

Part VI:  
What's Next?

# Outline

## Deep Generative Models



- Richard Feynman: “What I cannot create, I do not understand.”

## Deep Reinforcement Learning

- The technique makes Alpha Go better than professional players.

# Creation

This is Unsupervised Learning.



Draw something!

# Creation

Now



V.S.



If machine can draw a cat ...



# Deep Generative Models

Component-by-component


Auto-encoder

Adversarial Generative Network (GAN)

# Component-by-component

- To create an image, generating a pixel each time

E.g. 3 x 3 images



Can be trained with a large collection of images without any annotation

# Pokémon Creation




- Small images of 792 Pokémon's
  - Can machine learn to create new Pokémons?

**Don't catch them! Create them!**

- Source of image:  
[http://bulbapedia.bulbagarden.net/wiki/List\\_of\\_Pok%C3%A9mon\\_by\\_base\\_stats\\_\(Generation\\_VI\)](http://bulbapedia.bulbagarden.net/wiki/List_of_Pok%C3%A9mon_by_base_stats_(Generation_VI))

Original image is 40 x 40

Making them into 20 x 20




# Pokémon Creation

- Each pixel is represented by 3 numbers (corresponding to RGB)



R=50, G=150, B=100

- Each pixel is represented by a 1-of-N encoding feature



Clustering the similar color  167 colors in total





Real  
Pokémon



Never seen  
by machine!

Cover 50%

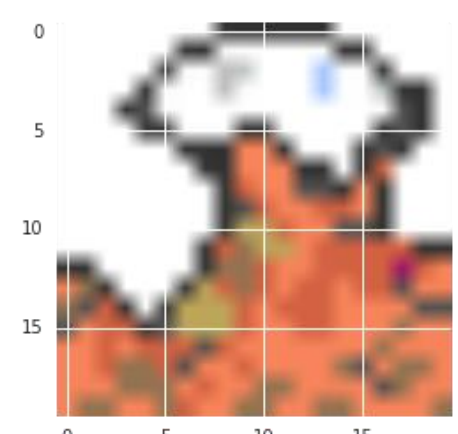





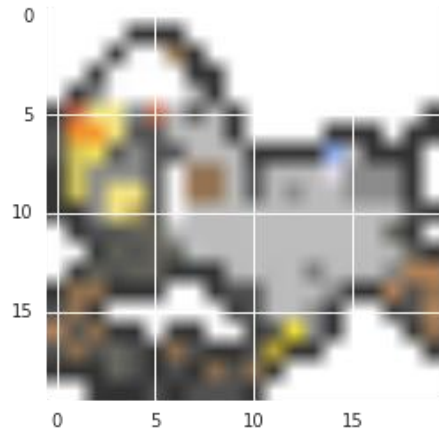
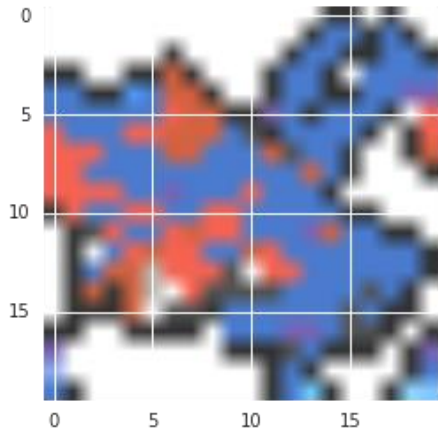


Cover 75%



# Pokémon Creation

Drawing from scratch  
Need some randomness



# PixelRNN


Ref: Aaron van den Oord, Nal Kalchbrenner, Koray Kavukcuoglu, Pixel Recurrent Neural Networks, arXiv preprint, 2016



Real  
World



# More than images .....



Audio: Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, Koray Kavukcuoglu, WaveNet: A Generative Model for Raw Audio, arXiv preprint, 2016

Video: Nal Kalchbrenner, Aaron van den Oord, Karen Simonyan, Ivo Danihelka, Oriol Vinyals, Alex Graves, Koray Kavukcuoglu, Video Pixel Networks , arXiv preprint, 2016

# Deep Generative Models

Criticism: do not consider the generation process globally



Component-by-component

Auto-encoder

Dimension reduction

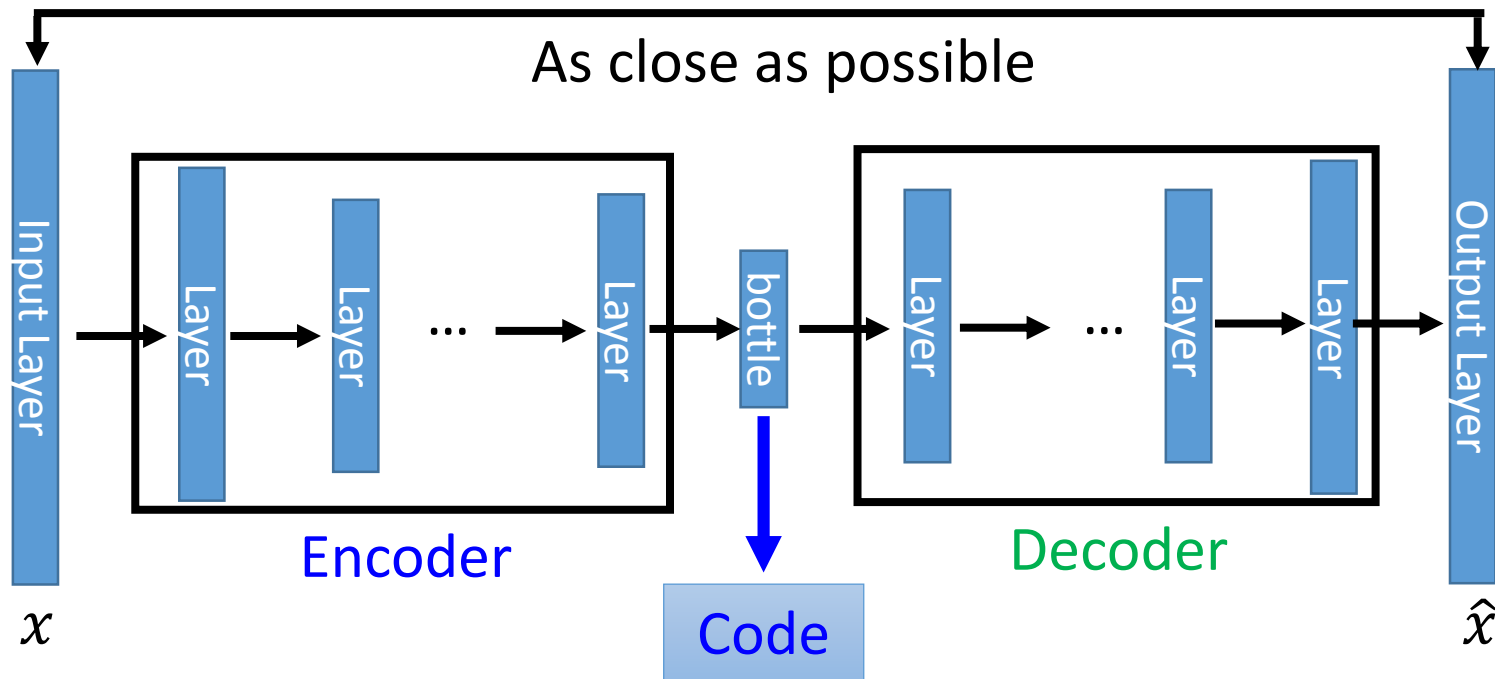
Adversarial Generative Network (GAN)

# Auto-encoder




# Deep Auto-encoder

- NN encoder + NN decoder = a deep network



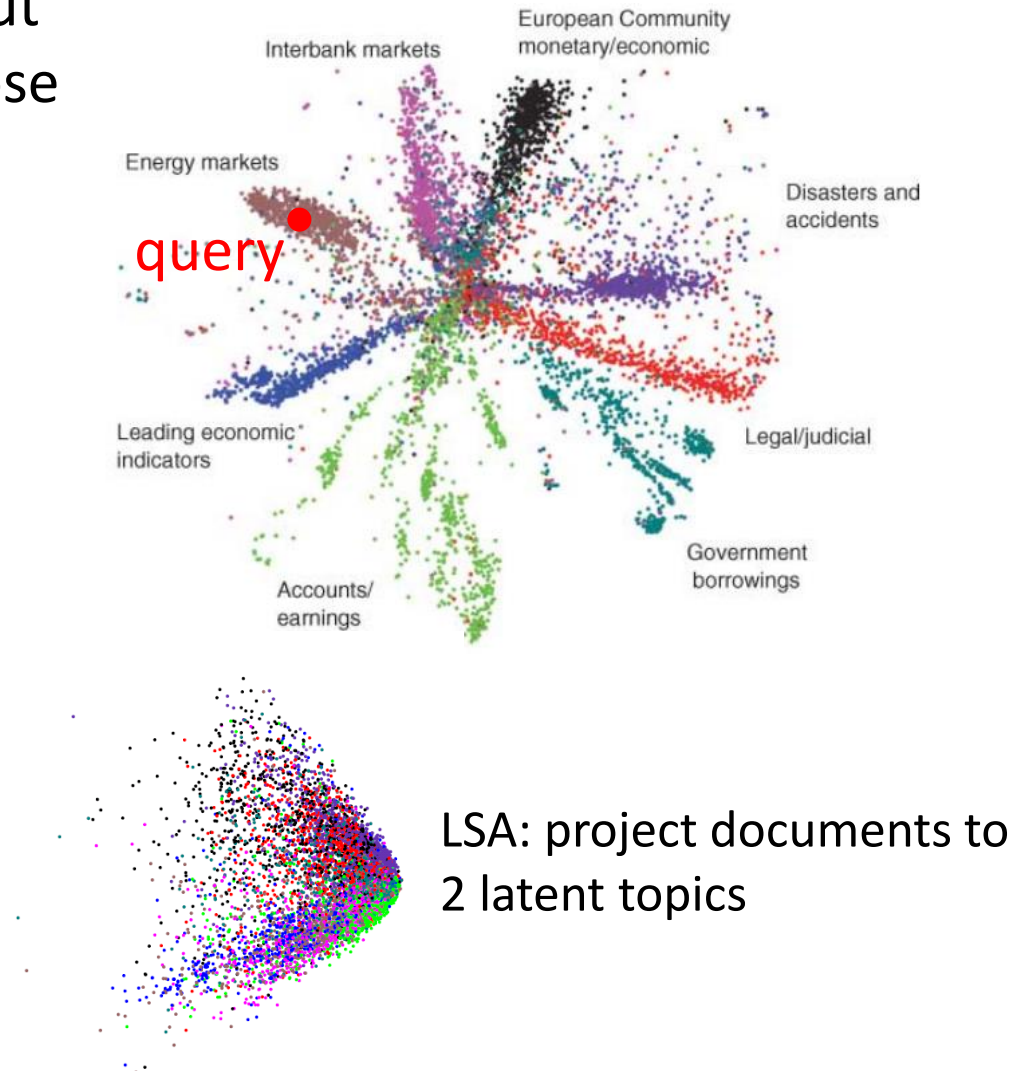

Reference: Hinton, Geoffrey E., and Ruslan R. Salakhutdinov. "Reducing the dimensionality of data with neural networks." *Science* 313.5786 (2006): 504-507





# Auto-encoder – Text Retrieval

The documents talking about the same thing will have close code.




## More: Contractive auto-encoder

# Auto-encoder

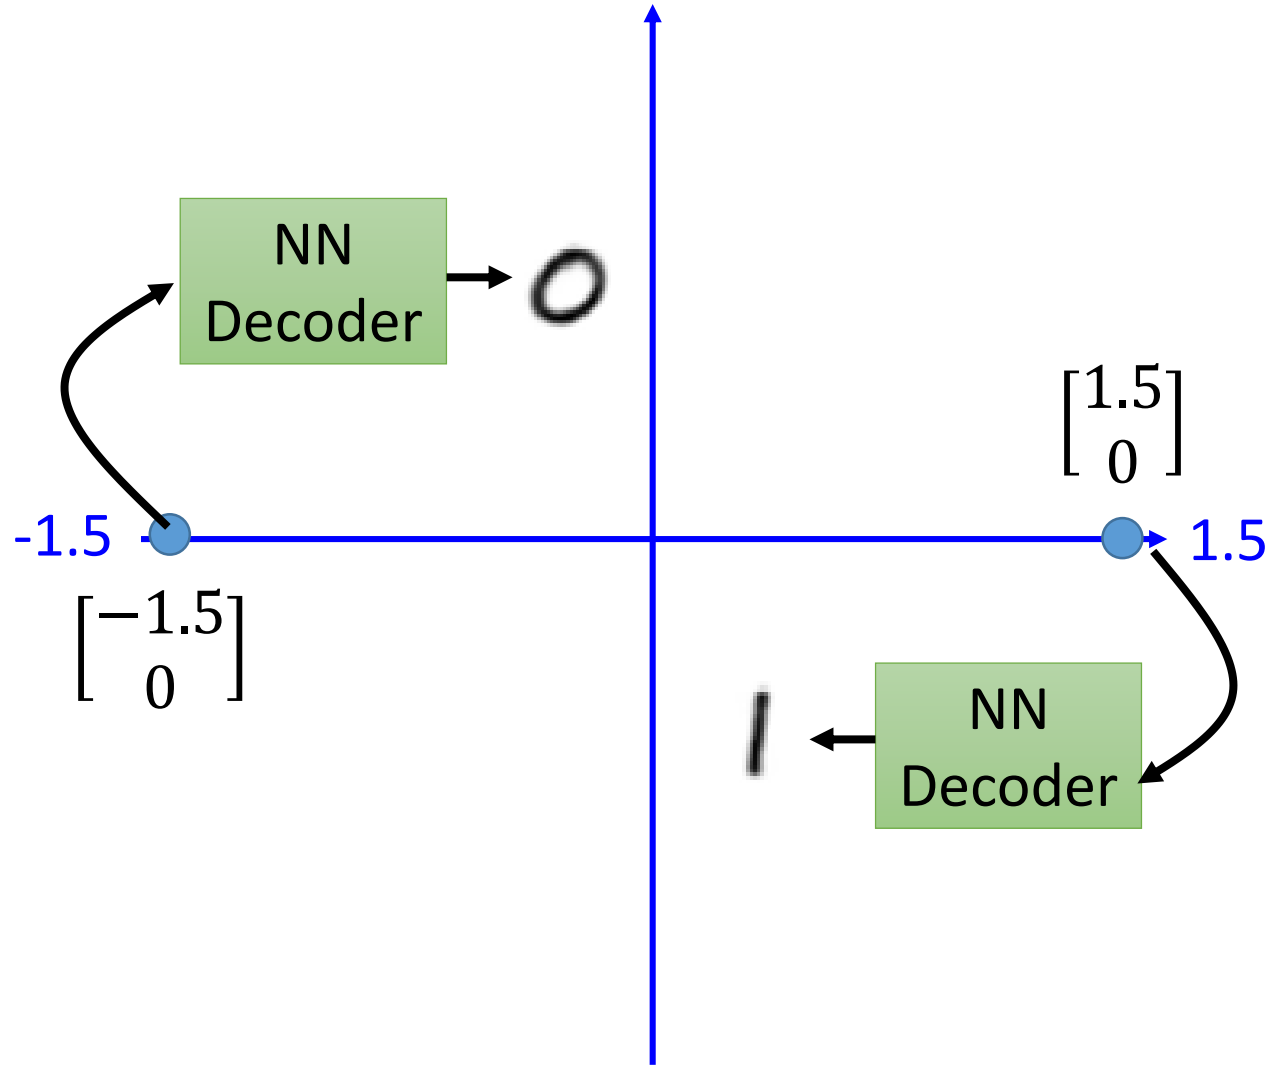

Ref: Rifai, Salah, et al. "Contractive auto-encoders: Explicit invariance during feature extraction." *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*. 2011.

- De-noising auto-encoder





Vincent, Pascal, et al. "Extracting and composing robust features with denoising autoencoders." *ICML*, 2008.


# Generation



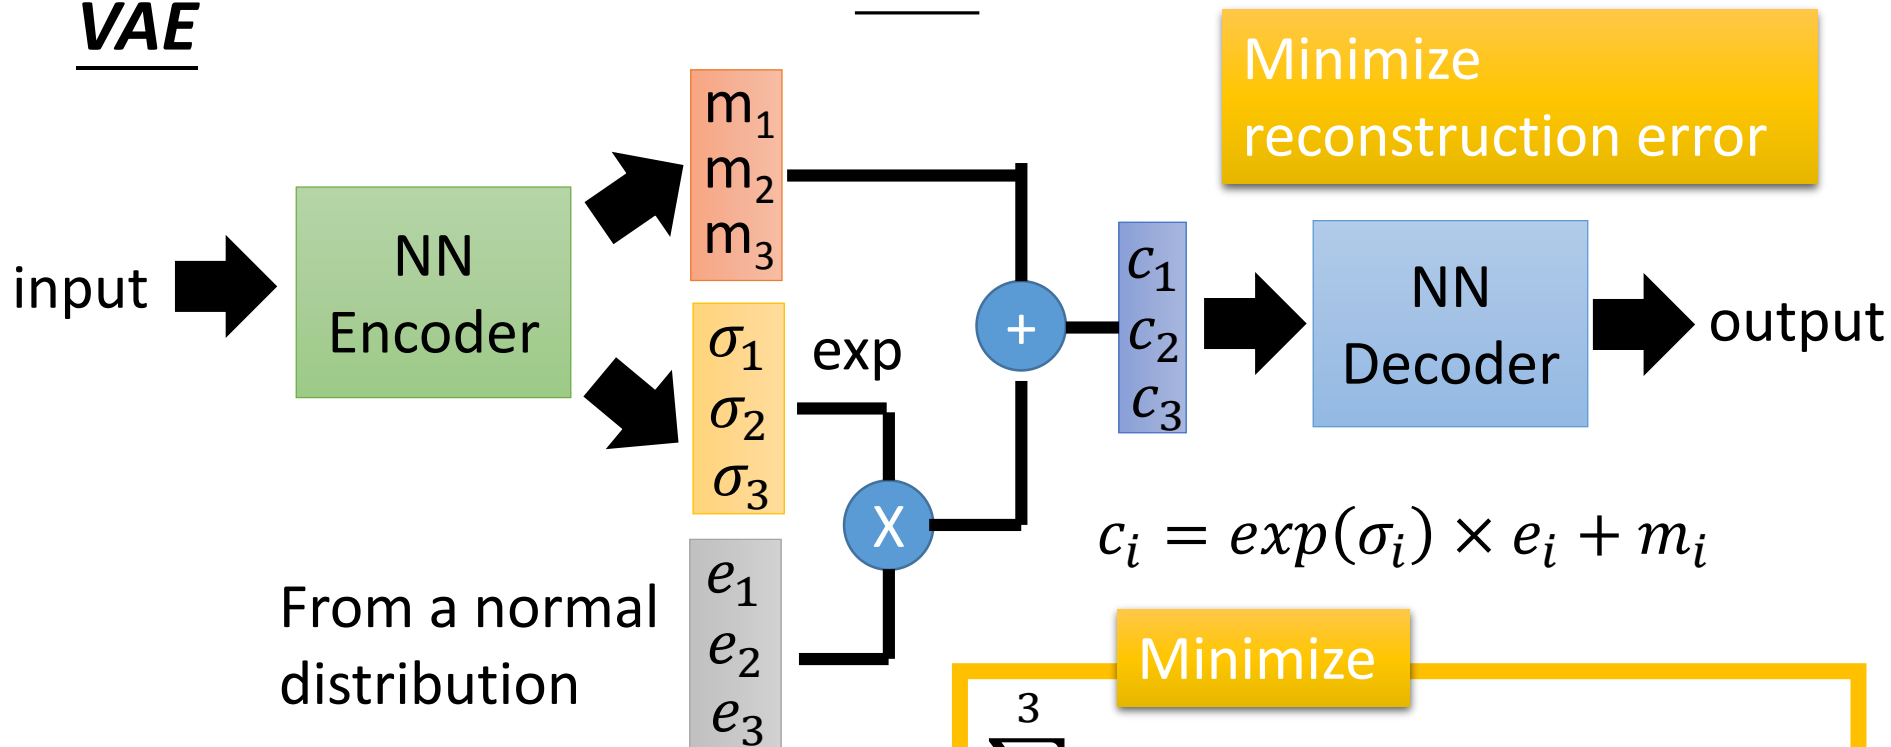
# Generation



# Auto-encoder



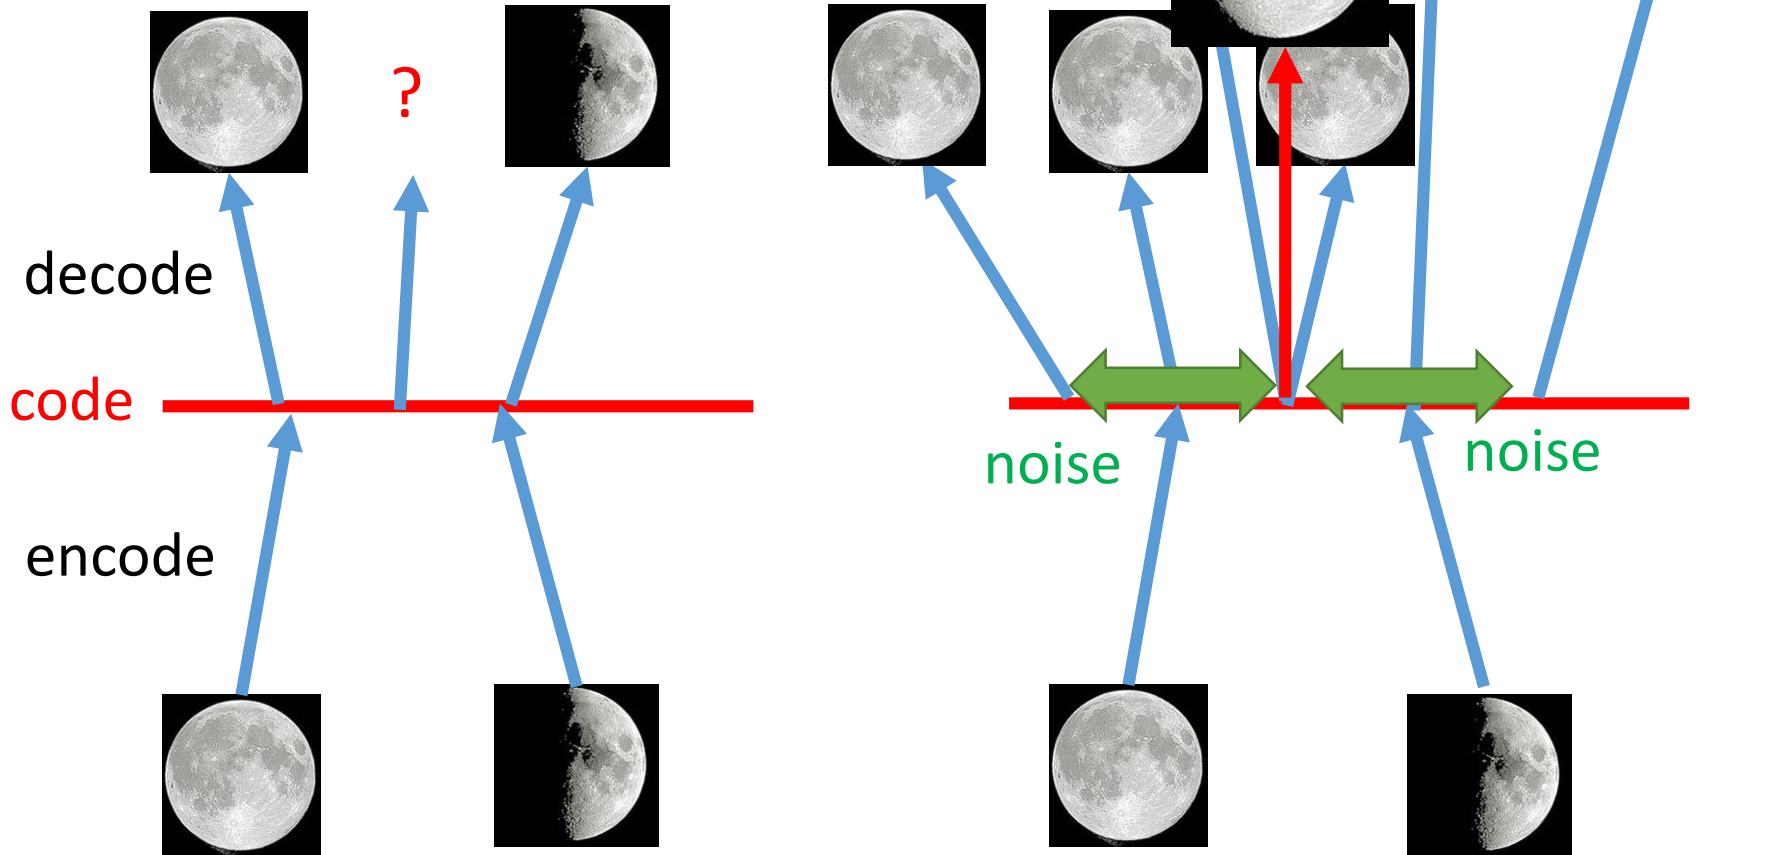
# VAE



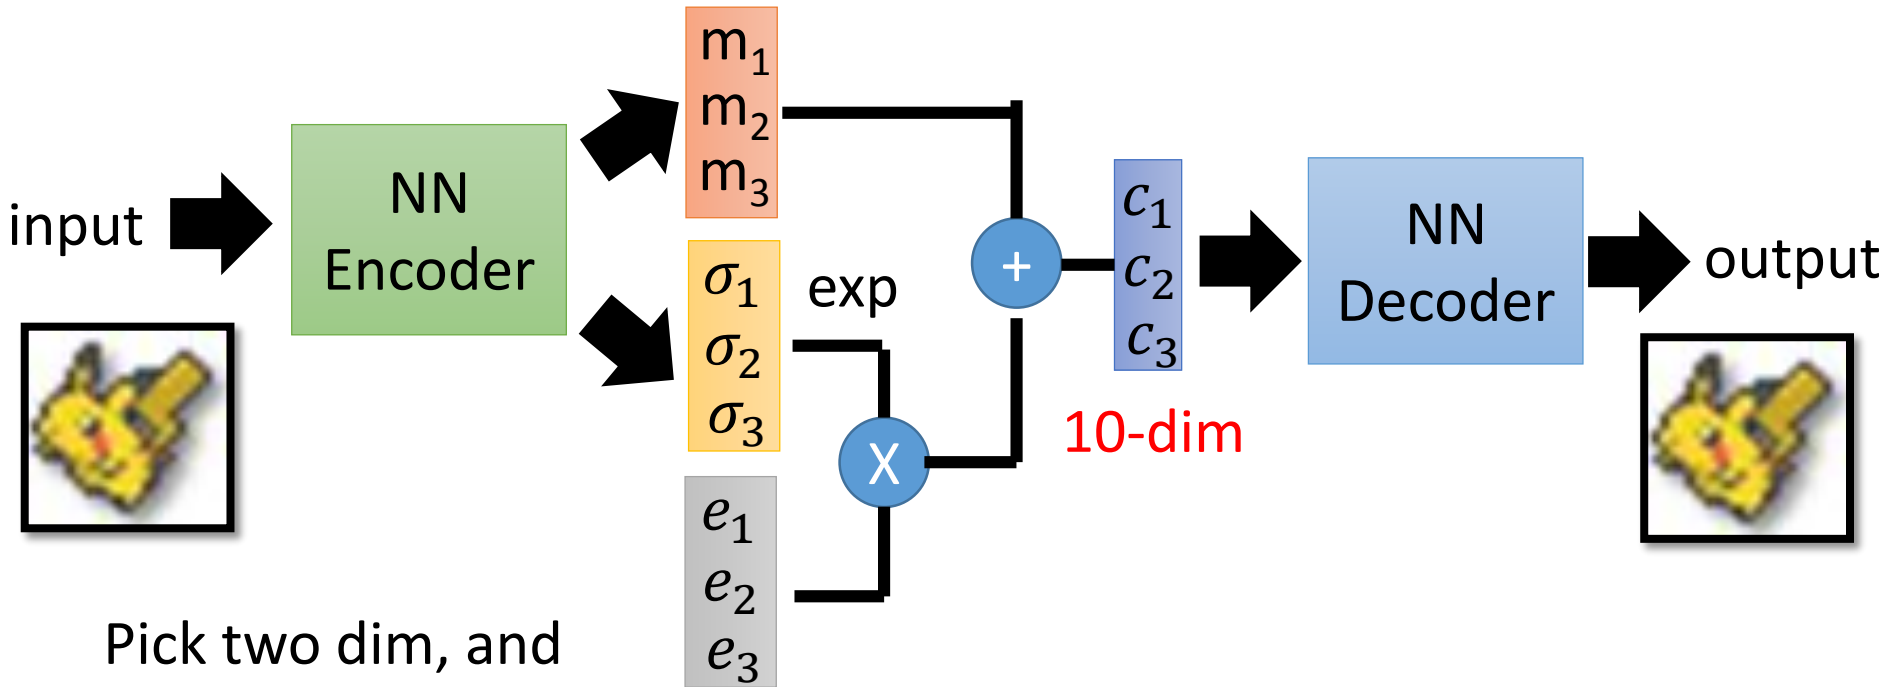
Auto-Encoding Variational Bayes,  
<https://arxiv.org/abs/1312.6114>

# Why VAE?


## Intuitive Reason




# Pokémon Creation



Pick two dim, and fix the rest eight

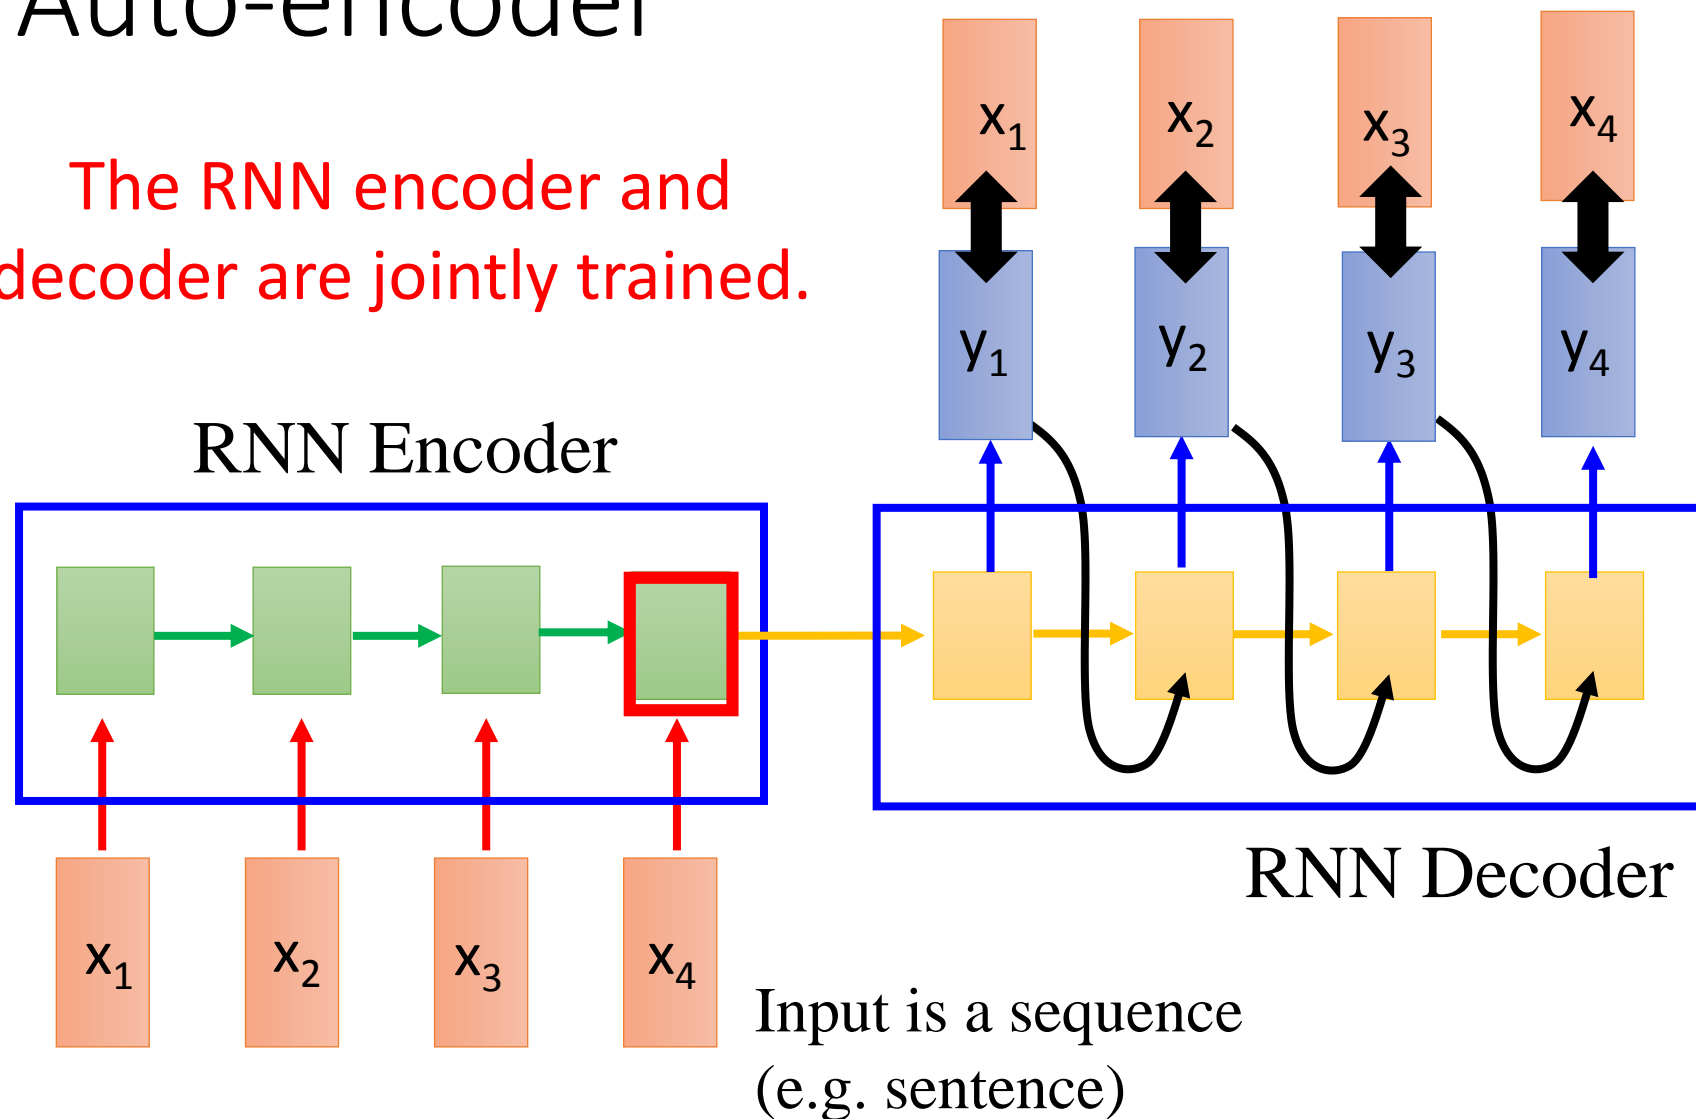




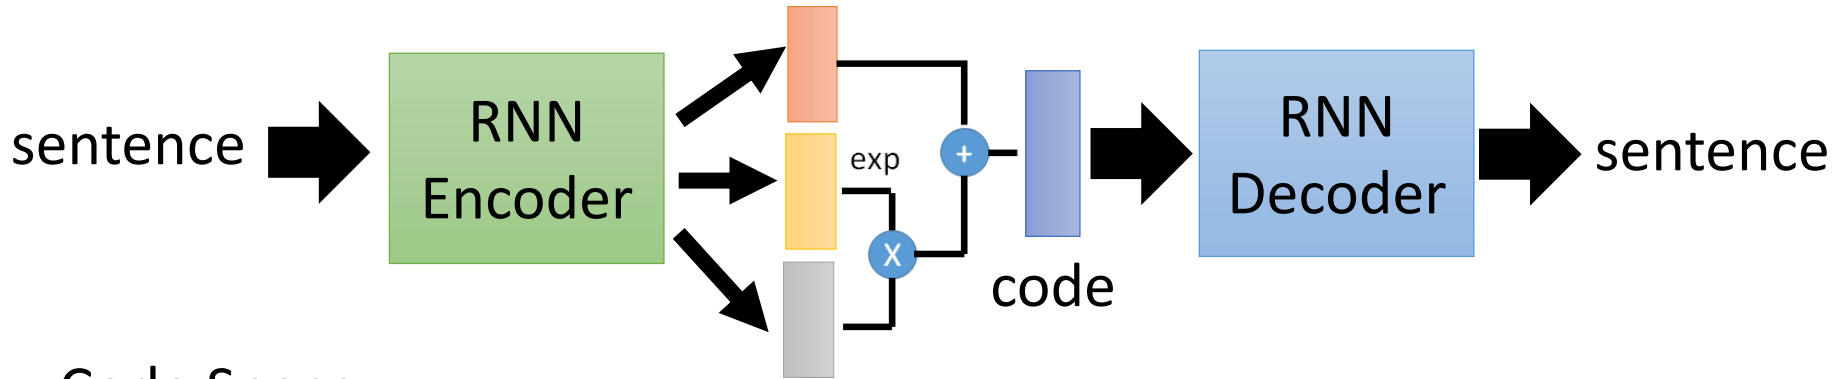


# Sequence-to-sequence Auto-encoder

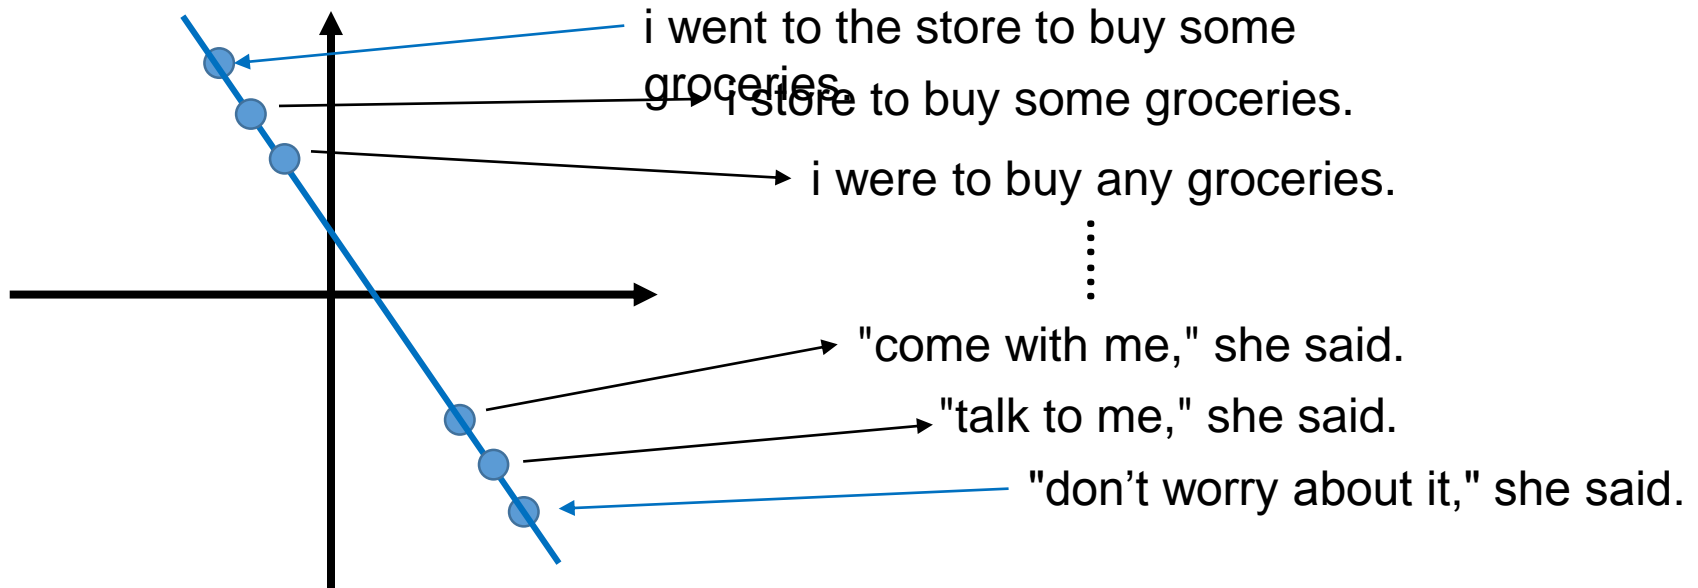
The RNN encoder and decoder are jointly trained.



# VAE – Sentence Generation



## Code Space



Ref: <http://www.wired.co.uk/article/google-artificial-intelligence-poetry>

Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Jozefowicz, Samy Bengio, Generating Sentences from a Continuous Space, arXiv preprint, 2015

# To learn more ...

- Carl Doersch, Tutorial on Variational Autoencoders
- Diederik P. Kingma, Danilo J. Rezende, Shakir Mohamed, Max Welling, “Semi-supervised learning with deep generative models.” *NIPS*, 2014.
- Sohn, Kihyuk, Honglak Lee, and Xinchen Yan, “Learning Structured Output Representation using Deep Conditional Generative Models.” *NIPS*, 2015.
- Xinchen Yan, Jimei Yang, Kihyuk Sohn, Honglak Lee, “Attribute2Image: Conditional Image Generation from Visual Attributes”, *ECCV*, 2016
- Cool demo:
  - [http://vdumoulin.github.io/morphing\\_faces/](http://vdumoulin.github.io/morphing_faces/)
  - <http://fvae.ail.tokyo/>

# Deep Generative Models


Component-by-component

Auto-encoder

Adversarial Generative Network (GAN)

# Problems of VAE

- It does not really try to simulate real images



VAE may just memorize the existing images, instead of generating new images

# Yann LeCun's comment

## What are some recent and potentially upcoming breakthroughs in deep learning?



**Yann LeCun**, Director of AI Research at Facebook and Professor at NYU

Written Jul 29 · Upvoted by [Joaquin Quiñero Candela](#), [Director Applied Machine Learning at Facebook](#) and [Nikhil Garg](#), [I lead a team of Quora engineers working on ML/NLP problems](#)



.....

The most important one, in my opinion, is adversarial training (also called GAN for Generative Adversarial Networks). This is an idea that was originally proposed by Ian Goodfellow when he was a student with Yoshua Bengio at the University of Montreal (he since moved to Google Brain and recently to OpenAI).

This, and the variations that are now being proposed is the most interesting idea in the last 10 years in ML, in my opinion.

<https://www.quora.com/What-are-some-recent-and-potentially-upcoming-breakthroughs-in-deep-learning>



# Yann LeCun's comment

## What are some recent and potentially upcoming breakthroughs in unsupervised learning?



**Yann LeCun**, Director of AI Research at Facebook and Professor at NYU

Written Jul 29 · Upvoted by Joaquin Quiñonero Candela, [Director Applied Machine Learning at Facebook](#) and Huang Xiao



Adversarial training is the coolest thing since sliced bread.

I've listed a bunch of relevant papers in a previous answer.

Expect more impressive results with this technique in the coming years.

What's missing at the moment is a good understanding of it so we can make it work reliably. It's very finicky. Sort of like ConvNet were in the 1990s, when I had the reputation of being the only person who could make them work (which wasn't true).

<https://www.quora.com/What-are-some-recent-and-potentially-upcoming-breakthroughs-in-unsupervised-learning>



# Evolution

<http://peellden.pixnet.net/blog/post/40406899-2013-%E7%AC%AC%E5%9B%9B%E5%AD%A3%EF%BC%8C%E5%86%AC%E8%9D%B6%E5%AF%82%E5%AF%A5>



Brown



veins

Butterflies are not brown



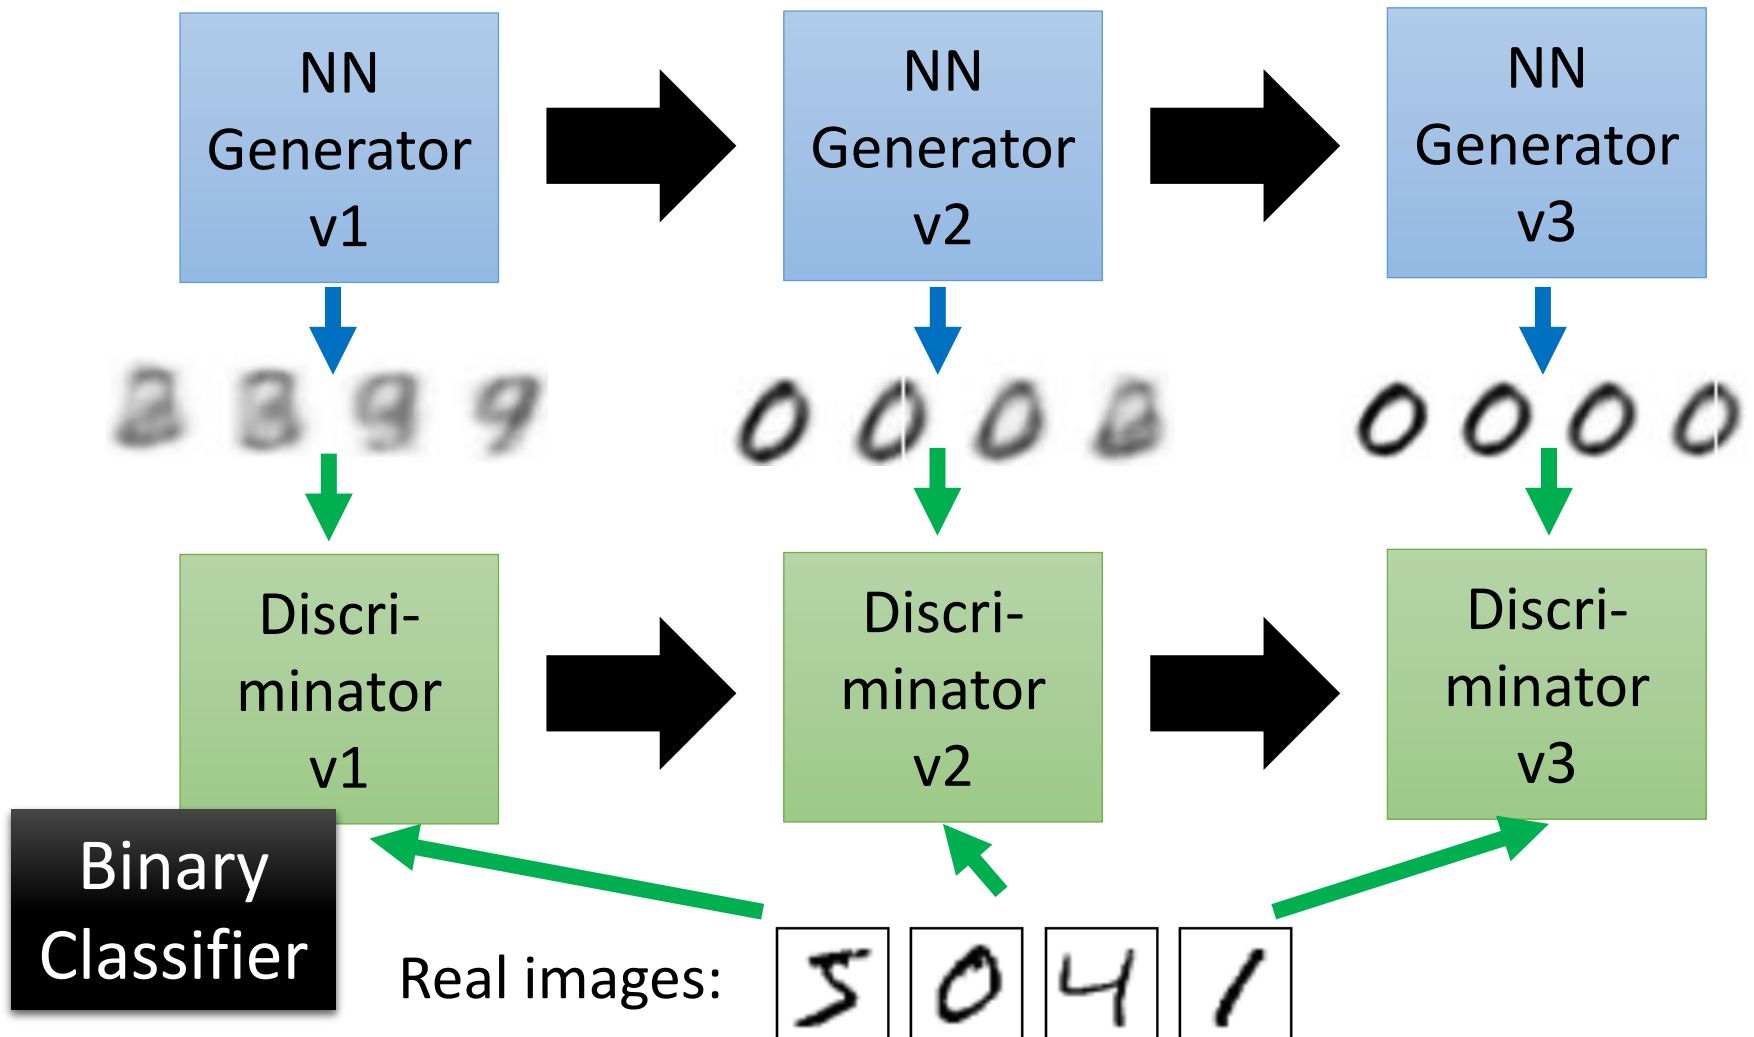
Butterflies do not have veins



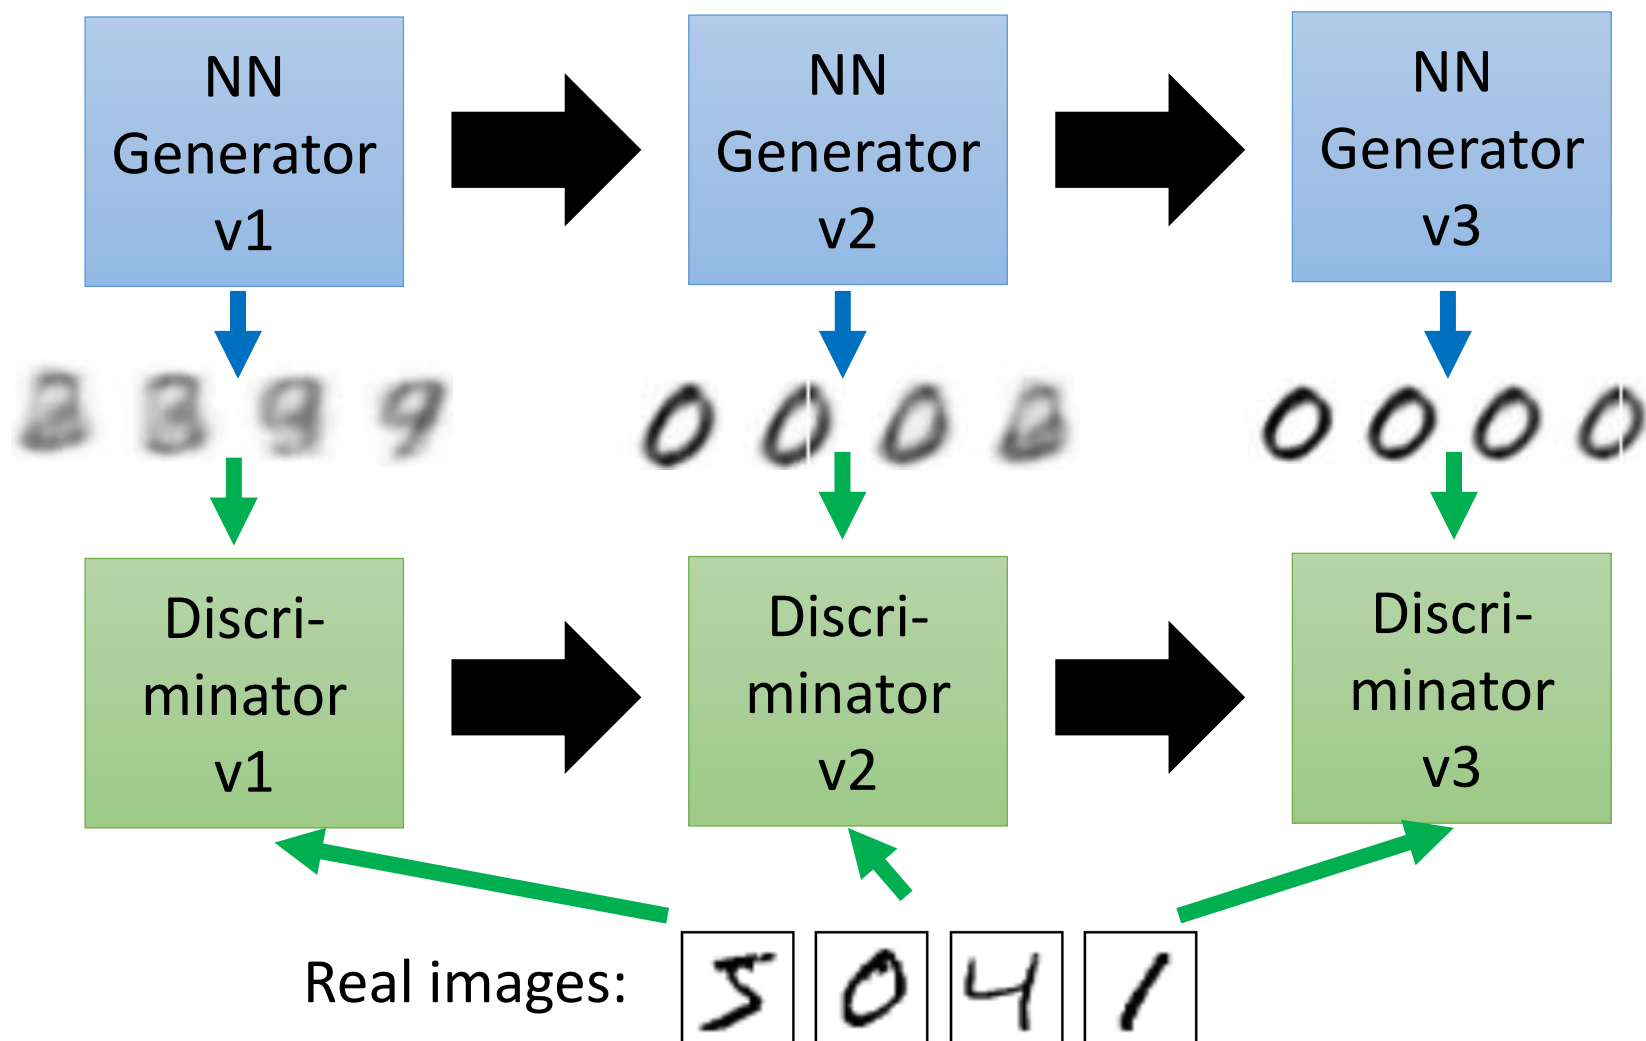
.....



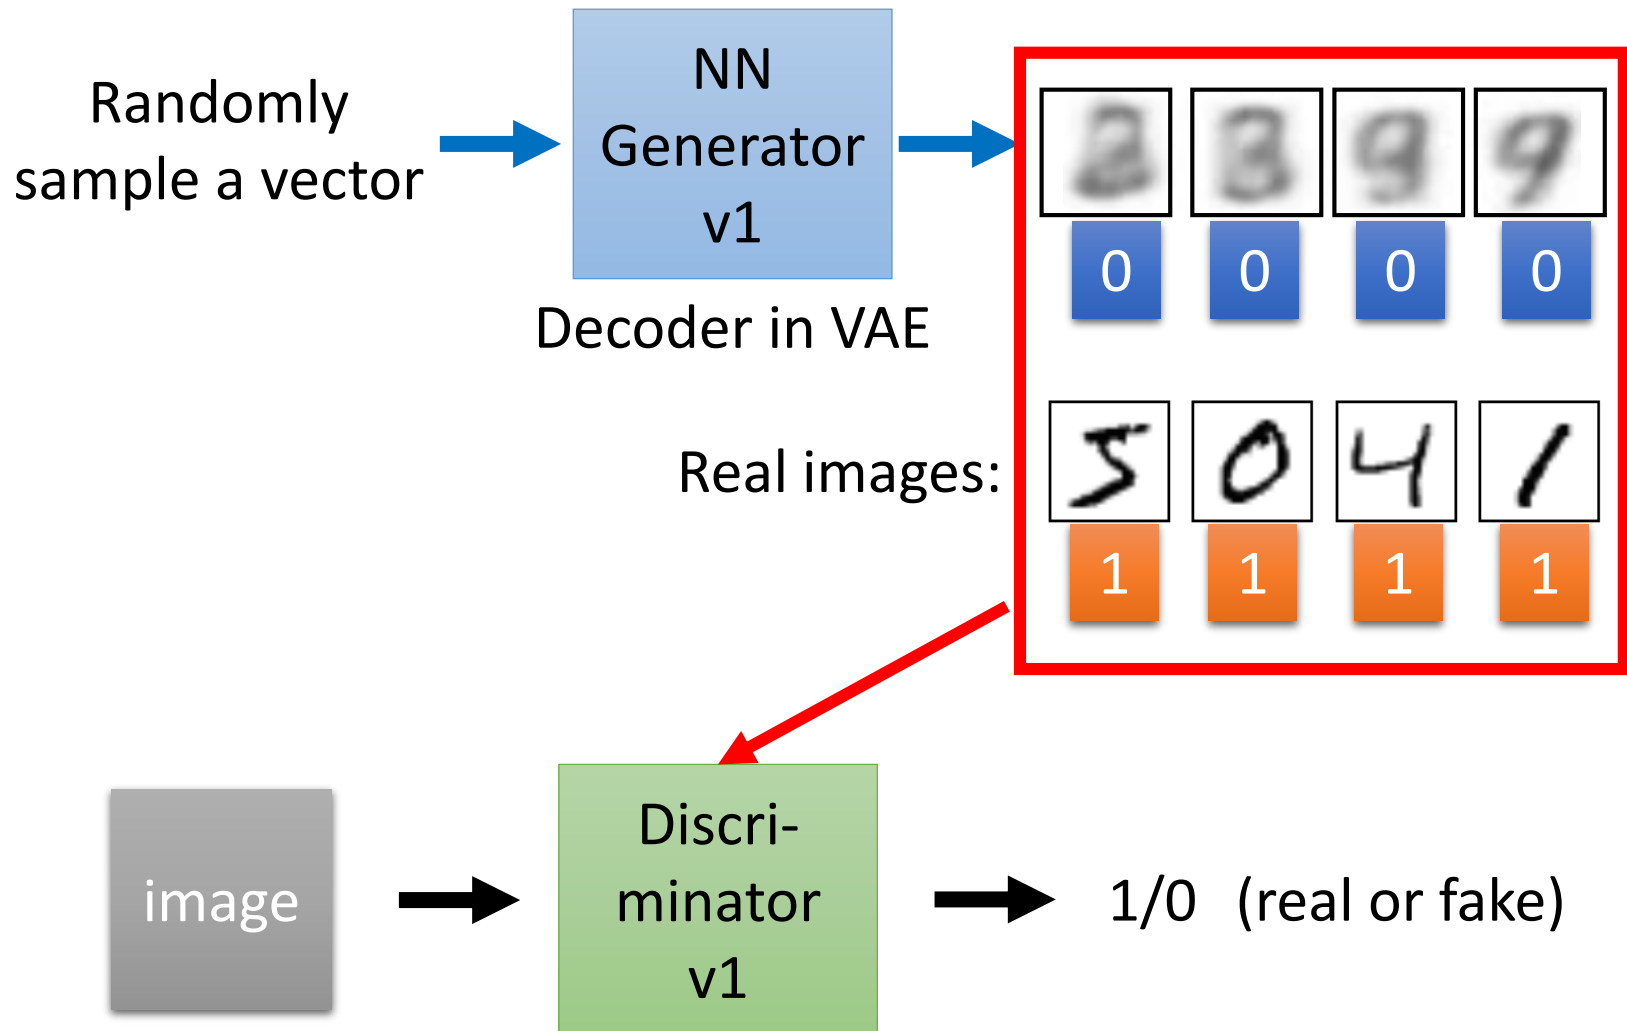
# The evolution of generation



# The evolution of generation



# GAN - Discriminator




# GAN - Generator

Updating the parameters of generator

➔ The output be classified as “real” (as close to 1 as possible)



Generator + Discriminator  
= a network

Using gradient descent to update the parameters in the generator, but fix the discriminator



# Cifar-10

- Which one is machine-generated?






Ref: <https://openai.com/blog/generative-models/>



# Cartoon

- Ref: <https://github.com/mattyachainer-DCGAN>



# Cartoon

- Ref: <http://qiita.com/mattya/items/e5bfe5e04b9d2f0bbd47>



元画像

-赤髪+金髪


-赤目+青目

+制服+セーラー

+笑顔+口開き

+青背景

長髪化ベクトル



一番左のキャラクターが元画像で、  
右に行くほど長髪化ベクトルを強く足している



# To learn more ...

- “Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks”
- “Improved Techniques for Training GANs”
- “Autoencoding beyond pixels using a learned similarity metric”
- “Deep Generative Image Models using a Laplacian Pyramid of Adversarial Network”
- “Super Resolution using GANs”
- “Generative Adversarial Text to Image Synthesis”

# Outline

## Deep Generative Models

- Richard Feynman: “What I cannot create, I do not understand.”

## Deep Reinforcement Learning

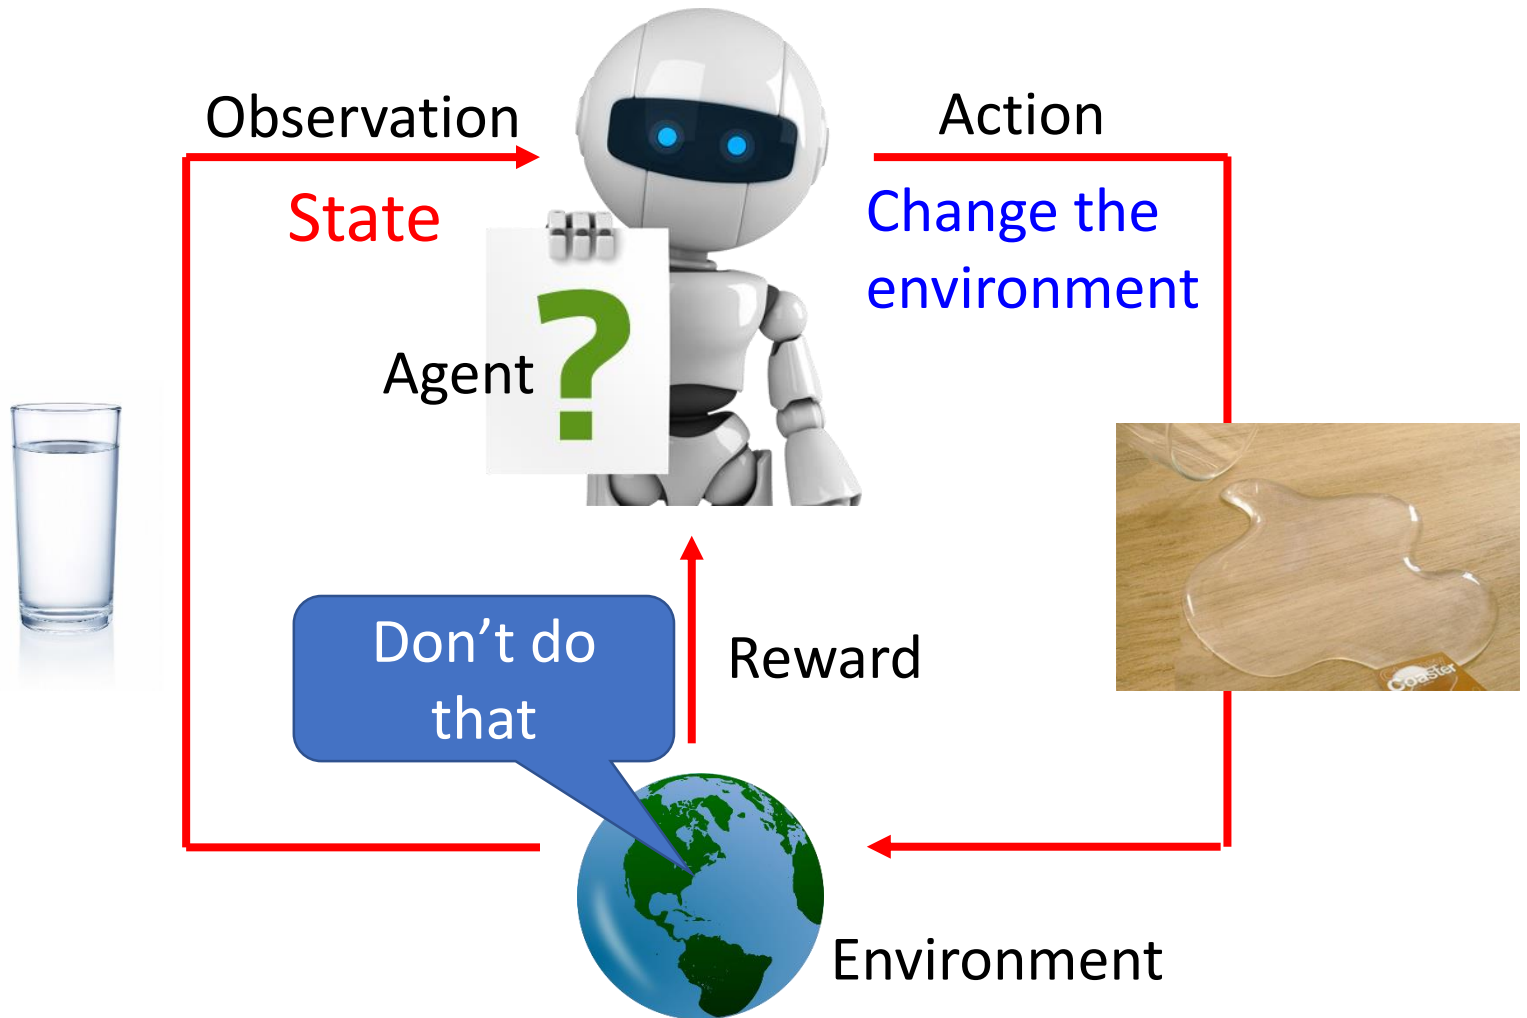
- The technique makes Alpha Go better than professional players.

# Deep Reinforcement Learning



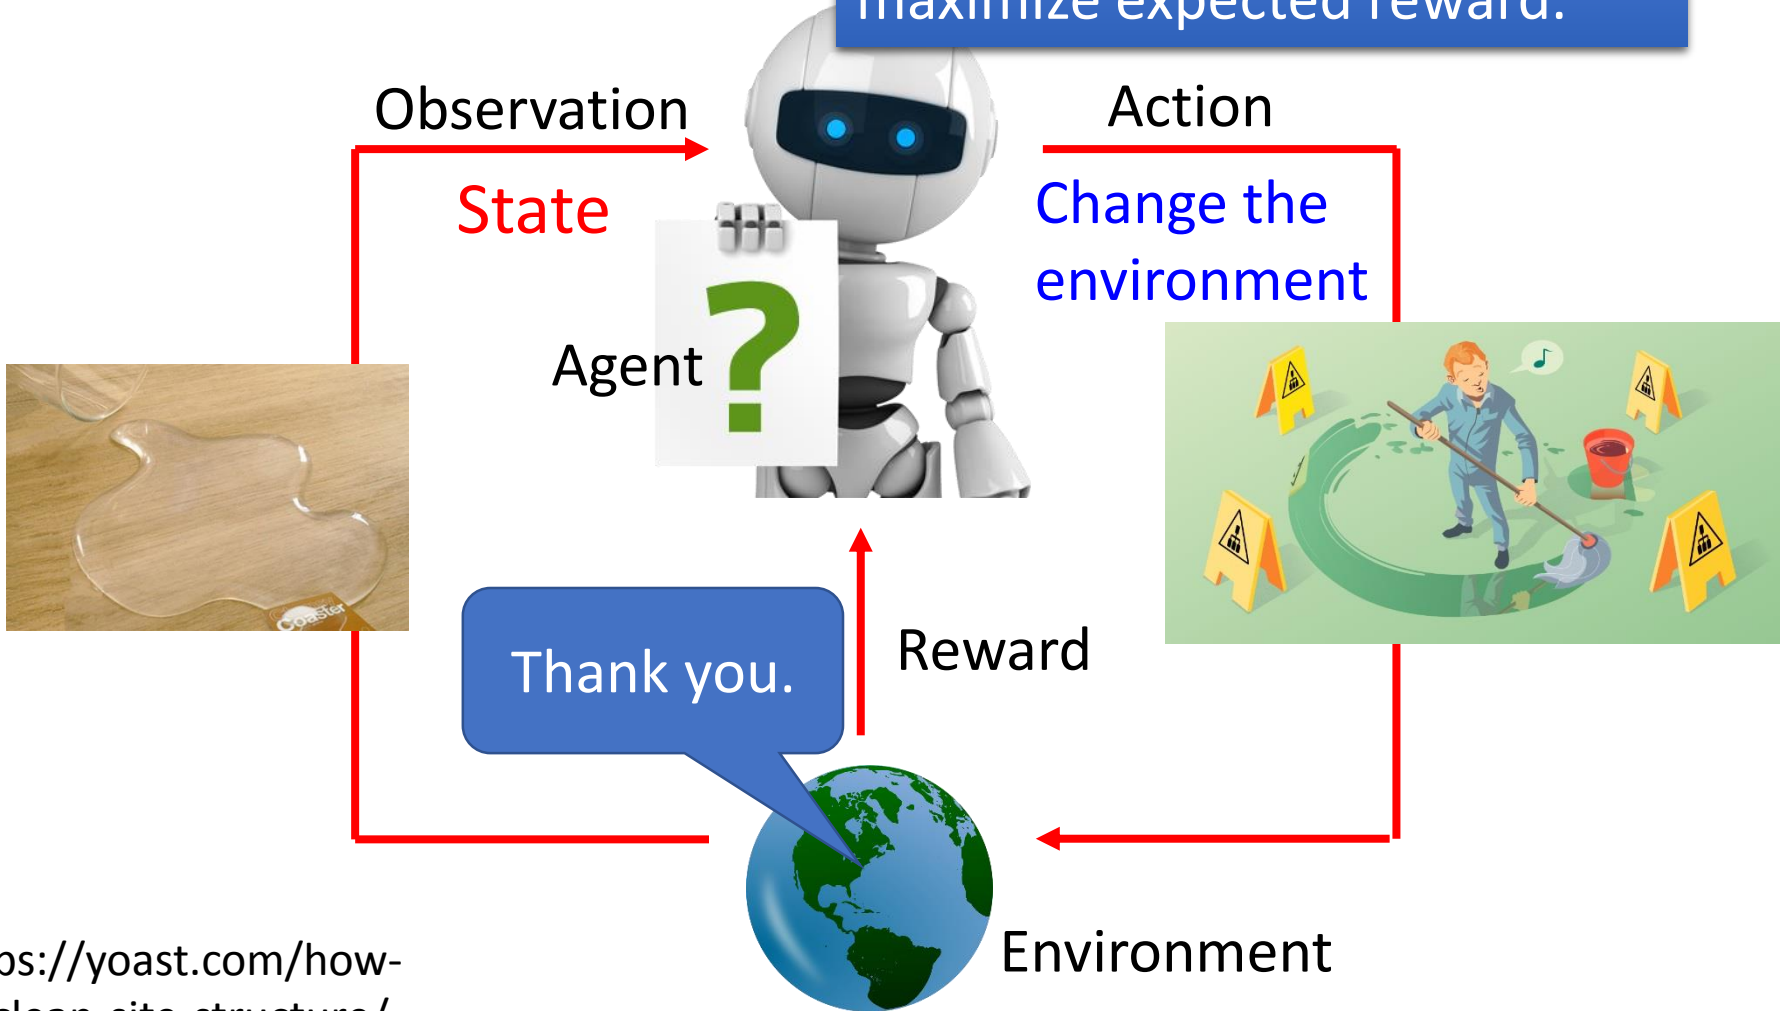
Deep Reinforcement Learning:  $AI = RL + DL$

# Scenario of Reinforcement Learning

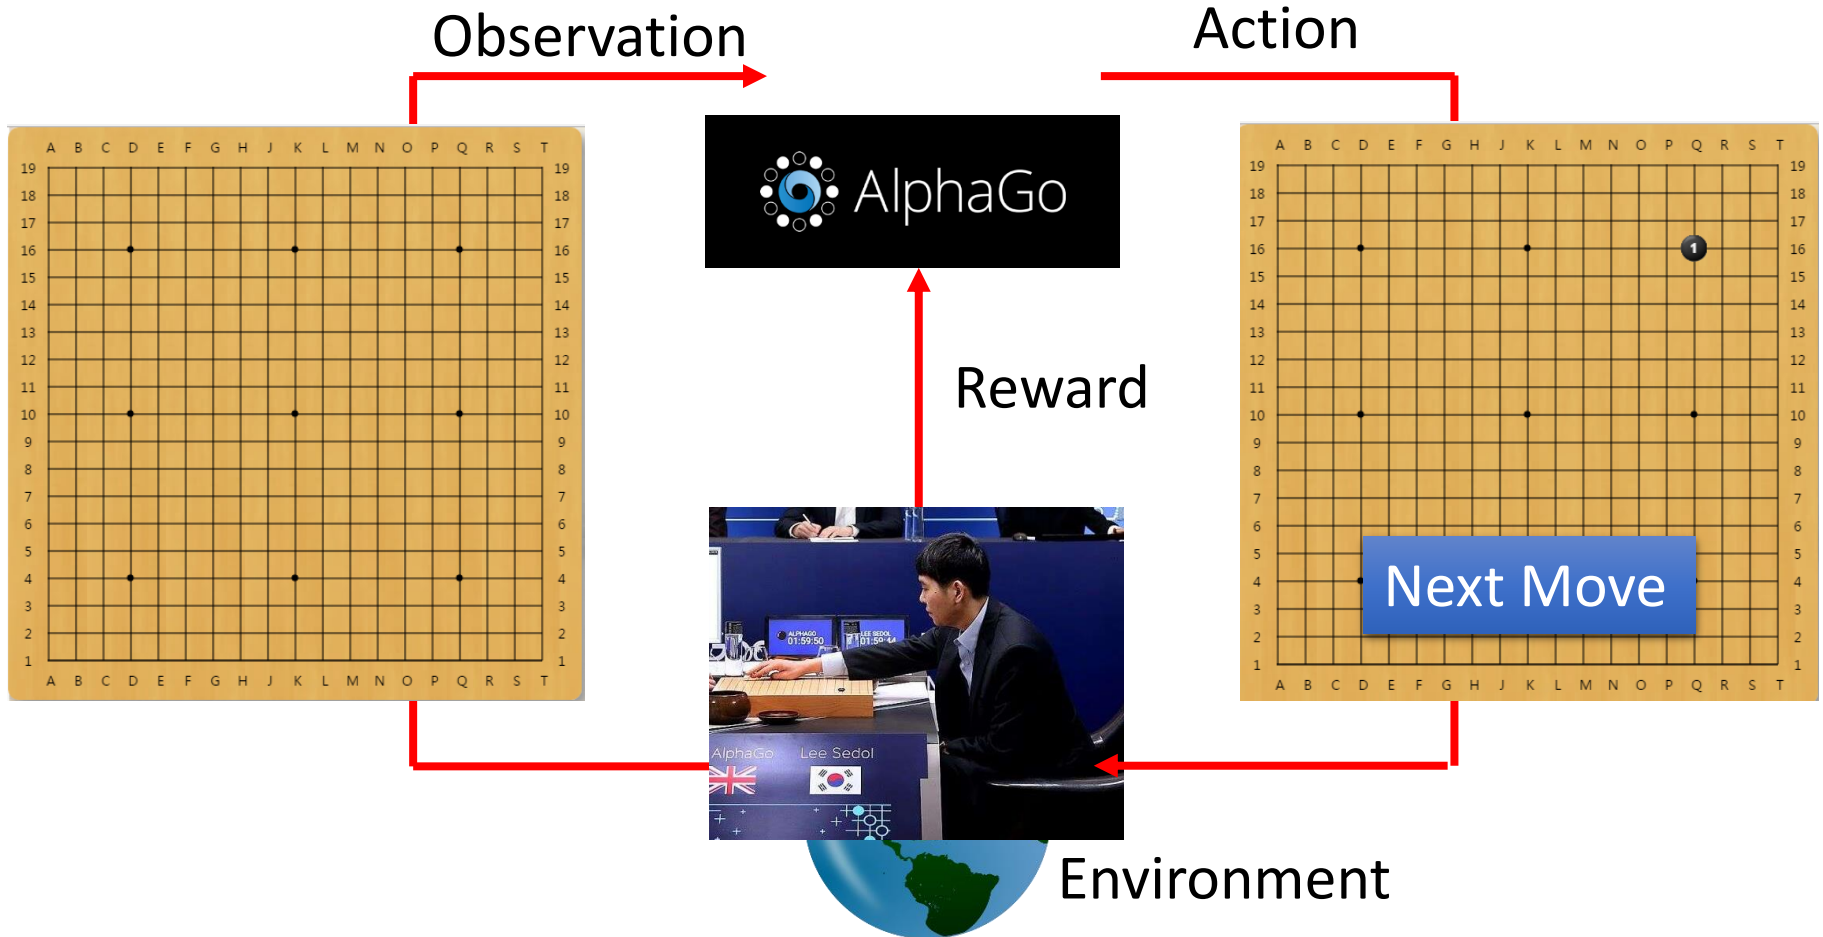


# Scenario of Reinforcement Learning

Agent learns to take actions to maximize expected reward.

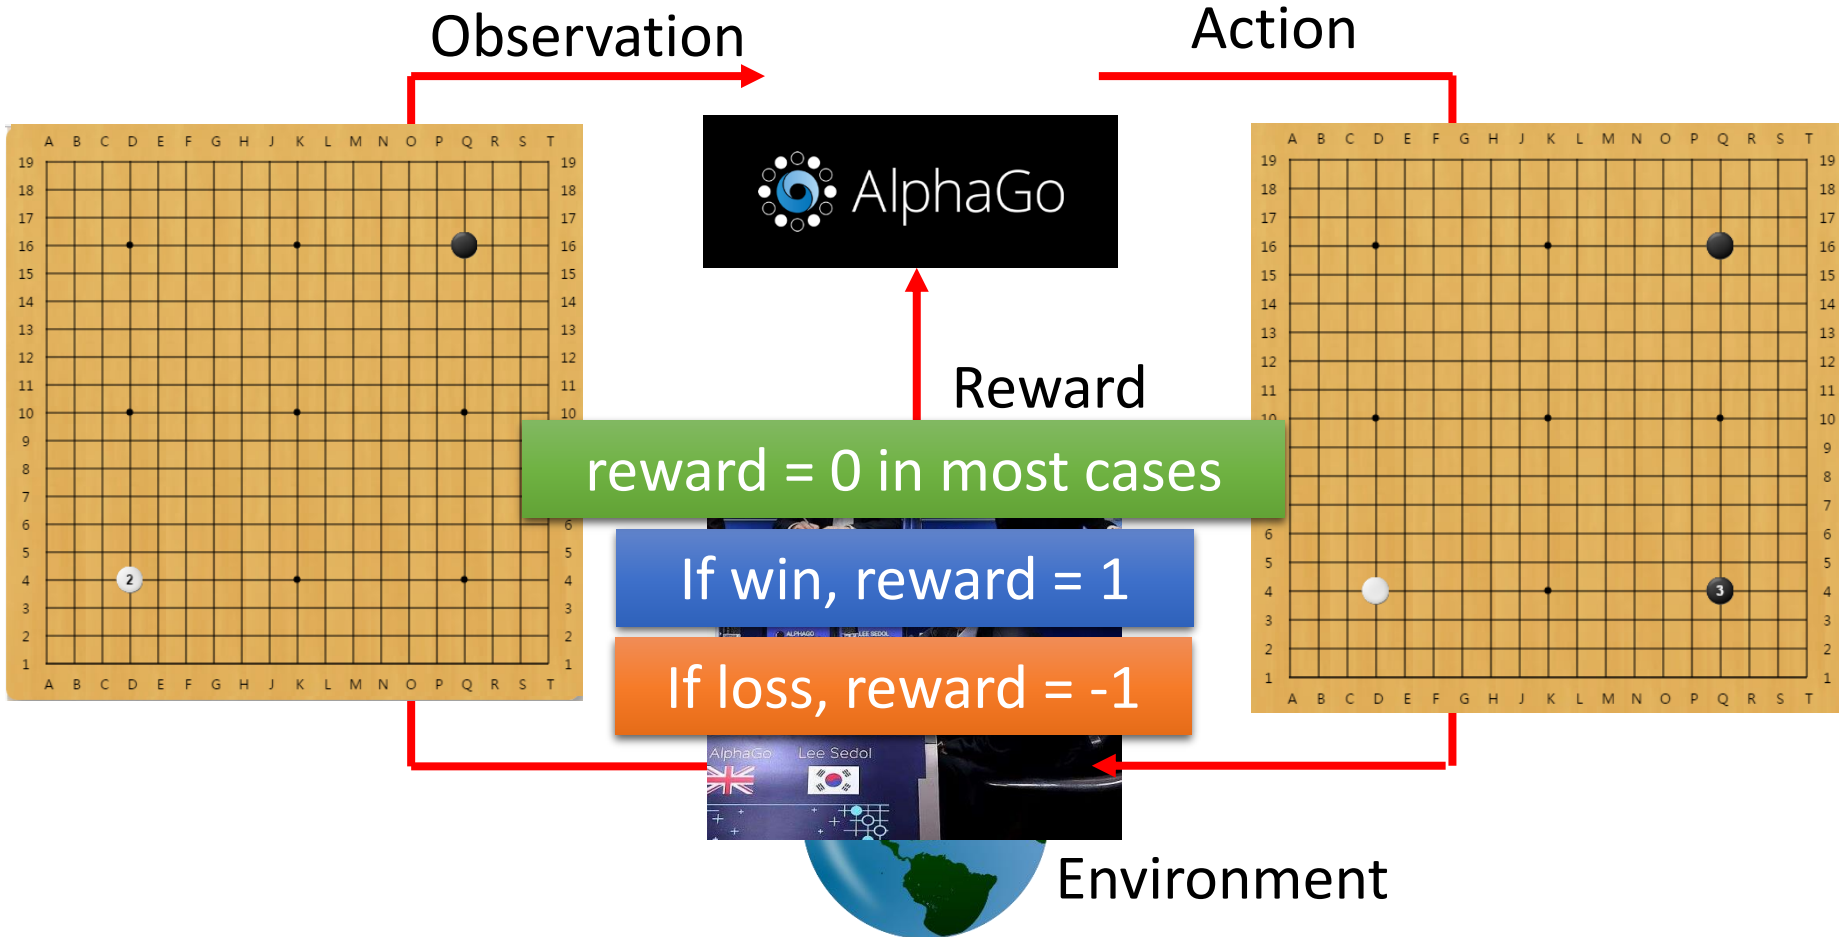


# Learning to play Go



# Learning to play Go

Agent learns to take actions to maximize expected reward.





# Learning to play Go


## - Supervised v.s. Reinforcement

- Supervised: Learning from teacher

Cannot be better than its teacher



Next move:  
"5-5"



Next move:  
"3-3"

- Reinforcement Learning Learning from experience

First move → ..... many moves ..... → Win!

(Two agents play with each other.)

Alpha Go is supervised learning + reinforcement learning.



# Learning a chat-bot

## - Supervised v.s. Reinforcement

- Supervised





Say "Hi"




Say "Good bye"


- Reinforcement



.....



Agent



Agent



Bad

# Learning a chat-bot

## - Reinforcement Learning

- Let two agents talk to each other (sometimes generate good dialogue, sometimes bad)



How old are you?



See you.



How old are you?



I am 16.



See you.



See you.



I thought you were 12.




What make you think so?

# Learning a chat-bot

## - Reinforcement Learning

- By this approach, we can generate a lot of dialogues.
- Use some pre-defined rules to evaluate the goodness of a dialogue

Machine learns from the evaluation



# Application: Interactive Retrieval

- Interactive retrieval is helpful.


[Wu & Lee, INTERSPEECH 16]



“Deep Learning” related to Machine Learning?  
“Deep Learning” related to Education?

# Deep Reinforcement Learning

- Different network depth



# More applications

- Flying Helicopter
  - <https://www.youtube.com/watch?v=0JL04JJjocc>
- Driving
  - <https://www.youtube.com/watch?v=0xo1Ldx3L5Q>
- Google Cuts Its Giant Electricity Bill With DeepMind-Powered AI
  - <http://www.bloomberg.com/news/articles/2016-07-19/google-cuts-its-giant-electricity-bill-with-deepmind-powered-ai>
- Text generation
  - Hongyu Guo, “Generating Text with Deep Reinforcement Learning”, NIPS, 2015
  - Marc'Aurelio Ranzato, Sumit Chopra, Michael Auli, Wojciech Zaremba, “Sequence Level Training with Recurrent Neural Networks”, ICLR, 2016

# To learn deep reinforcement learning .....

- Textbook: Reinforcement Learning: An Introduction
  - <https://webdocs.cs.ualberta.ca/~sutton/book/the-book.html>
- Lectures of David Silver
  - <http://www0.cs.ucl.ac.uk/staff/D.Silver/web/Teaching.html> (10 lectures, 1:30 each)
  - [http://videolectures.net/rldm2015\\_silver\\_reinforcement\\_learning/](http://videolectures.net/rldm2015_silver_reinforcement_learning/) (Deep Reinforcement Learning )
- Lectures of John Schulman
  - [https://youtu.be/aUrX-rP\\_ss4](https://youtu.be/aUrX-rP_ss4)