April 17, 2024

List of research themes

ID	Research Area	研究テーマ
1	World Model	Research on model merging of world models
2	World Model	Basic research on meta-reinforcement learning and policy gradient methods
3	World Model	Construction of Large World Models
4	Algorithm	Backpropagation-free Algorithm
5	Algorithm	Research on distillation methods for diffusion models
6	Algorithm	Research on Transformer architecture improvement
7	Algorithm	Research on Strong Lottery Hypothesis
8	LLM	Research on social risks of Large Language Models（bias, halting, anti-learning, etc)
9	LLM	Understanding the principles of operation of Large Language Models
10	LLM	Research on model development with Large Language Models
11	Robotics	City-scale 3D Vision-Language Using Foundation Models
12	Robotics	Developing Foundation Models for Robotics
13	Robotics	Application of Foundation Models in Robotics and Their Transfer Learning
14	Robotics	Motion Planning and Modeling Environments using Fast, Differentiable Simulators and Renderers
15	Social implementation	Research on low-cost application methods in the era of foundation models
16	Social implementation	Research on solving real-world problems by utilizing LLM
17	Brain model	Brain-based AI modeling based on brain reference architecture

1: Research on model merging of world models

Research Area: World Model
Keywords: World Model, Foundation models, Model merging

Overview

In deep learning research, model merging has been studied to improve performance by merging models after learning. In this research, we aim to learn agents that can be applied to a wider range of environments by merging multiple world models that have learned the dynamics of different environments.

Required Skills

- Knowledge of deep learning (level of completion of “Deep Learning Fundamentals” course)
- Research experience with deep learning frameworks such as PyTorch

Welcome skills

- Knowledge and implementation experience with world models (level of completion of “World Models and Intelligence” course)
- Experience with reading and implementing replication of the latest research papers

2: Basic research on meta-reinforcement learning and policy gradient methods

Research Area: World Model
Keywords: algorithm, meta-learning, reinforcement learning

Overview

In reinforcement learning, the type of exploration method is crucial for performance. Existing exploration methods are mostly independent of the environment and are rule-based. However, if there is some structure in the environment, and if we can make good use of it, we should be able to explore more efficiently. In this research theme, we aim to meta-learn efficient exploration strategies that take advantage of the structure in the environment.

Required Skills

- Fundamental knowledge of reinforcement learning and deep learning
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Research experience with the JAX framework

3: Scaling up world models

Research Area: World model
Keywords: World model, video prediction, multimodal

Overview

This project will scale up video prediction models such as Dreamer V3 and Video GPT and improve data collection and architecture for this purpose. We plan to collect not only video data but also text-conditional data such as GAIA-1. We will examine whether the scale law is valid in the world model and, if so, what properties it has.

Required Skills

- Knowledge of deep learning (level of completion of "Deep Learning Fundamentals" course)
- Research experience with deep learning frameworks such as PyTorch
- Must be able to read and understand the papers in reference independently

Welcome Skill

- Experience with model-based reinforcement learning such as Dreamer and video prediction models such as VideoGPT
- Experience with multi-node parallel computing and libraries such as Deep Speed
- Submission/acceptance to top International Conference

■ Reference

[1] VideoGPT: Video Generation using-VAE and Transformers, https://arxiv.org/abs/2104.10157
[2] Mastering Diverse Domains through World Models, https://arxiv.org/abs/2301.04104
[3] Physics of Language Models: Part 3.1, Knowledge Storage and Extraction, https://arxiv.org/abs/2309.14316
[4] GAIA-1: A Generative World Model for Autonomous Driving, https://arxiv.org/abs/2309.170802

4: Backpropagation-free algorithm

Research Area: Algorithms
Keywords: Backpropagation-free, Energy-Based Model, Predictive Coding, handling time

Overview

Conventional deep learning algorithms rely on the backpropagation method, which requires the propagation of global error information throughout the entire network. This method is questioned in terms of its biological plausibility, and in practical applications, it is time-consuming to train large-scale models. In this study, we aim to develop a new learning algorithm that learns network parameters from local information.

Required Skills

- Knowledge of deep learning (level of completion of “Deep Learning Fundamentals” Course)
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Experience with reading and replicating the latest research papers
- Knowledge and implementation experience with energy-based models and predictive coding

5: Research on distillation methods for diffusion models

Research Area: Algorithm
Keywords: Algorithm, Diffusion Model, Distillation

Overview

Diffusion models have been successful, especially in areas such as image generation, but they are time-consuming to generate samples. In this research, we develop a method to distill a learned diffusion model with another fast-generating model. In particular, we focus on methods that do not require training data during distillation.

Required Skills

- Knowledge of deep learning (level of completion of “Deep Learning Fundamentals” course)
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Experience in implementing diffusion models
- Knowledge of the mathematical background of diffusion models

■ Reference

[1] Song, Yang, et al. "Score-Based Generative Modeling through Stochastic Differential Equations." International Conference on Learning Representations. 2020.
[2] Song, Yang, et al. "Consistency Models." International Conference on Machine Learning. PMLR, 2023.

6: Research on transformer architecture improvement

Research Area: Algorithm
Keywords: Algorithm, Transformer

Overview

The Transformer is becoming the de facto standard architecture in a wide range of deep learning domains, including large language models. However, it is known to suffer from the following problems: (1) its computational cost grows quadratically in terms of the sequence length, and (2) its training tends to be unstable. In the research, we aim to propose a new architecture that solves these problems.

Required Skills

- Knowledge of deep learning (level of completion of the “Deep Learning Fundamentals” course)
- Research experience with deep learning frameworks such as PyTorch

Welcome skills

- Experience implementing Transformer from scratch (you may use a deep learning framework such as PyTorch)
- Knowledge of recent NN architectures

■ Reference

[1] Tay, Yi, et al. "Efficient transformers: A survey." ACM Computing Surveys 55.6 (2022): 1-28.

[2] He, Bobby, and Thomas Hofmann. "Simplifying Transformer Blocks." The Twelfth International Conference on Learning Representations. 2023.

7: Research on Strong Lottery Ticket Hypothesis

Research Area: Algorithm
Keywords: Algorithm, Lottery hypothesis, Grokking, Pruning

Overview

The Strong Lottery Ticket Hypothesis states that “among randomly initialized networks, there exists a network that achieves good performance without learning at all.” Although it is stated as a hypothesis, there are networks on WideResNet that achieve performance equivalent to ResNet34, and theoretical proof of existence using SubsetSum has also been provided. This phenomenon is interesting in the sense that it suggests the importance of the “structure optimization” perspective in addition to the conventional interpretation of deep learning as “weight optimization”.

In this project, the following research will be conducted based on this phenomenon.

- Developing better lottery detection algorithms
- Investigations of the relationship between different phenomena known as deep learning behavior such as Grokking and Double Descent

Required Skills

- Knowledge of deep learning (level of completion of “Deep Learning Fundamentals” Course)
- Research experience with deep learning frameworks such as PyTorch
- Interest in the phenomenon of the strong lottery hypothesis

Welcome Skills

- Experience in related areas such as the lottery hypothesis, the strong lottery hypothesis, Pruning, NAS, etc.
- Experience implementing ideas based on multiple papers

- Experience with submissions/adoption to top machine learning conferences

■ Reference

[1] What's Hidden in a Randomly Weighted Neural Network?, CVPR2020, https://arxiv.org/abs/1911.13299

[2] A Study on the Ramanujan Graph Property of Winning Lottery Tickets, ICML2022, https://proceedings.mlr.press/v162/pal22a.html

[3] Grokking: Generalization Beyond Overfitting on Small Algorithmic Datasets, https://arxiv.org/abs/2201.02177

8: Research on social risks of Large Language Models (bias, halting, anti-learning, etc.)

Research Area: LLM
Keywords: Large Language Models, bias, hallucination, Unlearning

Overview

As the social implementation of large language models advances, it is becoming increasingly important to study the social risks (e.g., bias, and hallucination) inherent in them. This project aims to observe and reduce these risks.

Required Skills

- Knowledge of deep learning (level of completion "Deep Learning Fundamentals" course)
- Research experience with deep learning frameworks such as PyTorch
- Knowledge of LLM (level of completion of "Large Language Models course") https://weblab.t.u-tokyo.ac.jp/llm_contents/

Welcome Skills

- Experience with multi-GPU or multi-node language model training using distributed learning frameworks (e.g., deep speed)
- Research achievements including main papers in the field of natural language processing at international or domestic conferences (e.g. **ACL, EMNLP, JSAI, NLP)

■ Reference

[1] Survey of Hallucination in Natural Language Generation, https://arxiv.org/abs/2202.03629
[2] Bias and Fairness in Large Language Models: A Survey, https://arxiv.org/abs/2309.00770
[3] Machine Unlearning for Traditional Models and Large Language Models: A Short Survey, https://arxiv.org/pdf/2404.01206.pdf

9: Understanding the principles of operation of Large Language Models

Research Area: LLM
Keywords: Large Language Models, in-context learning, Probing, Chain-of-Thought

Overview

This project focuses on understanding the principles of operation in order to pursue the reasons for the dramatic performance development of Large Language Models. Specifically, we will theoretically and experimentally elucidate the principles of In-text Learning and Prompting, develop Prompting strategies such as Chain of Thought, utilize and construct methods for analyzing the internal state of the Transformer by probing and intervention, and analyze the relationship between learning data and the ability o express the ability.

Required Skills

- Knowledge of deep learning (level of completion of “Deep Learning Fundamentals” course)
- Research experience with deep learning frameworks such as PyTorch
- Knowledge of LLM (level of completion of “Large Language Models course”) https://weblab.t.u-tokyo.ac.jp/llm_contents/

Welcome Skills

- Experience with multi-GPU and multi-node language model training using distributed learning frameworks (e.g., deepspeed)
- Research achievements including main papers in the field of natural language processing at international or domestic conferences (e.g. **ACL, EMNLP, JSAI, NLP)

■ Reference

[1] A Survey of Large Language Models, https://arxiv.org/abs/2303.18223
[2] Chain-of-Thought Prompting Elicits Reasoning in Large Language

Models, https://arxiv.org/abs/2201.11903
[3] Do Llamas Work in English? On the Latent Language of Multilingual Transformers, https://arxiv.org/abs/2402.10588
[4] Why Can GPT Learn In-Context? Language Models Implicitly Perform Gradient Descent as Meta-Optimizers, https://arxiv.org/abs/2212.10559
[5] Attention is Not Only a Weight: Analyzing Transformers with Vector Norms, https://arxiv.org/abs/2004.10102

10: Research on model development with Large Language Models

Research Area: LLM
Keywords: Large Language models

Overview

Advances in the language understanding and language generation capabilities of Large Language Models are expected to open up new possibilities for automated model development methods.
Existing methods, such as NAS, are built on non-verbal algorithms, but this project proposes an automated development method that leverages natural language capabilities to explore new possibilities. For example, one of the research directions is to automatically generate development code from natural language.

Required Skills

- Knowledge of deep learning (level of completion of “Deep Learning Fundamentals” course)
- Research experience with deep learning frameworks such as PyTorch
- Knowledge of LLM (level of completion of “Large Language Models course”) https://weblab.t.u-tokyo.ac.jp/llm_contents/

Welcome Skills

- Experience with multi-GPU and multi-node language model training using distributed learning frameworks (e.g., deepspeed)
- Research achievements including main papers in the field of natural language processing at international or domestic conferences (e.g. **ACL, EMNLP, JSAI, NLP)

■ Reference

[1] Code Llama: Open Foundation Models for Code,

https://arxiv.org/abs/2308.12950
[2] Neural Architecture Search with Reinforcement Learning,
https://arxiv.org/abs/1611.01578

11: City-scale 3D vision-Language using foundation models

Research Area: Robotics
Keywords: Foundation models, 3D Vision-Language

Overview

Developments in urban 3D scanning technologies have enabled the creation of accurate and photorealistic large-scale 3D scene datasets. While city-level photorealistic 3D scans are now practical in various fields, such as autonomous navigation and disaster risk assessment, understanding city scenes through human-posed language remains challenging. This project aims to create city-scale 3D datasets by automatically adding detailed semantic information to 3D maps using vision-language foundation models and to develop methods for accessing geospatial information from such 3D data. Connecting language with 3D maps enables the realization of a dialogue system that can freely access 3D geospatial information and aerial navigation systems that can automatically identify destinations and routes in the 3D map.

Required Skills

- Knowledge on deep learning (level of completion of “Deep Learning Fundamentals” course)
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Research and development experience in the field of computer vision
- Experience in 3D data processing
- Experience in creating research datasets

■ Reference

[1] ScanQA https://arxiv.org/pdf/2112.10482.pdf
[2] CityRefer https://arxiv.org/abs/2310.18773

12: Developing foundation models for robotics

Research Area: Robotics
Keywords: Robotics, Foundation Models

Overview

Parallel with the development of LLM (Large Language Model) and VLM (Vision Language Model), research has been conducted in the field of robot control, such as Google's RT-X, collecting large-scale datasets and learning large models. This theme will explore model training and data collection methods for developing such robot foundation models.

Specifically, we focus on the following topics:

- Validation of scaling laws of the learned policy of the robot
- Validation of in-context learning in robot control
- Transfer learning of LLM/VLM to robot control data
- Data collection and model validation on large-scale real robots in collaboration with multiple laboratories
- Research on efficient data collection methods

Required Skills

- Fundamental knowledge of reinforcement learning and deep learning
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Research experience using the JAX framework (not necessary, but better to have)
- Experience in controlling real robots, especially in developing robot systems using ROS
- Experience in submitting to international conferences on robotics and robot learning, such as ICRA/RSS/CoRL/IROS

■ Reference

[1] RT-X (ICRA2024), https://arxiv.org/abs/2310.08864

[2] Octo, https://octo-models.github.io/

13: Application of foundation models in robotics and their transfer learning

Research Area: robotics
Keywords: Robotics, Foundation Models

Overview

What innovations should be made when incorporating recent foundation models (LLMs and VLMs) into real robot systems such as service robots? Also, what innovations are needed when adapting recent models of large-scale robot control policy, such as RT-X, which outputs end-to-end control, to new environments and settings?
This research theme will identify issues in effectively utilizing these foundation models in robotics systems and verify the solutions. Specifically, the following themes are focused on:

- Design of service robot system utilizing multiple foundation models
- Transfer learning and finetuning methods for new tasks, environments, and modalities (force, tactile, point clouds, etc.) of large-scale robot foundation models such as RT-X
- Verification of model merging for robot foundation models

Required Skills

- Fundamental knowledge of reinforcement learning and deep learning
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Experience in controlling real robots, especially in developing robot systems using ROS
- Experience in submitting to international conferences on robotics and robot learning, such as ICRA/RSS/CoRL/IROS

■ Reference

[1] RT-X (ICRA2024), https://arxiv.org/abs/2310.08864
[2] Octo, https://octo-models.github.io/
[3] Self-recovery Prompting (ICRA2024), https://arxiv.org/abs/2309.14425

14: Motion planning and modeling environments using fast, differentiable simulators and renderers

Research Area: Robotics
Keywords: Robotics, World Model, Differentiable Physics, Isaac Sim

Overview

Fast and differentiable physics simulators and renderers accelerated with GPUs, such as Isaac Sim and Omniverse, have become available in recent years. In this theme, we will conduct research on methods and frameworks for efficient robot motion planning, environment recognition, and modeling by utilizing these simulators and renderers.

Required Skills

- Fundamental knowledge of reinforcement learning and deep learning
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Experience implementing GPU-accelerated differentiable simulators such as IsaacSim / Omniverse, BRAX, Kaolin, wisp, etc.
- Experience in controlling real robots, especially in developing robot systems using ROS
- Experience in submitting to international conferences on robotics and robot learning, such as ICRA/RSS/CoRL/IROS
- Experience in any of the following:
  - Experience with TAMP frameworks (PDDL, PDDLStream, etc.)
  - Experience with motion planning libraries such as Moveit/CuRobo

- Sim2Real experience with reinforcement learning policies
- Experience in experimenting with 3D visual SLAM methods
- Experience in implementing neural renderers, such as NeRFs and Gaussian Splatting

■ Reference

[1] IsaacSim, https://docs.omniverse.nvidia.com/index.html
[2] GenDom (ICRA2024), https://arxiv.org/abs/2309.09051

15: Research on low-cost application methods in the era of foundation models

Research Area: Applied Machine/Deep Learning, Social Implementation
Keywords: Applied Machine/Deep Learning, Social Implementation, MLOps, Foundation Model, Large Language Models

Overview

In the new era of deep learning models, such as the foundation models / LLMs, high-performing models have been proposed in various application domains, they require huge computational resources and data sets. However, those resources cannot be accessed/obtained in the wild for assorted reasons.
In this research theme, we will study to develop applied methodologies that work in those restricted environments.

Required Skills

- Knowledge of deep learning (level of completion of “Deep Learning Basic” course)
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Experience with reading the latest research papers and reproducing its experiment results
- Knowledge and implementation experience in transfer learning

16: Research on tackling real-world / Social issues by utilizing LLM

Research Area: Applied Machine / Deep Learning, Social Implementation
Keywords: Social Implementation, Social Issues, Large Language Models

Overview

LLM technology has made remarkable progress, and it demonstrates effectiveness not only in traditional natural language processing tasks but also in tasks related to decision-making and data analysis. In this research theme, we will explore the usage of LLMs not limited to NLP techniques in various social issues.

Required Skills

- Knowledge of deep learning (level of completion of “Deep Learning Basics” course)
- Research experience with deep learning frameworks such as PyTorch

Welcome Skills

- Experience with reading the latest research papers and reproducing its experiment results
- Knowledge and implementation experience in transfer learning
- Experience in coding with LLM

17: Brain-based AI modeling based on the brain reference architecture

Research Area: Brain Model
Keywords: Brain, World Mode, Large Language Models

Overview

We will build computational models of various brain regions in accordance with the mesoscopic-level anatomy of the brain. This is based on a development methodology called BRA-driven development using brain reference architecture (BRA). Examples of specific research include the parallel creation of hypothetical component diagrams connecting computational components in the brain and functional realization graphs, which are hierarchical descriptions of the computational functions of the brain, based on the BRA data format. In this process, we will streamline the design process while utilizing AI techniques such as Large Language Models.

Required Skills

- The candidate should have a deep understanding of articles in neuroscience and be interested in discussing computational models based on anatomical structure based on this knowledge.
- The candidate should be able to discuss modeling complex biological systems.

Welcome Skills

- Being able to have an in-depth understanding of neuroscience publications
- Experience in modeling complex biological systems