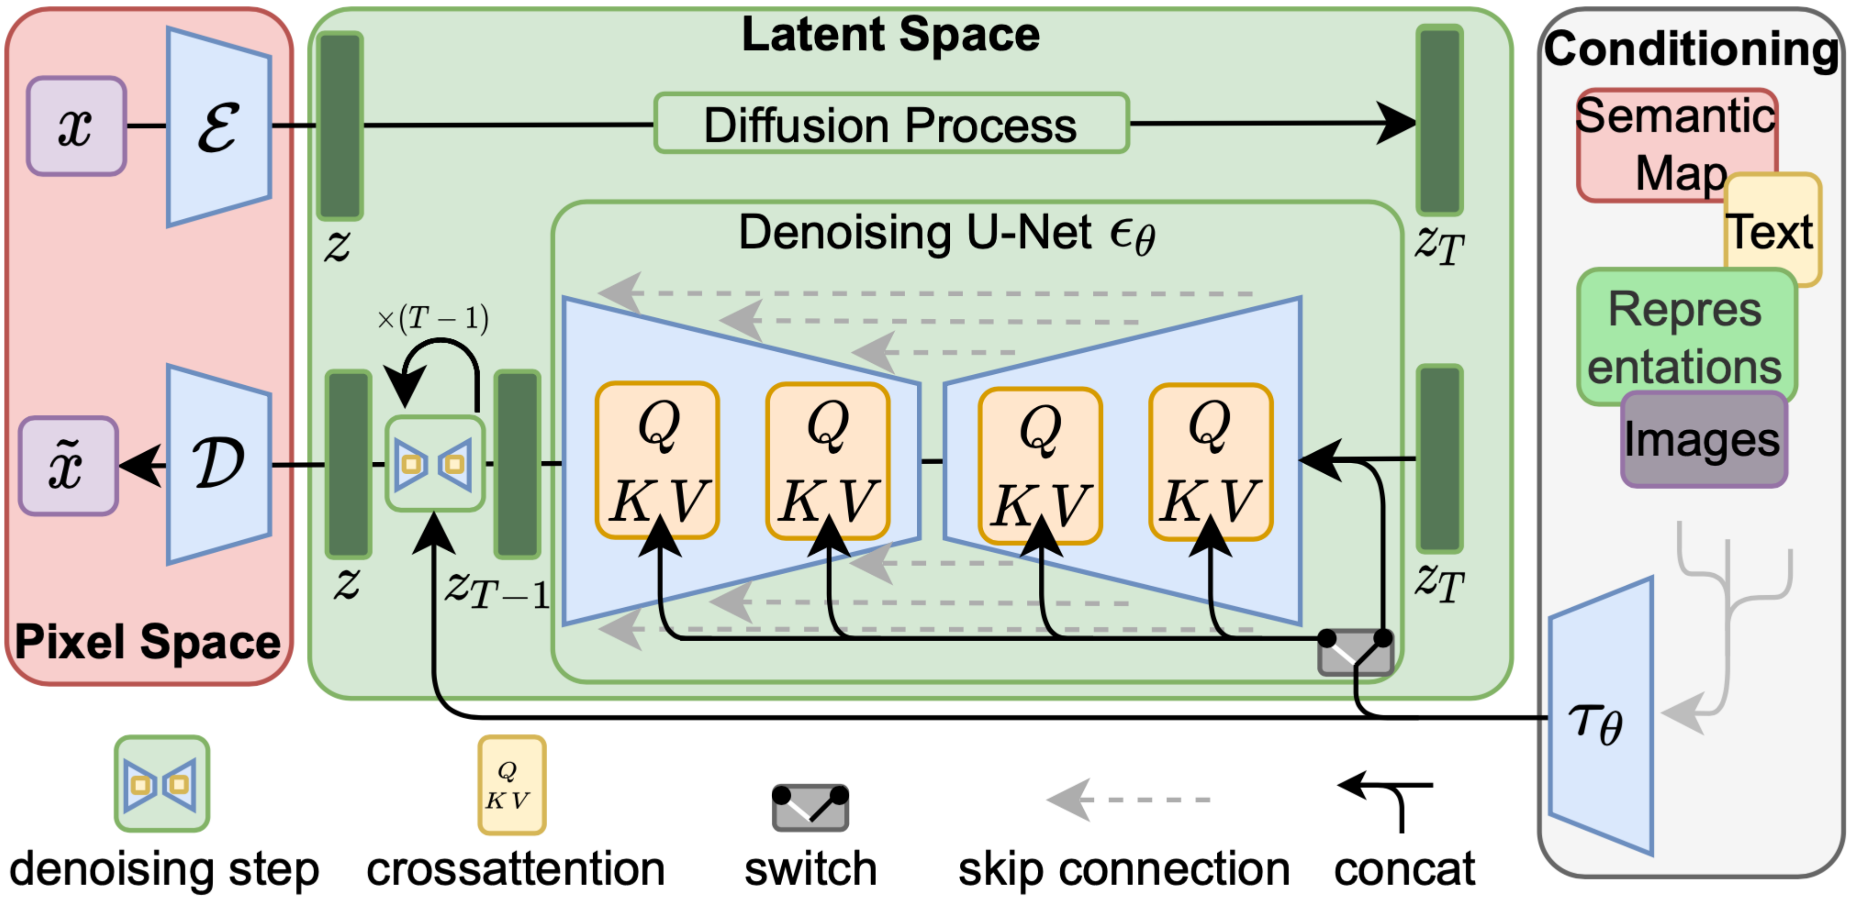

Low loss, stable, ~1000×expensive to run than GANs

https://lilianweng.github.io/posts/2021-07-11-diffusion-models/

Main application of diffusion models: image and video generation

https://www.superannotate.com/blog/diffusion-models

In materials, diffusion models can be used for structure generation

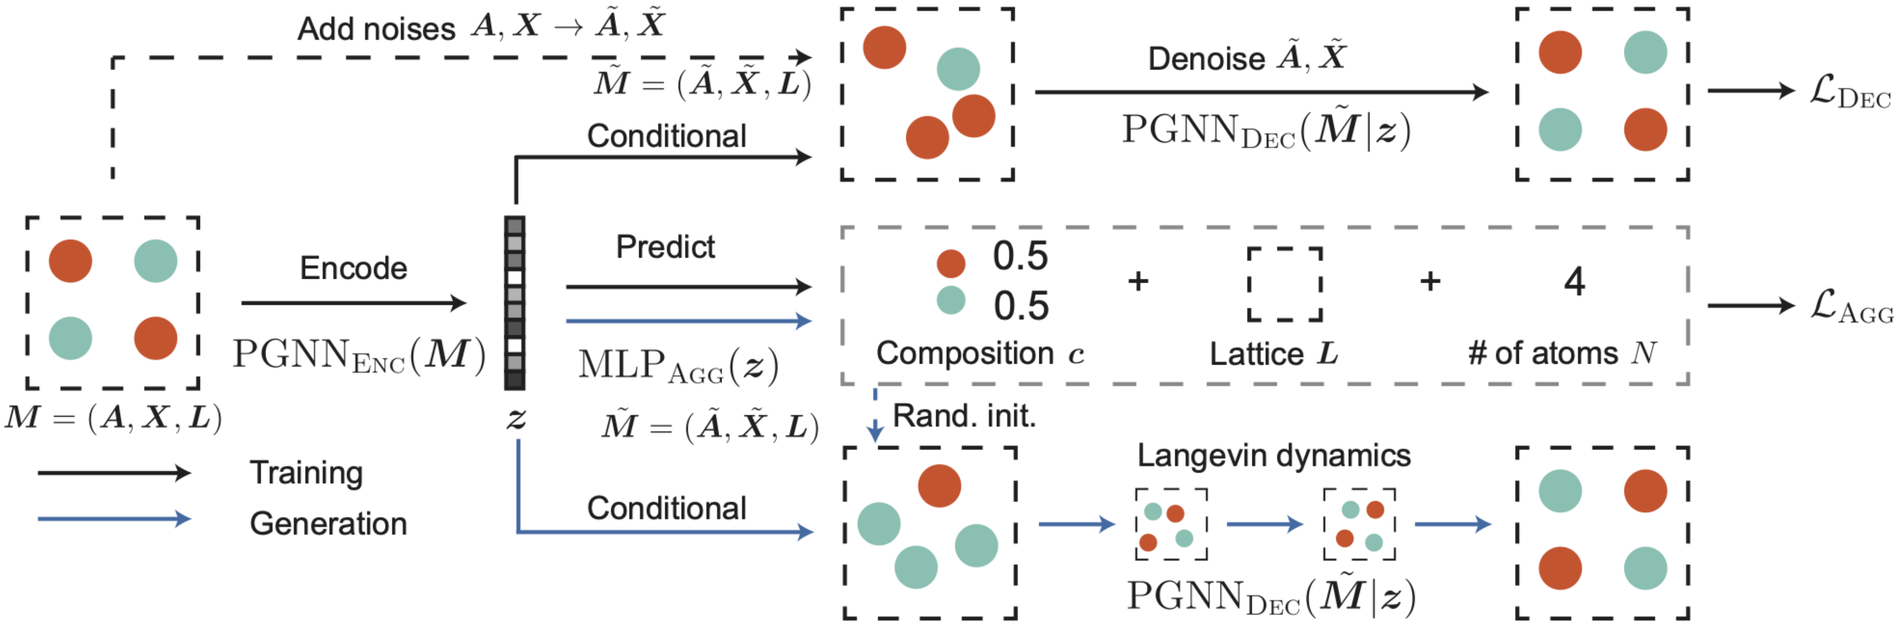

Crystal diffusion variational autoencoder (CDVAE)

- One of the first diffusion models to be developed for structure prediction
- Periodic graph networks for encoding a latent space and denoising
- Property predictor: for composition, lattice, and number of atoms from latent space
- Langevin dynamics: final structure

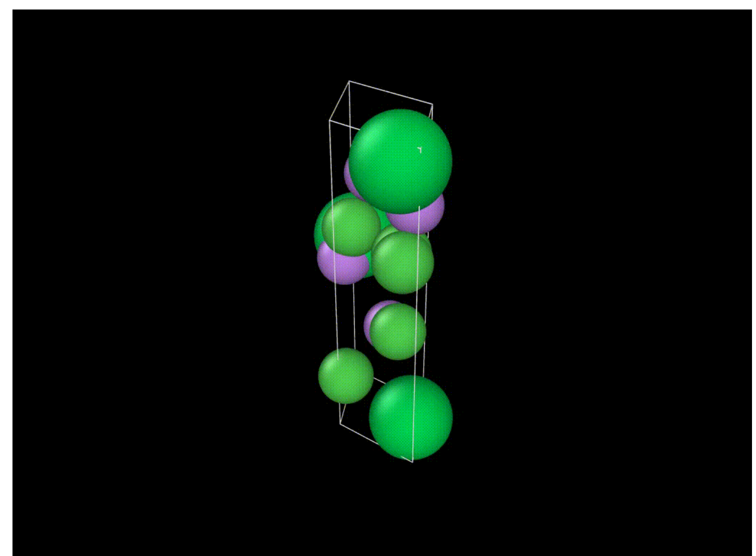

Xie et al., ICLR 2022; https://github.com/txie-93/cdvae

Diffusion models in action

Inverse design of new superconductors

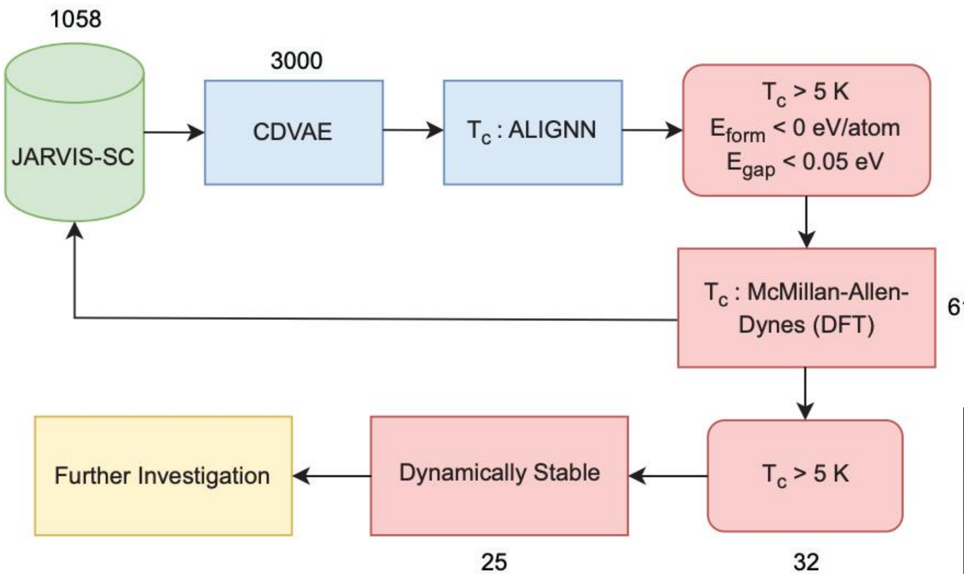

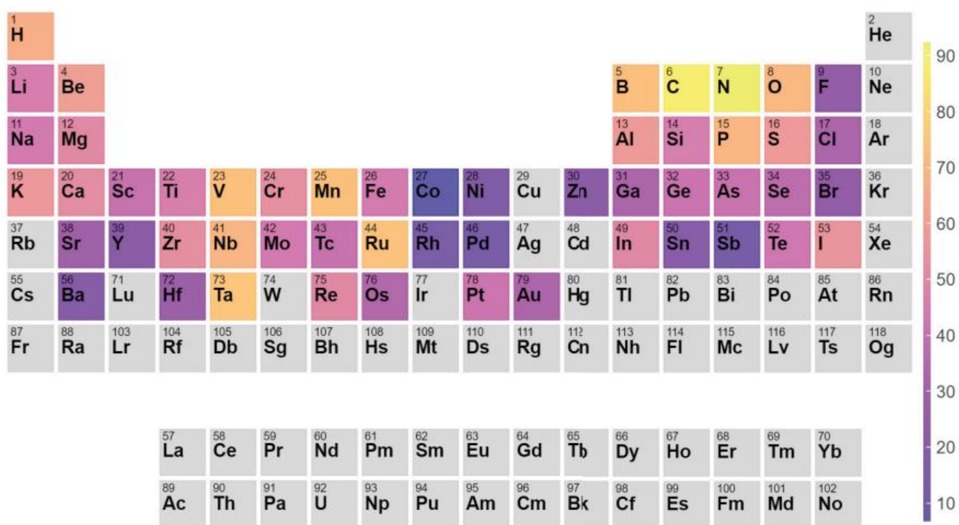

Wines et al., arXiv, 2304.08446v4 (2023)

UniMat + Diffusion (Google Deepmind)

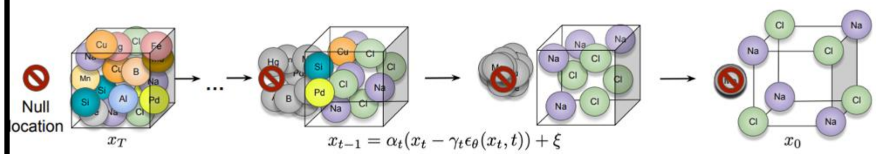

Yang et al., arXiv, 2311.09235v2 (2023)

MatterGen (Microsoft)

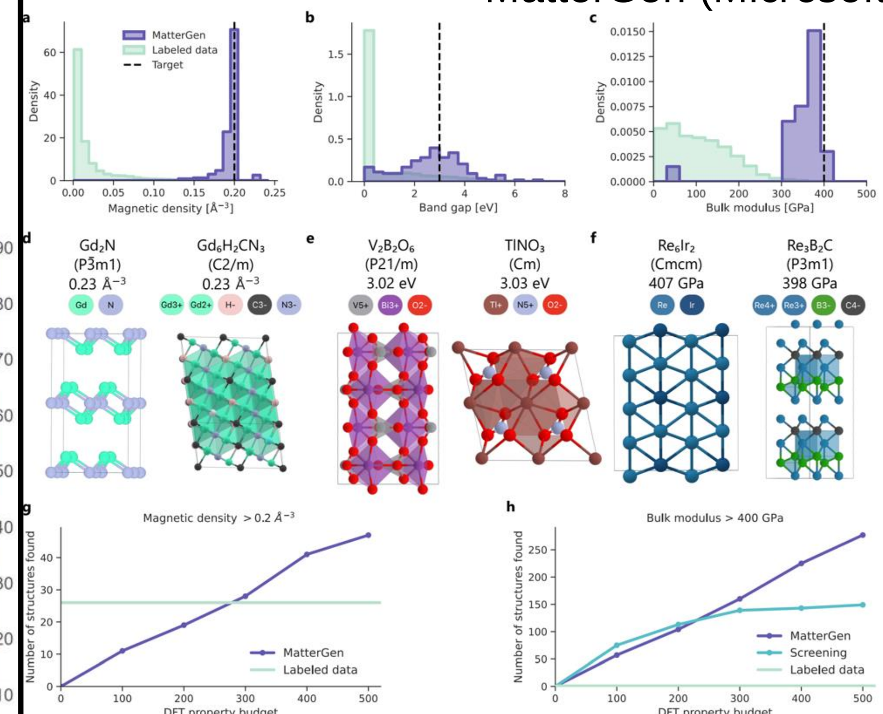

Zeni et al., arXiv, 2312.03687v2 (2024)

Evaluating diffusion models

- Validity
  - Overlapping atoms? Large lattice parameters? Charge-neutral?
- Uniqueness
  - How diverse are the generated structures (in a randomly chosen sample size)?
- Structure matching
  - Do one of the ‘best’ generated structures include the known ground state?
- (Meta)stability
  - How (meta or un)stable are the generated structures?
  - Evaluated using density functional theory or a foundational interatomic potential

Park et al., https://doi.org/10.26434/chemrxiv-2024-rw8p5

Chemeleon: text + structure

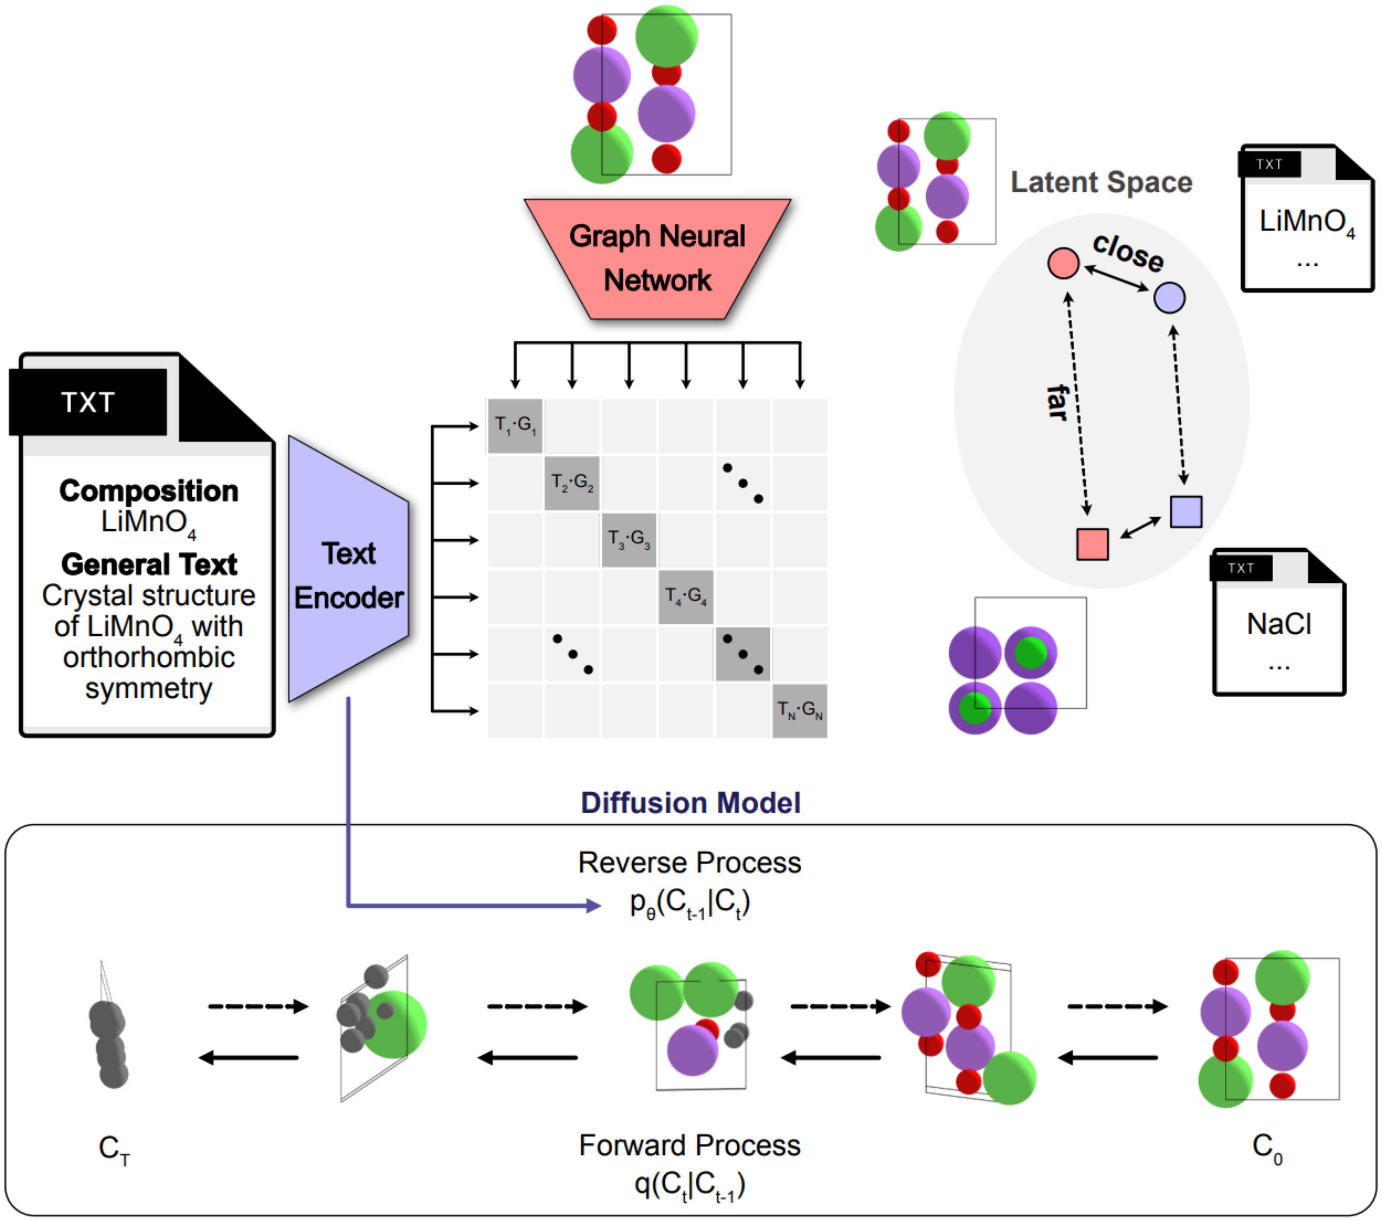

Text is aligned with graph embeddings for a given structure (contrastive learning): Crystal CLIP

Text used to predict noise during denoising

- Can eliminate need for a secondary model for property prediction

Park et al., https://doi.org/10.26434/chemrxiv-2024-rw8p5

Text+graphs: better at structure classification

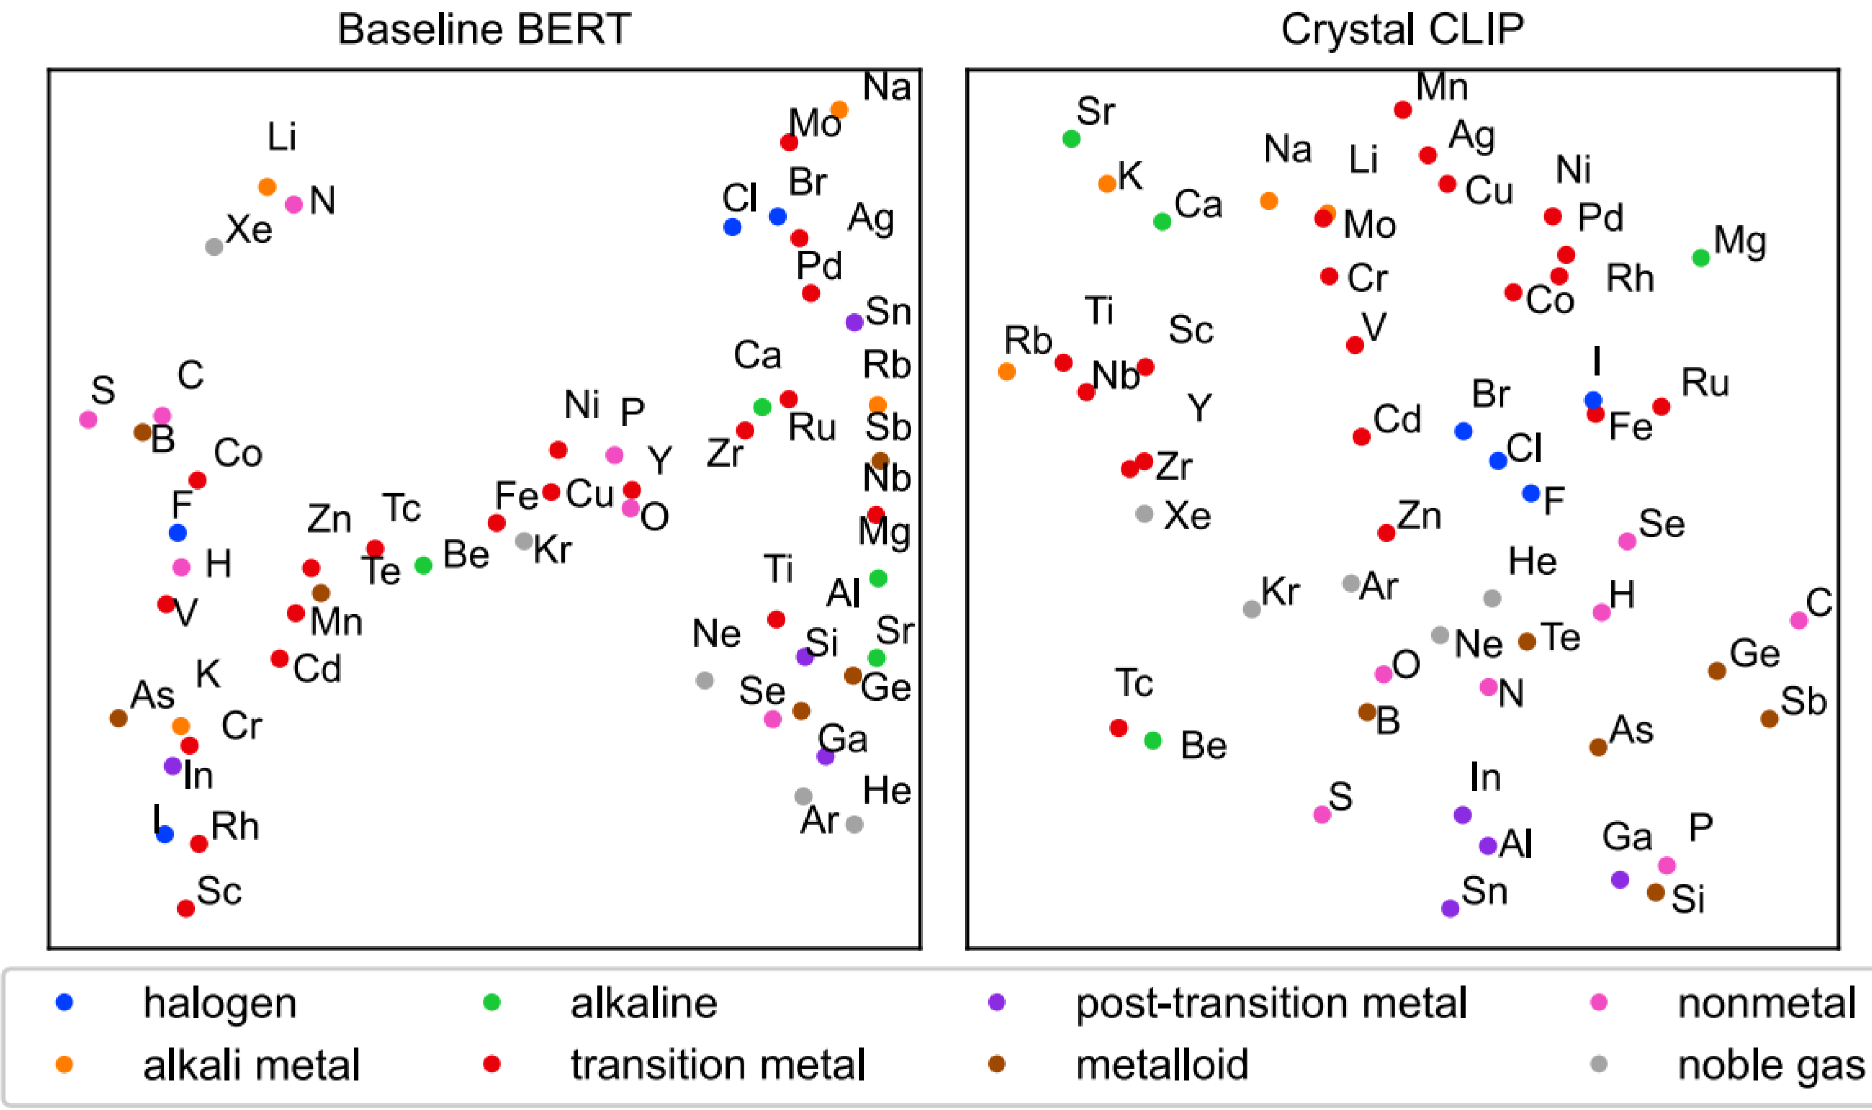

Park et al., https://doi.org/10.26434/chemrxiv-2024-rw8p5

Chemeleon in action

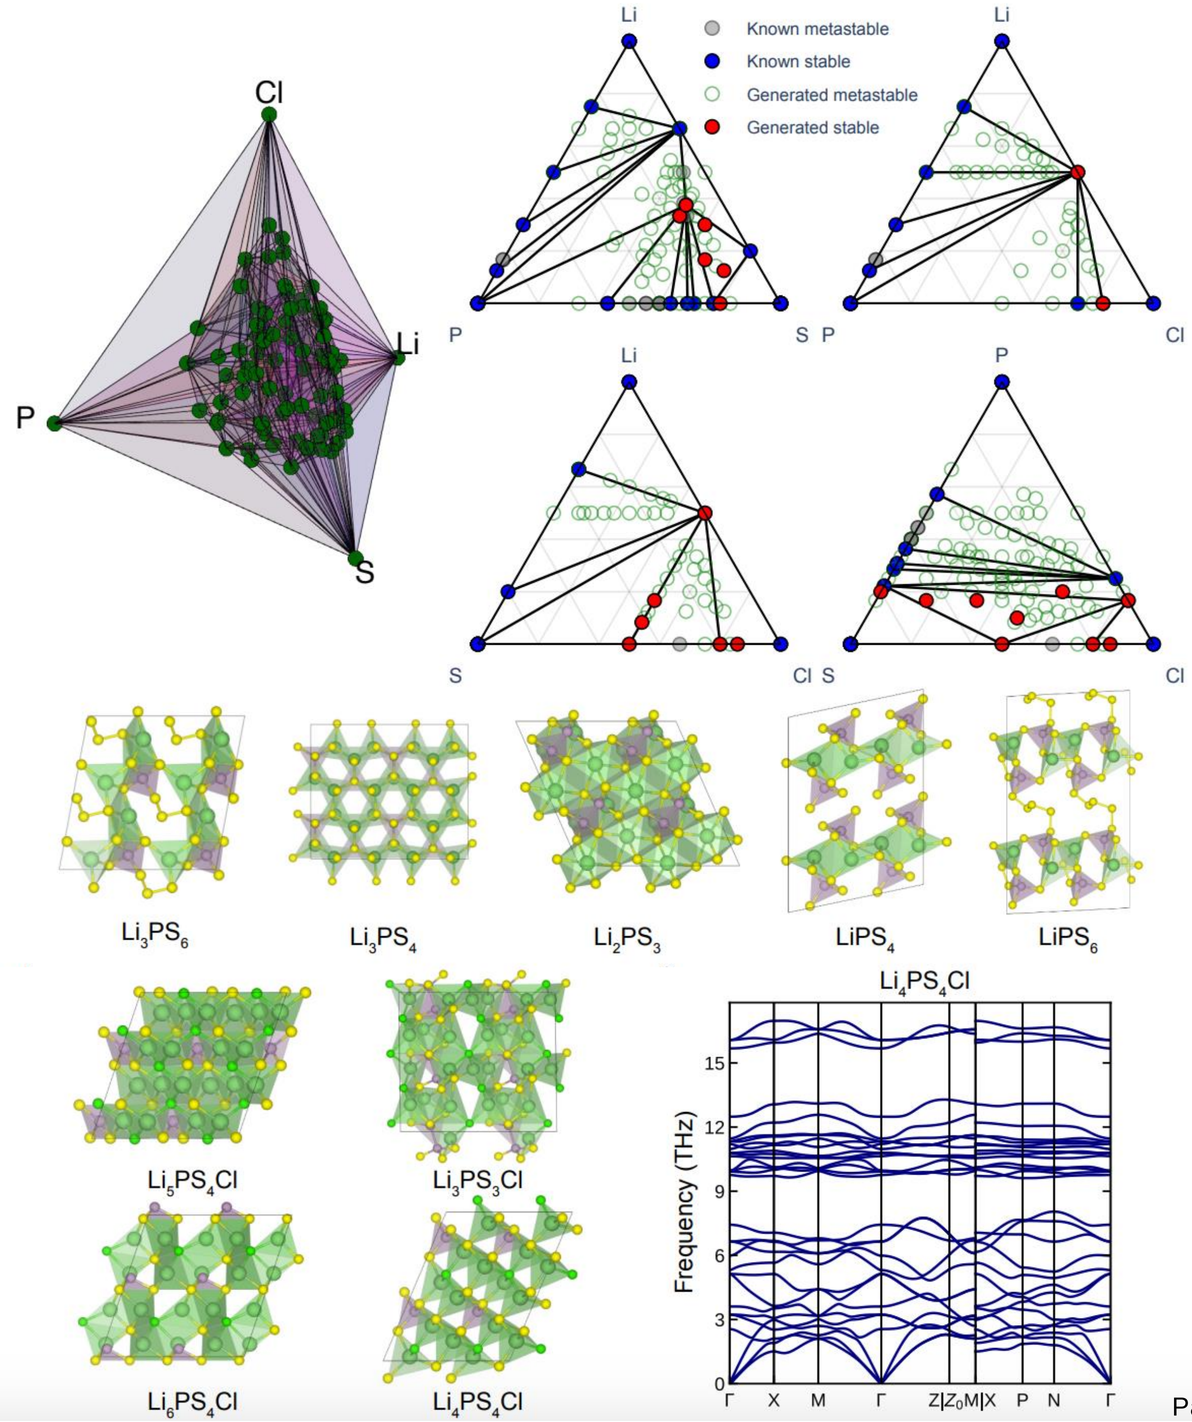

Scanning the Li-P-S-Cl quaternary system

- 2400 possible compositions for a max coefficient of 6
- Use charge neutrality to restrict compositions to 781

Some generated structures are stable

- Verified using density functional theory

Park et al., https://doi.org/10.26434/chemrxiv-2024-rw8p5

Hands—on session?

Generate some structures?! With Chemeleon

Set the generation parameters

Here we generate just one sample, to run quickly. But you can increase this later. Since we are using the composition only model, it can only take elements as a prompt.

In [17]:
```
# Set parameters
n_samples = 3
n_atoms = 5
prompt = "Ba Ti O"
```

In [ ]:
```
%%time
# Generate crystal structures
atoms_list = composition_model.sample(prompt, n_atoms, n_samples)
```

In [ ]:
```
# Visualise
visualizer = Visualizer(atoms_list)
visualizer.view(index=0)
```

Summary

- Generative models can facilitate structure generation/enumeration
  - Identify structures beyond simple human intuition
- Classical: autoencoders, modern: diffusion
- Diffusion: learn probability distributions associated with noising/denoising
- Nascent stage: can produce ‘bad’ structures or ‘incorrect’ compositions
  - More chemical constraints?
  - Experimental validation?

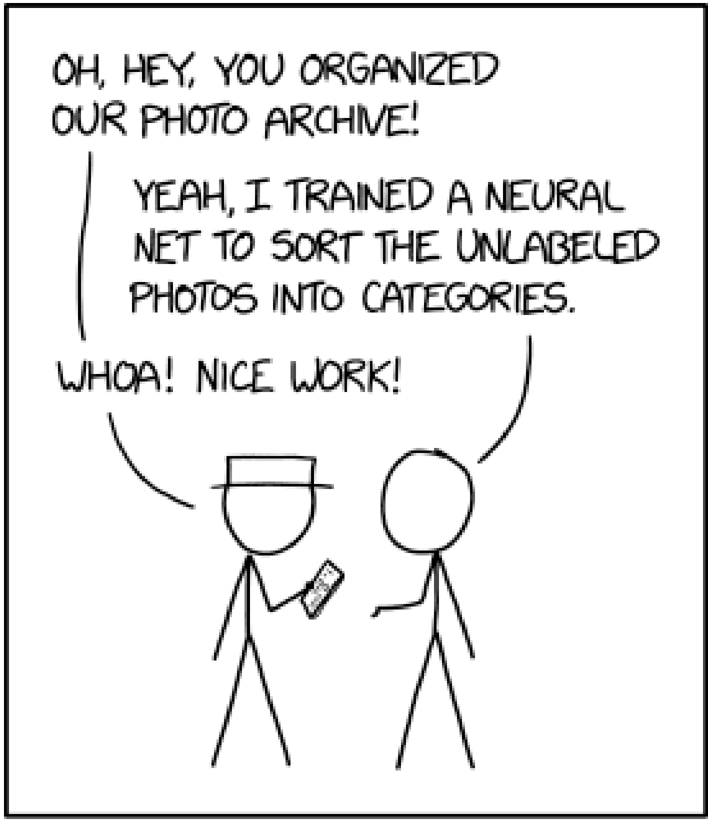

ENGINEERING TIP:
WHEN YOU DO A TASK BY HAND, YOU CAN TECHNICALLY SAY YOU TRAINED A NEURAL NET TO DO IT.