

HOMework 623

PAPER PRESENTATION *

10-423/10-623 GENERATIVE AI
<http://423.mlcourse.org>

OUT: 03/28/24

DUE: 04/09/24

TAs:

1 Paper Presentation

For this assignment you will read a recent generative AI paper and video record a (short, 5-10 minute) presentation about it.

We will follow the role-playing methodology from Collin Raffel and Alec Jacobson.

<https://colinraffel.com/blog/role-playing-seminar.html>

You may select one of the following roles and present your paper in that format. (The instructions below were copied and adapted from 10-719's use of the role-playing paper presentation approach.)

1. **Scientific Peer Reviewer:** Complete a full review of the paper as if it were submitted to a conference. Follow the [guidelines for NeurIPS reviewers](#) to produce your review. In particular, please answer Questions 1 to 10 under "Review Form", including assigning an overall score.
2. **Archaeologist:** Determine where this paper sits in the context of previous and subsequent work. Find and briefly report on both: (1) a prior paper that substantially influenced the current paper, and (2) a more recent paper that cites this current paper.
3. **Academic Researcher:** You're an academic researcher working on a new project in this area. Propose an imaginary follow-up project that builds on the current paper. Pretend that this new project has been successful, and write up a brief introduction for a paper about your project using the five-point structure provided [here](#) (under "The Introduction"). You do not need to actually *write* the introduction, but instead should present in the style of an introduction to this new project.
4. **Industry Practitioner:** You work at a company or organization developing an application or product of your choice (one that has not already been suggested in a prior session). Describe the application/product in detail, and bring a convincing pitch for why you should be paid to implement the method in the paper for this particular application.
5. **Private Investigator:** You are a detective who needs to run a background check on one of the paper's authors. Where have they worked? What did they study? What previous projects might have led to working on this one? What motivated them to work on this project?

*Compiled on Thursday 28th March, 2024 at 21:55

6. **Social Impact Assessor:** Identify how this paper self-assesses its positive or negative impact on the world. Have any additional positive social impacts been left out? What are possible negative social impacts that were overlooked or omitted? Please read [this short paper](#) to see examples.

A few notes:

- For each of the above roles, your presentation *must* start with a summary of the paper.
- You are welcome to record with your webcam turned off (Khan Academy style).
- After submission, we will share all the HW623 video presentations with the rest of class via Piazza, so others can learn from you as well.
- You will make slides and submit them alongside your video recording.

2 Select a paper

To select a paper, you should pick of the ones below. If you would like to present an generative AI paper not on this list, send a Private Note to the "Instructors" with the paper and 1-2 sentences about why you are interested in it.

- [1] Yuntao Bai et al. *Constitutional AI: Harmlessness from AI Feedback*. 2022. arXiv: [2212.08073 \[cs.CL\]](#).
- [2] Keqin Chen et al. *Shikra: Unleashing Multimodal LLM's Referential Dialogue Magic*. 2023. arXiv: [2306.15195 \[cs.CV\]](#).
- [3] Zhengyang Geng, Ashwini Pople, and J. Zico Kolter. *One-Step Diffusion Distillation via Deep Equilibrium Models*. 2023. arXiv: [2401.08639 \[cs.CV\]](#).
- [4] Rohit Girdhar et al. *ImageBind: One Embedding Space To Bind Them All*. 2023. arXiv: [2305.05665 \[cs.CV\]](#).
- [5] Yutong He, Ruslan Salakhutdinov, and J. Zico Kolter. *Localized Text-to-Image Generation for Free via Cross Attention Control*. 2023. arXiv: [2306.14636 \[cs.CV\]](#).
- [6] Zahra Kadkhodaie, Florentin Guth, Eero P. Simoncelli, and Stéphane Mallat. *Generalization in diffusion models arises from geometry-adaptive harmonic representations*. 2024. arXiv: [2310.02557 \[cs.CV\]](#).
- [7] Diederik P. Kingma and Prafulla Dhariwal. *Glow: Generative Flow with Invertible 1x1 Convolutions*. 2018. arXiv: [1807.03039 \[stat.ML\]](#).
- [8] Aditya Kusupati et al. *Matryoshka Representation Learning*. 2024. arXiv: [2205.13147 \[cs.LG\]](#).
- [9] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. *BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models*. 2023. arXiv: [2301.12597 \[cs.CV\]](#).
- [10] Ji Lin et al. *VILA: On Pre-training for Visual Language Models*. 2024. arXiv: [2312.07533 \[cs.CV\]](#).
- [11] Shuming Ma et al. *The Era of 1-bit LLMs: All Large Language Models are in 1.58 Bits*. 2024. arXiv: [2402.17764 \[cs.CL\]](#).
- [12] William Peebles and Saining Xie. *Scalable Diffusion Models with Transformers*. 2023. arXiv: [2212.09748 \[cs.CV\]](#).
- [13] Dustin Podell et al. *SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis*. 2023. arXiv: [2307.01952 \[cs.CV\]](#).
- [14] Rafael Rafailov et al. *Direct Preference Optimization: Your Language Model is Secretly a Reward Model*. 2023. arXiv: [2305.18290 \[cs.LG\]](#).

- [15] Yang Song et al. *Score-Based Generative Modeling through Stochastic Differential Equations*. 2021. arXiv: [2011.13456](https://arxiv.org/abs/2011.13456) [cs.LG].
- [16] Zineng Tang et al. *CoDi-2: In-Context, Interleaved, and Interactive Any-to-Any Generation*. 2023. arXiv: [2311.18775](https://arxiv.org/abs/2311.18775) [cs.CV].
- [17] Yilun Xu et al. *PFGM++: Unlocking the Potential of Physics-Inspired Generative Models*. 2023. arXiv: [2302.04265](https://arxiv.org/abs/2302.04265) [cs.LG].
- [18] Jianwei Yang et al. *Set-of-Mark Prompting Unleashes Extraordinary Visual Grounding in GPT-4V*. 2023. arXiv: [2310.11441](https://arxiv.org/abs/2310.11441) [cs.CV].
- [19] Shunyu Yao et al. *Tree of Thoughts: Deliberate Problem Solving with Large Language Models*. 2023. arXiv: [2305.10601](https://arxiv.org/abs/2305.10601) [cs.CL].
- [20] Lvmin Zhang and Maneesh Agrawala. *Transparent Image Layer Diffusion using Latent Transparency*. 2024. arXiv: [2402.17113](https://arxiv.org/abs/2402.17113) [cs.CV].
- [21] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. *Adding Conditional Control to Text-to-Image Diffusion Models*. 2023. arXiv: [2302.05543](https://arxiv.org/abs/2302.05543) [cs.CV].

3 Video Recording and Submission

Video Recording For the presentation, you will record yourself and then send us the link to the recording. Here are the steps you should take:

1. Open Zoom and start your "Personal Meeting". Be sure to turn on your microphone.
2. Click "Record" and then "Record to the Cloud".
3. Introduce yourself by name. Then go ahead and present for 5-10 minutes.
4. Extremely important: Now click "End the Meeting". Your recording will be uploaded to the cloud and processed.
5. Go to <https://cmu.zoom.us/recording> and wait (10 minutes) for your recording to finish processing.
6. Click the "Share..." button on the left of the recording, then click "Copy Sharing Information" and paste the clipboard text into Gradescope.

Submit to Gradescope You should submit your presentation video and slides to Gradescope.

7. Open the HW623 assignment in Gradescope.
8. Paste the video recording information that you copy/pasted above into the free-text field on the assignment.
9. Then upload your slides as a PDF on the file-upload question.
10. In Gradescope, save and submit this assignment.