

---

---

# ML 2026 Spring HW4

# Training Transformers

TA: 劉建豐、馮柏翰、陳品睿

[ntu-ml-2026-spring-ta@googlegroups.com](mailto:ntu-ml-2026-spring-ta@googlegroups.com)

Deadline: 2026/**04/16** 23:59:59 (UTC+8)

---

---

# Outline

- Task Description
- Dataset
- Metric
- Baseline
- Submission & Grading
- Hints

# Useful Links

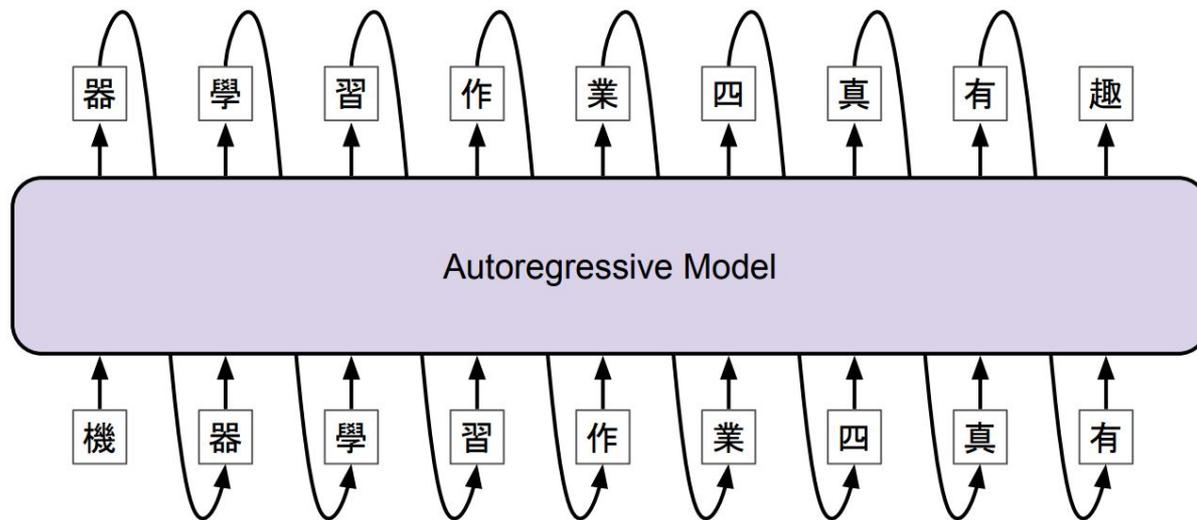
- [Judgeboi](#)
- [Sample code \(colab\)](#)
- [Sample code \(kaggle\)](#)
- Self-attention: [link1](#), [link2](#)
- Transformers: [link1](#), [link2](#)
- [Pytorch Tutorial](#)

# Task Description

- Using a transformer decoder-only model for training, focusing on next-token prediction with Pokémon images.
- Goal: Learn how to use current LM architecture to do next token prediction.

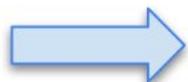
# Task Description

- Next-token prediction is a fundamental concept in language modeling, involving the prediction of the most likely next word (or token) in a sequence based on the preceding context



# Task Description

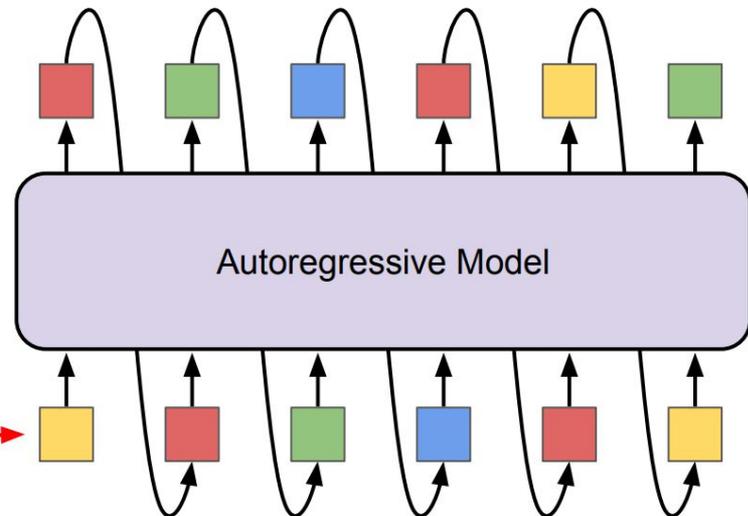
- Training: Next token prediction



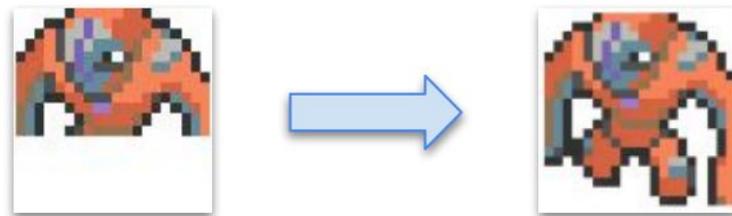
Pixel



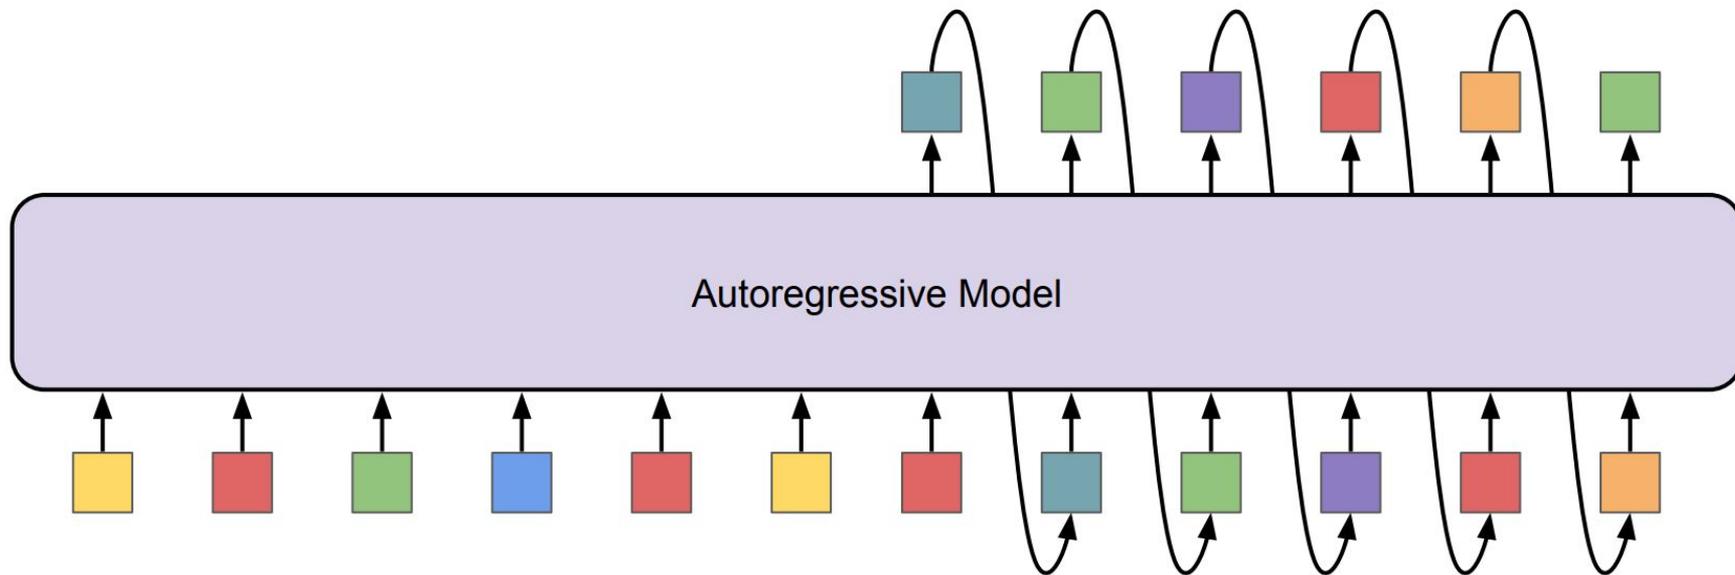
Pixel



# Task Description



- Testing: Given 60% of an image, predict the remaining part.



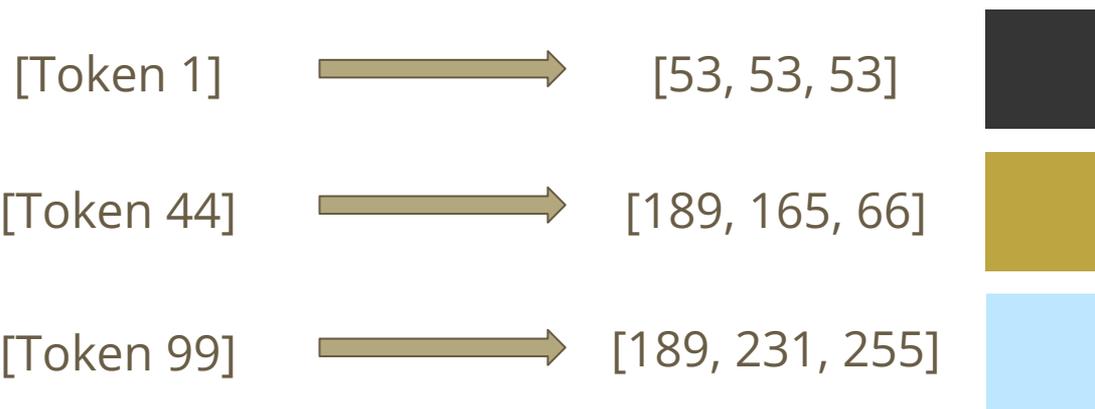
# Dataset

- 792 Pokémon's small images



# Dataset

- For each image in dataset, it is consist of  $20 * 20$  numbers (tokens)
- Each number (token) can map a color pixel via color map
  - For example, token 0 represents a white pixel

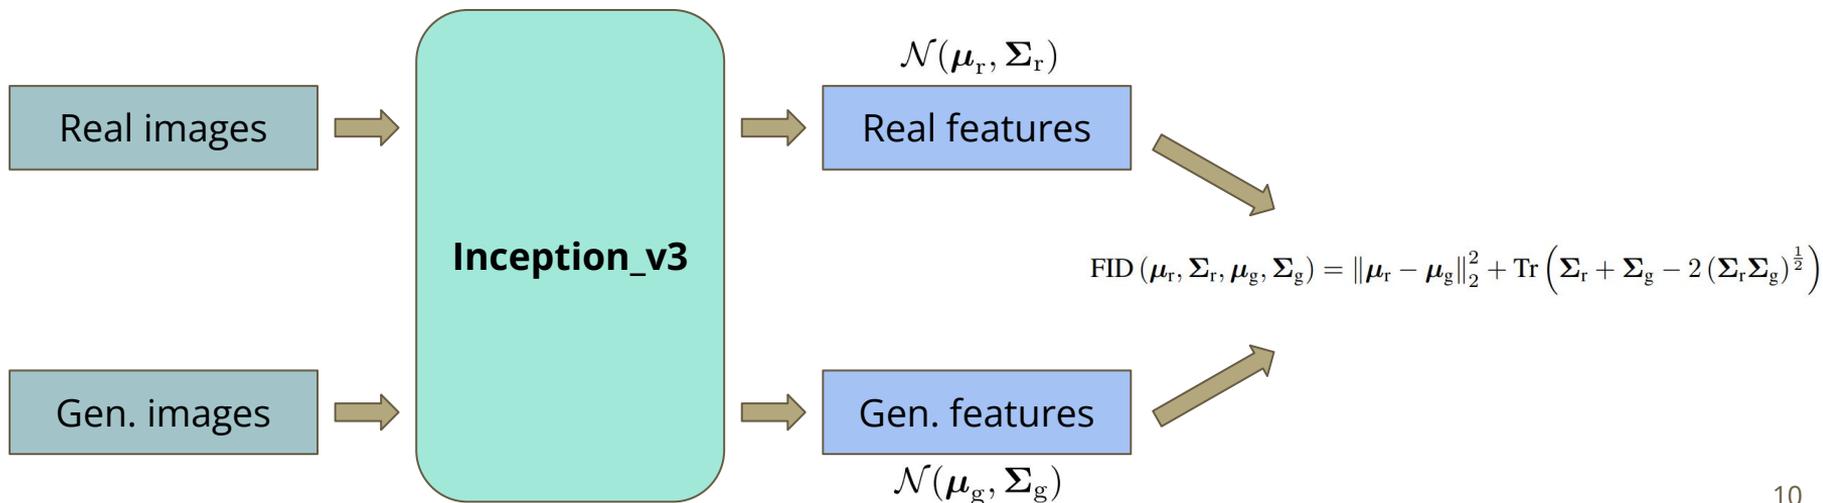


# Dataset

- Total Images: 792
- Dataset Split:
  - Train: 632
  - Validation: 80
  - Test: 80
- Image Size:  $20 \times 20 = 400$  pixels
- Number of Classes (Pixel Colors): 167

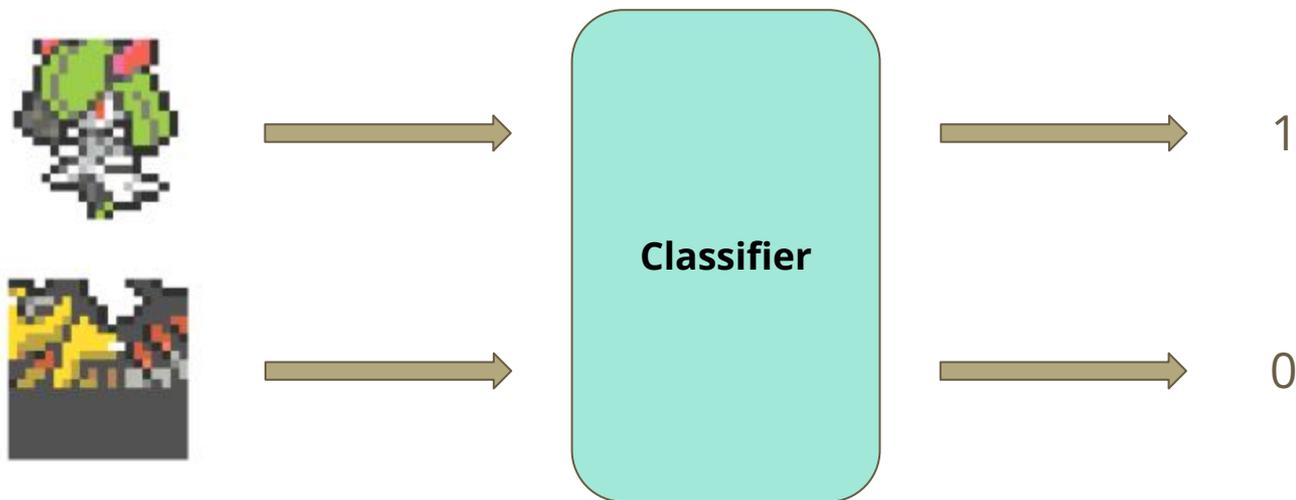
# Metric

- Fréchet inception distance (FID)
  - Use inception model to capture the mean and standard deviation of features
  - Calculate the Fréchet distance between distribution of real images and generated ones



# Metric

- Pokémon Detection Rate (PDR)
  - Use a classifier (trained by TAs) to determine whether the generated images are pokemon or not



# Baseline

	FID	PDR	Estimated Training Time	Score
Public simple baseline	$\leq 86.00$	$\geq 0.1$	~ 10 mins	+1pt
Private simple baseline	$\leq 86.00$	$\geq 0.1$		+1pt
Public medium baseline	$\leq 80.00$	$\geq 0.5$	~ 20 mins	+1pt
Private medium baseline	$\leq 80.00$	$\geq 0.5$		+1pt
Public strong baseline	$\leq 73.00$	$\geq 0.85$	30 ~ 40 mins	+1pt
Private strong baseline	$\leq 73.00$	$\geq 0.85$		+1pt
Code submission	—	—	—	+4pt

**Reminder: You need to pass FID and PDR at the same time to pass baseline.**

## Hint – Simple Baseline

- By executing the sample code, you can build a decoder-only transformer-based model, GPT-2, to perform next token prediction for generating pokemon image

## Hint – Medium Baseline

- Adjust the number of epochs and learning rate in the sample code
- Modify the hyperparameters in the Model configuration in the sample code
  - Number of attention head
  - Dimension of embedding layer
  - Number of embedding layers

## Hint – Strong Baseline

- Try different model architectures like Llama and Mistral. (Recommended)
- Train a classifier to determine whether an image looks like a Pokémon.

(Optional)

- The sample code saves the checkpoint with the lowest training loss, since the reconstruction accuracy on the validation set does not directly indicate the model's image generation capability.
- Discriminator in GAN

## Hint – Strong Baseline

- Import the corresponding Config class from the Transformers package and set the appropriate hyperparameters for training.
  - [Llama](#)
  - [Mistral](#)

## Hint – Strong Baseline

- The sample code saves the checkpoint with the lowest training loss since validation set reconstruction accuracy doesn't directly reflect the model's image generation ability.
- One solution is to train a classifier (discriminator) to determine whether an image looks like a Pokémon and use it to select the best checkpoints.



construct binary classification dataset



Classifier

## Hint – Parameters Range

Number of epochs	30 ~ 150 (default: 50)
Learning rate	1e-5 ~ 1e-2 (default: 1e-3)
Batch size	8 ~ 64 (default: 16)
Weight decay	0.1 ~ 1e-5 (default: 0.1)
Atten. head number	1 ~ 12 (default: 2)
Emb. dimension	32 ~ 512 (default: 64)
Emb. layer number	1 ~ 12 (default: 2)

# Submission – Judgeboi

- You should submit only **one .txt** file to the Judgeboi
  - Other types of filename extension (like .zip, .jpg or .png) are not allowed
- **5** submission quota per day, reset at **23:59 (UTC+8)**.
- Each submission uploaded to JudgeBoi will be evaluated with a time limit of 10 minutes.
  - If your submission isn't evaluated due to time limitations, you can submit your file later



# Submission – Code Submission

- Submit **ML2026Spring\_hw4.ipynb** to NTU COOL
- **Deadline: 4/16 (Thu.) 23:59**
- Compress your code into zip if not using .ipynb file
- **No late submission is allowed**
- We can only see your last submission.
- Do not submit model weights or dataset.
- **If your code is not reasonable or reproducible, you will receive 0 points for this homework.**

# Submission & Deadline

- Submit your source code to **NTU Cool**
- 2026/**04/16** 23:59:59 (UTC+8)
- No late submission is allowed

# Grading Release Date

- The grading of the homework will be released by 2026/**04/19** 23:59:59 (UTC+8)

# Grading - Regulations

- You should NOT plagiarize, if you use any other resource, you should cite it in the reference.
- Do NOT share codes or prediction files with any living creatures.
- Do NOT use any approaches to submit your results more than 5 times a day.
- Do NOT search for or use additional data for training or the answers for the testing data.
- Do NOT use closed-source LLM APIs like GPT-4, Gemini, etc.
- You should NOT modify your input file or prediction files manually.
- Make sure that TAs can reproduce the predictions using the code you submit. (Fix the random seed)
- Your final grade  $\times 0.9$  and get a score 0 for that homework if you violate any of the above rules first time (within a semester).
- You will get F for the final grade if you violate any of the above rules multiple ( $> 1$ ) times.
- Prof. Lee & TAs preserve the rights to change the rules & grades.

# If You Have Any Questions

- NTU Cool **HW4** 作業討論區
  - 如果同學的問題不涉及作業答案或隱私, 請**一律使用** NTU Cool 討論區
  - 助教們會優先回答NTU Cool討論區上的問題
- Email: [ntu-ml-2026-spring-ta@googlegroups.com](mailto:ntu-ml-2026-spring-ta@googlegroups.com)
  - Title should start with [ML 2026 Spring **HW4**]
  - Email with the wrong title will be moved to trash automatically
- TA Hours
  - Each Friday before / after class:
    - (Fri.) 13.20 ~ 14.10 / 17:30~18:00
    - Location: 博理112