Homework 4
Training Transformer

TAs: 李晨安 林毓翔 林熙哲

Email: ntu-ml-2025-spring-ta@googlegroups.com

Deadline: 2025/4/11 23:59:59 (UTC+8)

Outline

- Task Description
- Dataset
- Baselines
- Submission & Grading
- Appendix

Task Description

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.

Next token prediction

- 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.

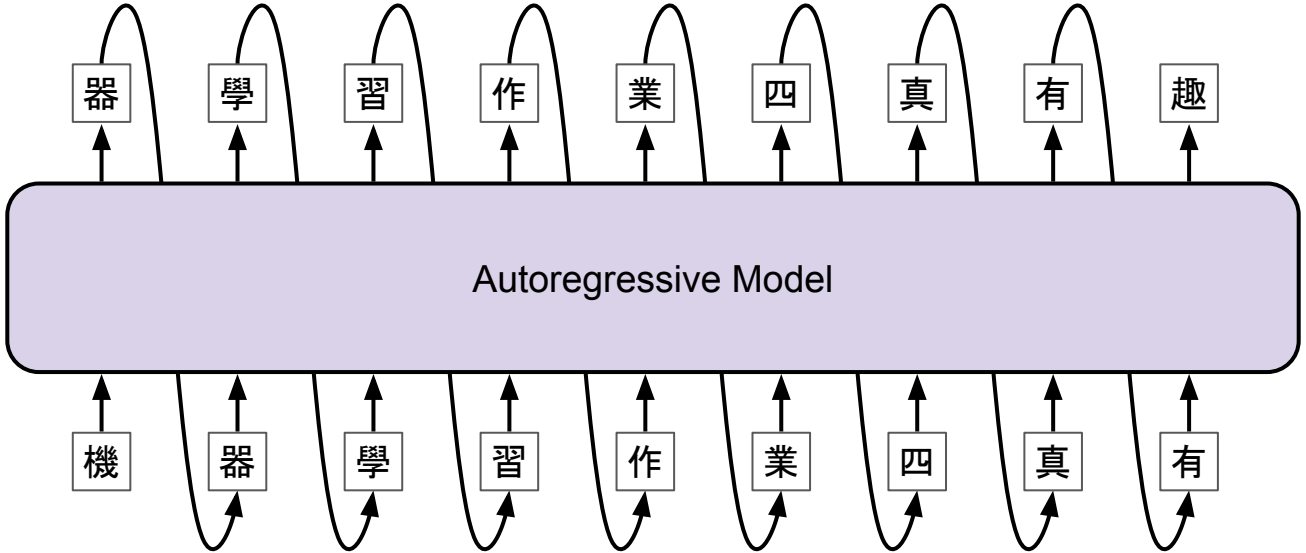

Pokémon Creation

- Training: Next token prediction

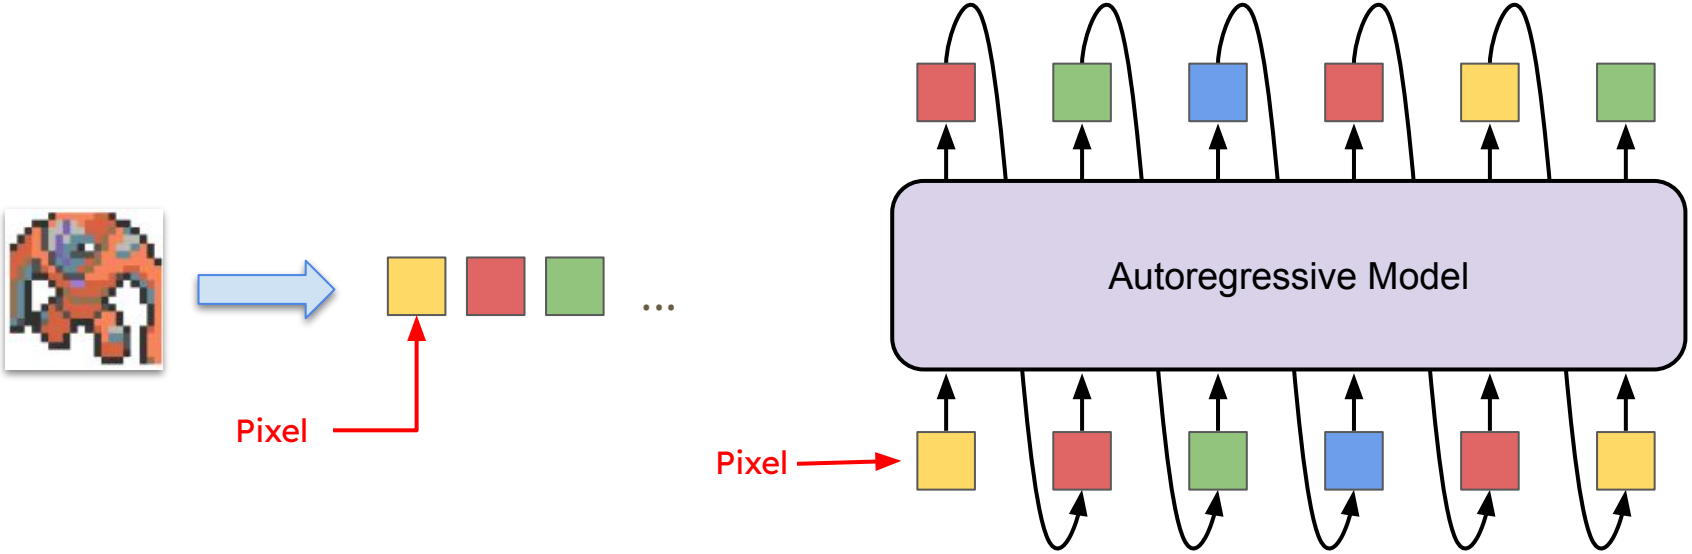

Pokémon Creation

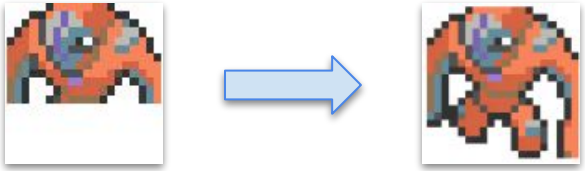

- Testing: Given 60% of an image, predict the remaining part.

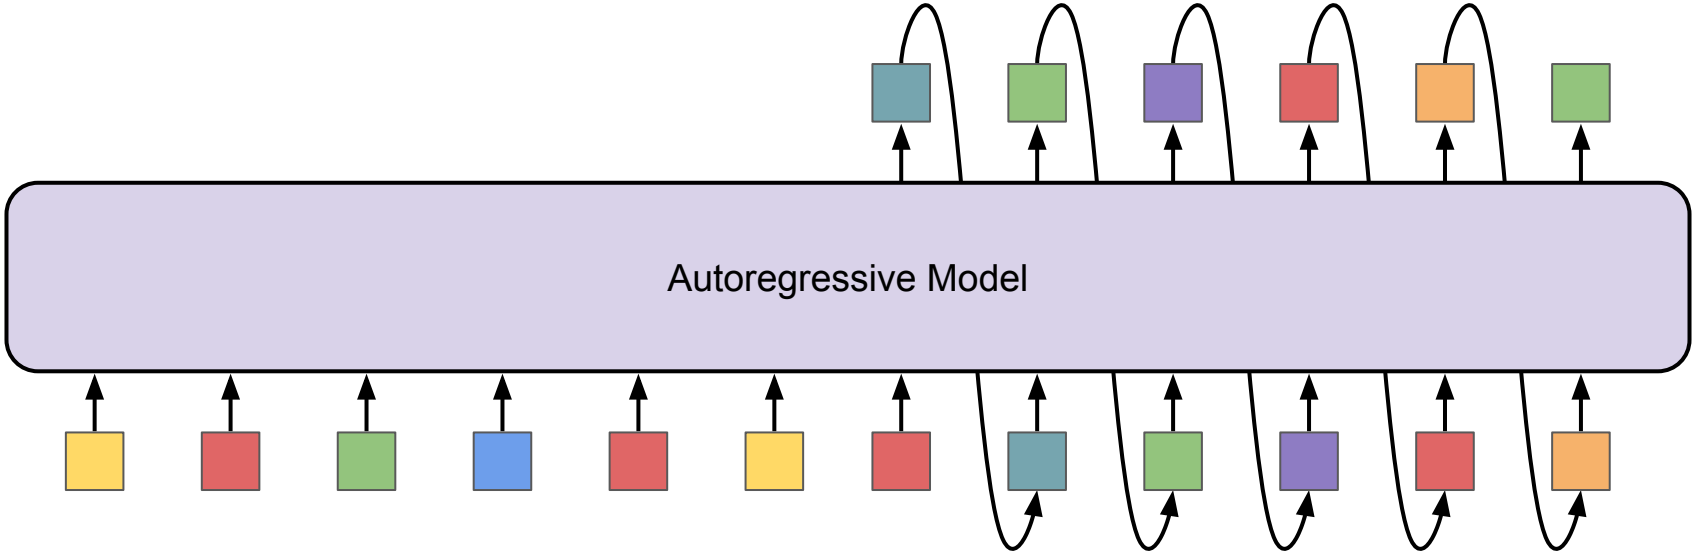

Metrics

FID (Frechet Inception Distance)

- Use Inception Net to create features for real and fake images
- Calculate the Frechet distance between distribution of two features, the lower the better.

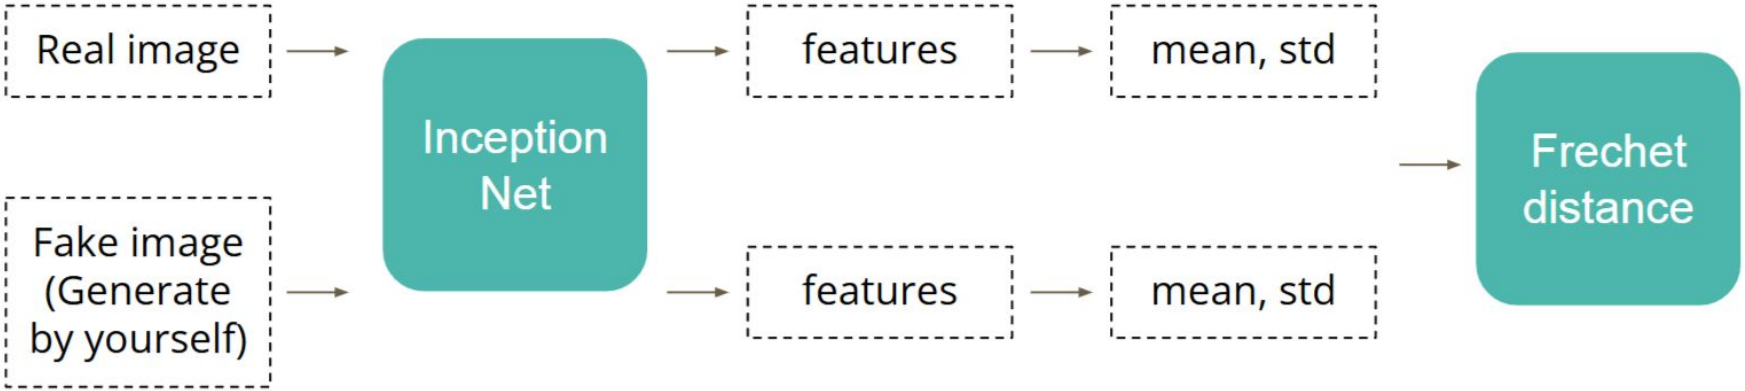

Metrics

PDR (Pokémon Detection Rate)

- Use a classifier to count how many images in your submission look like a Pokémon, the higher the better.

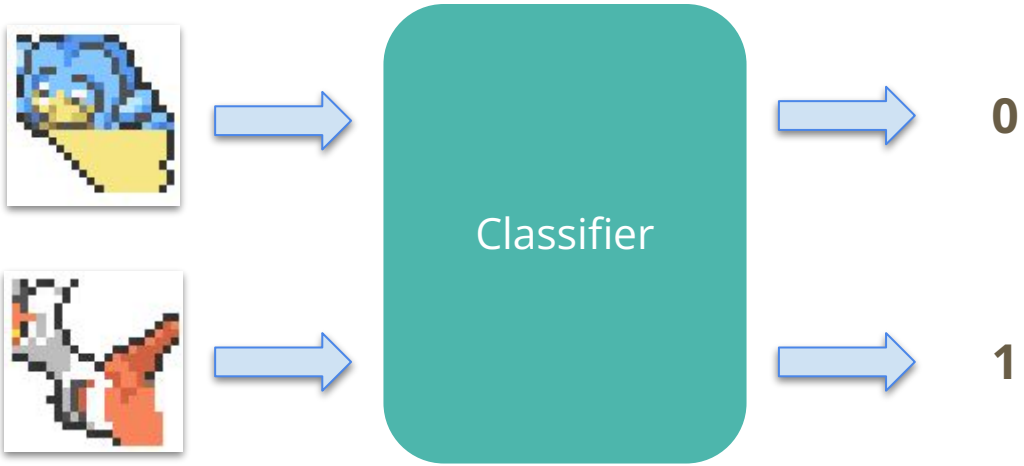

Dataset

Dataset

- Dataset: Small images of 792 Pokémon's (20x20)

Dataset

- Each image has 400 numbers (20 × 20), with each number representing a color.

id string · *lengths* 7⇔8 1.6%	pixel_color sequence · *lengths* 400 100%
train_0	[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 2, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 3, 4, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 5, 6, 4, 0, 0, 0, 0, 0, 0, 0, 0, 1, 2, 0, 0, 0, 0, 0, 0, 1, 3, 3, 7, 2, 8, 2, 0, 0, 0, 0, 1, 9, 10, 0, 0, 0, 0, 0, 0, 11, 3, 7, 12, 13, 14, 14, 8, 0, 0, 1, 9, 15, 4, 0, 0, 0, 0, 0, 0, 0, 1, 13, 13, 14, 14, 14, 14, 8, 2, 9, 0, 10, 0, 0, 0, 0, 0, 0, 0, 0, 16, 14, 14, 14, 9, 13, 14, 14, 17, 6, 18, 4, 0, 0, 0, 0, 0, 0, 0, 16, 12, 13, 13, 14, 19, 2, 14, 17, 12, 0, 0, 0, 10, 0, 0, 0, 0, 0, 1, 9, 0, 0, 0, 13, 13, 20, 14, 17, 6, 21, 21, 22, 0, 0, 0, 0, 0, 0, 23, 9, 24, 0, 0, 12, 13, 14, 14, 14, 17, 3, 4, 25, 2, 0, 0, 0, 0, 0, 0, 11, 9, 12, 12, 13, 13, 13, 13, 26, 0, 1, 27, 14, 14, 8, 2, 0, 0, 0, 0, 0, 11, 1, 13, 13, 13, 13, 13, 8, 0, 5, 14, 14, 14, 13, 13, 8, 0, 0, 0, 0, 0, 0, 1, 3, 18, 13, 13, 13, 8, 14, 14, 14, 13, 13, 13, 13, 8, 0, 0, 0, 0, 0, 23, 3, 28, 0, 13, 13, 13, 29, 13, 13, 13, 16, 5, 13, 13, 16, 0, 0, 0, 0, 0, 1, 9, 24, 0, 13, 13, 13, 13, 3, 16, 0, 0, 23, 5, 0, 0, 0, 0, 0, 0, 23, 3, 12, 24, 15, 18, 3, 6, 16, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 23, 3, 6, 6, 6, 16, 30, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 23, 1, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]

id int64 0 166	color sequence · *lengths* 3 R G B 3
0	[255, 255, 255]
1	[53, 53, 53]
2	[49, 49, 49]
3	[186, 186, 186]
4	[51, 51, 51]
5	[54, 54, 54]
6	[187, 187, 187]
7	[83, 83, 83]

lca0503/ml2025-hw4-pokemon

lca0503/ml2025-hw4-colormap

Data Statistics

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

Baselines

Simple Baseline

- Build a transformer decoder-only model for next-token prediction using Pokémon images with sample code.

Estimate training time:
10 mins using T4 on Colab

Medium Baseline

- Modify the hyperparameters in the Model configuration in the sample code.
- Adjust the number of epochs and learning rate in the sample code.

Estimate training time:
20 mins using T4 on Colab

Strong Baseline

- Try different model architectures like Llama, Mistral, or Qwen. (Recommended)
  - The sample code uses GPT-2.
- 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.

Estimate training time:
30 mins using T4 on Colab

Hints: Different model architectures (Recommended)

- Import the corresponding Config class from the Transformers package and set the appropriate hyperparameters for training.
- LlamaConfig
  - Link
- MistralConfig
  - Link
- Qwen2Config
  - Link

Hints: Train a classifier (Optional)

- 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 to determine whether an image looks like a Pokémon and use it to select the best checkpoints.
- This step is optional because you can still pick a good checkpoint through human evaluation, and training time is short.

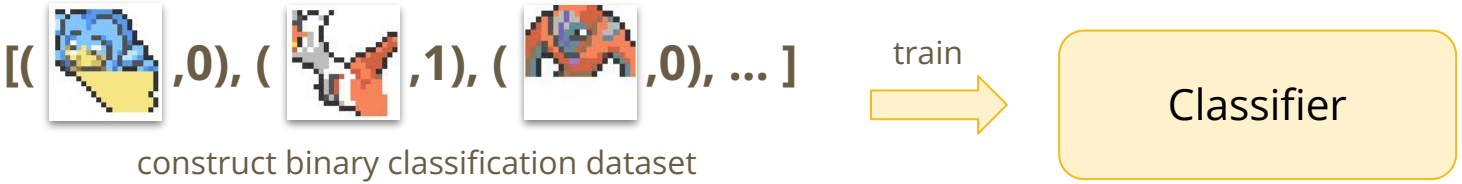

Submission & Grading

Submission - JudgeBoi

- Only *.txt file is 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.

Submission - JudgeBoi

- Each submission file must contain exactly 80 lines, with each line representing an image.
- Each image consists of 400 numbers.

Each line contains 20 x 20 = 400 numbers.

Each line represents an image.

Submission - NTU COOL

- Submit your code to NTU COOL
  - We can only see your **last submission**.
  - Do **NOT** submit the model checkpoint or dataset.
  - If your code is not reproducible, your final grade will be multiplied by **0.9**.
  - You should compress your code into a single zip file:
    - ex. b12901000_hw4.zip

<Student ID>_hw4.zip

Grading

	FID	PDR	Estimate training time	Score
Public Simple Baseline	**≤ 84.50**	**≥ 0.1**	**10 mins**	**+ 1pt**
Private Simple Baseline	**≤ 84.50**	**≥ 0.1**		**+ 1pt**
Public Medium Baseline	**≤ 81.00**	**≥ 0.5**	**20 mins**	**+ 1pt**
Private Medium Baseline	**≤ 81.00**	**≥ 0.5**		**+ 1pt**
Public Strong Baseline	**≤ 73.00**	**≥ 0.85**	**30 mins**	**+ 1pt**
Private Strong Baseline	**≤ 73.00**	**≥ 0.85**		**+ 1pt**
Code Submission	**-**	**-**	**-**	**+ 4pts**

***Your results must meet both FID and PDR requirements to pass the baseline.**

Deadline

- JudgeBoi: **2025/04/11 23:59 (UTC+8)**
- NTU COOL: **2025/04/11 23:59 (UTC+8)**

Regulation

- You should NOT plagiarize, if you use any other resource, you should cite it in the reference.
- You should NOT modify your prediction files manually.
- 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 or use additional data or pre-trained models.**
- **The generated image can only be used for this homework.**
- Your **final grade x 0.9** and **get a score 0 for that homework** if you violate any of the above rules **first time (within a semester)**.
- Your will **get F for the final grade** if you violate any of the above rules **multiple times (within a semester)**.
- Prof. Lee & TAs preserve the rights to change the rules & grades.

Appendix

Links

- Colab: link
- Dataset: link1, link2
- JudgeBoi: link

If any questions, you can ask us via...

- NTU COOL (Recommended)
- Email
  - ntu-ml-2025-spring-ta@googlegroups.com
  - The title should begin with “[HW4]”
- TA hour
  - Friday, 13:30 ~ 14:20
  - Friday, After Course ~ 18:00