Novelty Assessment Report

Paper: Binomial Gradient-Based Meta-Learning for Enhanced Meta-Gradient Estimation
PDF URL: https://openreview.net/pdf?id=mKgUAO41zf
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-29

Abstract

Meta-learning offers a principled framework leveraging task-invariant priors from related tasks, with which task-specific models can be fine-tuned on downstream tasks, even with limited data records. Gradient-based meta-learning (GBML) relies on gradient descent (GD) to adapt the prior to a new task. Albeit effective, these methods incur high computational overhead that scales linearly with the number of GD steps. To enhance efficiency and scalability, existing methods approximate the gradient of prior parameters (meta-gradient) via truncated backpropagation, yet suffer large approximation errors. Targeting accurate approximation, this work puts forth binomial GBML (BinomGBML), which relies on a truncated binomial expansion for meta-gradient estimation. This novel expansion endows more information in the meta-gradient estimation via efficient parallel computation. As a running paradigm applied to model-agnostic meta-learning (MAML), the resultant BinomMAML provably enjoys error bounds that not only improve upon existing approaches, but also decay super-exponentially under mild conditions. Numerical tests corroborate the theoretical analysis and showcase boosted performance with slightly increased computational overhead.

Core Task Landscape

This paper addresses: **Meta-Gradient Estimation in Gradient-Based Meta-Learning**

A total of **50 papers** were analyzed and organized into a taxonomy with **21 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Meta-Gradient Estimation Theory and Bias-Variance Analysis**
- **Efficient Meta-Gradient Computation Methods**
- **Variance Reduction and Debiasing Techniques**
- **Meta-Gradient Reinforcement Learning Applications**
- **Supervised and Few-Shot Learning Applications**
- **Specialized Meta-Gradient Techniques and Extensions**
- **Historical Perspectives and Surveys**

Complete Taxonomy Tree

- Meta-Gradient Estimation in Gradient-Based Meta-Learning Survey Taxonomy
- Meta-Gradient Estimation Theory and Bias-Variance Analysis
  - Bias and Variance Characterization (4 papers)
  - [5] A theoretical understanding of gradient bias in meta-reinforcement learning (Feng, 2022) View paper
  - [8] Settling the Bias and Variance of Meta-Gradient Estimation for Meta-Reinforcement Learning (Bo Liu (Benjamin Liu), 2021) View paper
  - [12] No DICE: An investigation of the bias-variance tradeoff in meta-gradients (R Vuorio, 2021) View paper
  - [42] An Investigation of the Bias-Variance Tradeoff in Meta-Gradients (Vuorio, 2022) View paper
  - Convergence Analysis and Iteration Complexity (3 papers)
  - [13] On the Iteration Complexity of Hypergradient Computation (Grazzi, 2020) View paper
  - [39] A New First-Order Meta-Learning Algorithm with Convergence Guarantees (Chayti, 2024) View paper
  - [48] Theoretical Convergence of Multi-Step Model-Agnostic Meta-Learning (Ji, 2020) View paper
- Efficient Meta-Gradient Computation Methods
  - Implicit Gradient and Hypergradient Approximation (3 papers)
  - [2] Meta-learning with implicit gradients (Rajeswaran, 2019) View paper
  - [9] Efficient Curvature-Aware Hypergradient Approximation for Bilevel Optimization (Yang Junfeng, 2025) View paper
  - [18] Scalable bayesian meta-learning through generalized implicit gradients (Yilang Zhang, 2023) View paper
  - Structural Exploitation and Mixed-Mode Differentiation (2 papers)
  - [28] Scalable Meta-Learning via Mixed-Mode Differentiation (Kemaev, 2025) View paper
  - [44] Continuous-Time Meta-Learning with Forward Mode Differentiation (Deleu, 2022) View paper
  - Long-Horizon and Multi-Step Meta-Gradient Methods ★ (3 papers)
  - [0] Binomial Gradient-Based Meta-Learning for Enhanced Meta-Gradient Estimation (Anon et al., 2026) View paper
  - [10] One step at a time: Pros and cons of multi-step meta-gradient reinforcement learning (Bonnet, 2021) View paper
  - [20] A Lazy Approach to Long-Horizon Gradient-Based Meta-Learning (Muhammad Abdullah Jamal, 2021) View paper
- Variance Reduction and Debiasing Techniques
  - Stochastic Meta-Gradient Variance Reduction (2 papers)
  - [1] A stochastic approach to Bi-Level optimization for hyperparameter optimization and meta learning (Kimï¼ Min-Young, 2025) View paper

- [25] Efficient Variance Reduction for Meta-Learning (Yang, 2022) View paper
- Meta-Gradient Debiasing Approaches (2 papers)
- [6] Meta-Reinforcement Learning with Evolving Gradient Regularization (Jiaxing Chen, 2025) View paper
- [14] Debiasing Meta-Gradient Reinforcement Learning by Learning the Outer Value Function (Bonnet, 2022) View paper

- Meta-Gradient Reinforcement Learning Applications
  - Hyperparameter and Return Adaptation (3 papers)
  - [17] Meta-Gradient Reinforcement Learning (Zhongwen Xu, 2022) View paper
  - [37] Meta-Gradient Search Control: A Method for Improving the Efficiency of Dyna-style Planning (Martin, 2024) View paper
  - [46] Meta-Gradient Reinforcement Learning with an Objective Discovered Online (Xu, 2022) View paper
  - Policy-Level Meta-Learning and Credit Assignment (3 papers)
  - [11] Promp: Proximal meta-policy search (Jonas Rothfuss, 2018) View paper
  - [15] Learning locomotion skills via model-based proximal meta-reinforcement learning (Qing Xiao, 2019) View paper
  - [35] Offline Meta-Reinforcement Learning with Evolving Gradient Agreement (Jiaxing Chen, 2024) View paper
  - Multi-Agent and Incentive Design (2 papers)
  - [21] Adaptive Incentive Design with Multi-Agent Meta-Gradient Reinforcement Learning (Yang, 2021) View paper
  - [40] A Meta-Gradient Approach to Learning Cooperative Multi-Agent Communication Topology (Zhang Qi, 2021) View paper
- Supervised and Few-Shot Learning Applications
  - MAML Variants and Enhancements (3 papers)
  - [7] LaANIL: ANIL with Look-Ahead Meta-Optimization and Data Parallelism (Vasu Tammisetti, 2024) View paper
  - [19] Improving Generalization in Meta-Learning via Meta-Gradient Augmentation (Wang Ren, 2025) View paper
  - [36] Meta Gradient Boosting Neural Networks (Manqing Dong, 2021) View paper
  - Alternative Gradient-Based Meta-Learning Frameworks (3 papers)
  - [3] A contrastive rule for meta-learning (Zucchet, 2022) View paper
  - [4] Meta-learning with warped gradient descent (Sebastian Flennerhag, 2019) View paper
  - [49] VIABLE: Fast Adaptation via Backpropagating Learned Loss (Feng, 2019) View paper
  - Prompt Tuning and Language Model Meta-Learning (1 papers)
  - [34] Self-supervised Meta-Prompt Learning with Meta-Gradient Regularization for Few-shot Generalization (Kaihang Pan, 2023) View paper
  - Graph and Vision Applications (4 papers)
  - [27] Towards Adaptive Meta-Gradient Adversarial Examples for Visual Tracking (Wei-Long Tian, 2025) View paper
  - [30] Edge Sparsification for Graphs via Meta-Learning (Guihong Wan, 2021) View paper
  - [41] Personalized Image Aesthetics Assessment via Meta-Learning With Bilevel Gradient Optimization. (Hancheng Zhu, 2022) View paper
  - [47] Graph Meta Learning via Local Subgraphs (Kexin Huang, 2020) View paper
- Specialized Meta-Gradient Techniques and Extensions
  - Optimism and Bootstrapping in Meta-Gradients (1 papers)
  - [33] Optimistic Meta-Gradients (Flennerhag, 2023) View paper
  - Meta-Gradient Regularization and Robustness (2 papers)
  - [29] Robust meta gradient learning for high-dimensional data with noisy-label ignorance (Ben Liu, 2023) View paper
  - [43] Semi-Supervised Learning with Meta-Gradient (Zhang, 2021) View paper
  - Continual and Lifelong Meta-Learning (2 papers)
  - [32] Reevaluating Meta-Learning Optimization Algorithms Through Contextual Self-Modulation (Barton, 2024) View paper
  - [50] Continual Meta Learning (Sharifnassab, n.d.) View paper
  - Adversarial Meta-Gradient Applications (2 papers)
  - [26] Topological Adversarial Attacks on Graph Neural Networks Via Projected Meta Learning (Mohammed Aburidi, 2024) View paper
  - [38] Meta Gradient Adversarial Attack (Zheng Yuan, 2021) View paper
  - Biologically-Inspired and Neuromorphic Meta-Learning (1 papers)
  - [16] Meta-learning spiking neural networks with surrogate gradient descent (Kenneth Stewart, 2022) View paper
  - Domain-Specific and Emerging Applications (2 papers)
  - [22] Collaborative Optimisation of Greyâ□□Markov Modelling and Adversarial Metaâ□□Learning for Resilient Intelligent Transportation Systems (J Jia, 2025) View paper
  - [31] Theory and Application of Meta Learning Techniques (Yang, 2024) View paper
- Historical Perspectives and Surveys (3 papers)
  - [23] A History of Meta-gradient: Gradient Methods for Meta-learning (Sutton, 2022) View paper
  - [24] Meta-learning & Compositional Generalization in Neural Networks (Schug, 2025) View paper
  - [45] Stateless neural meta-learning using second-order gradients (M. Huisman, 2022) View paper

Narrative

Core task: meta-gradient estimation in gradient-based meta-learning. The field addresses how to compute gradients of outer-level objectives with respect to inner-level parameters or hyperparameters, a central challenge in bi-level optimization for meta-learning. The taxonomy reveals several major branches: theoretical investigations into bias-variance trade-offs and convergence properties (e.g., Gradient Bias Theory[5], Bias Variance Meta-Gradient[8]); efficient computation methods that reduce the cost of unrolling long inner optimization trajectories (e.g., Implicit Gradients[2], Multi-Step Meta-Gradient[10]); variance reduction and debiasing techniques to stabilize noisy gradient estimates (e.g., Efficient Variance Reduction[25], Debiasing Outer Value[14]); and diverse application domains spanning reinforcement learning (Meta-Gradient RL[17], Proximal Meta-Policy[11]) and supervised or few-shot learning settings. Additional branches cover specialized extensions—such as meta-gradient methods for adversarial robustness, continual learning, and neural architecture search—as well as historical surveys that trace the evolution of these ideas (History of Meta-Gradient[23]).

A particularly active line of work focuses on long-horizon and multi-step meta-gradient methods, which seek to balance computational efficiency with the fidelity of gradient estimates over extended inner-loop trajectories. Binomial Gradient Meta-Learning[0] sits within this branch, proposing a novel binomial sampling strategy to approximate multi-step meta-gradients more efficiently. It shares thematic concerns with Multi-Step Meta-Gradient[10], which also tackles the challenge of propagating gradients through many inner updates, and with Lazy Long-Horizon[20], which explores lazy evaluation techniques to defer expensive computations. Compared to these neighbors, Binomial Gradient Meta-Learning[0] emphasizes a probabilistic sampling perspective that aims to reduce variance while maintaining low bias, contrasting with the deterministic unrolling strategies or lazy caching approaches seen in related works. Across the field, open questions remain about the optimal trade-offs between computational cost, gradient bias, and variance, especially as meta-learning scales to more complex tasks and longer adaptation horizons.

Related Works in Same Category

The following **2 sibling papers** share the same taxonomy leaf node with the original paper:

1. One step at a time: Pros and cons of multi-step meta-gradient reinforcement learning

Authors: Bonnet, ClÃ©ment, Caron, Paul, ClÃ©ment Bonnet, et al. (13 authors total) | **Year/Venue**: 2021 | **URL**: View paper

Abstract

Self-tuning algorithms that adapt the learning process online encourage more effective and robust learning. Among all the methods available, meta-gradients have emerged as a promising approach. They leverage the differentiability of the learning rule with respect to some hyper-parameters to adapt them in an online fashion. Although meta-gradients can be accumulated over multiple learning steps to avoid myopic updates, this is rarely used in practice. In this work, we demonstrate that whilst mult...

Relationship Analysis

Both papers belong to the Long-Horizon and Multi-Step Meta-Gradient Methods category, addressing the challenge of computing meta-gradients over extended inner-loop optimization horizons. The original paper (BinomGBML) focuses on reducing approximation error in meta-gradient estimation through a truncated binomial expansion that enables parallel computation, while the candidate paper analyzes the bias-variance tradeoff inherent in multi-step meta-gradients and proposes a mixing strategy to balance these competing factors. The key difference is that BinomGBML emphasizes computational efficiency and error reduction through mathematical reformulation, whereas the candidate paper investigates the fundamental statistical properties (bias and variance) of multi-step meta-gradients in reinforcement learning contexts.

---

2. A Lazy Approach to Long-Horizon Gradient-Based Meta-Learning

Authors: Muhammad Abdullah Jamal, Liqiang Wang, Boqing Gong | **Year/Venue**: 2021 | **URL**: View paper

Abstract

Gradient-based meta-learning first trains task-specific models by an inner loop and then backpropagates meta-gradients through the loop to update the meta-model. To avoid high-order gradients, existing methods either take a small number of inner steps or approximate the meta-updates for the situations that the meta-model and task models lie in the same space. To enable long inner horizons for more general meta-learning problems, we instead propose an intuitive teacher-student strategy. The key i...

Relationship Analysis

Both papers belong to the Long-Horizon and Multi-Step Meta-Gradient Methods category, addressing the computational challenges of meta-gradient estimation over extended inner-loop optimization horizons. The original paper (BinomGBML) proposes a truncated binomial expansion approach to improve meta-gradient estimation accuracy through parallel HVP computations, while the candidate paper introduces a teacher-student strategy where a student network explores the task-specific model space and a teacher performs a "leap" to define a lightweight computational graph. The key difference is that BinomGBML focuses on mathematical reformulation of the backpropagation process for efficiency, whereas the candidate employs a distinct architectural approach with separate student-teacher networks to avoid high-order gradients entirely.

Contributions Analysis

Overall novelty summary. The paper proposes BinomGBML, a method using truncated binomial expansion to estimate meta-gradients in gradient-based meta-learning, specifically applied to MAML. It resides in the 'Long-Horizon and Multi-Step Meta-Gradient Methods' leaf, which contains only three papers total including this one. This is a relatively sparse research direction within the broader taxonomy of 50 papers across 21 leaf nodes, suggesting the specific problem of multi-step meta-gradient approximation via binomial expansions has received limited prior attention compared to other meta-gradient estimation strategies.

The taxonomy reveals that BinomGBML's leaf sits within 'Efficient Meta-Gradient Computation Methods', alongside sibling branches addressing implicit differentiation (3 papers), structural exploitation (2 papers), and the current long-horizon methods. Neighboring branches include variance reduction techniques (4 papers) and theoretical bias-variance analysis (7 papers). The scope notes clarify that this leaf focuses on extended inner-loop horizons, excluding single-step approximations handled by implicit gradient methods. The paper's binomial expansion approach appears to bridge computational efficiency concerns with the theoretical error analysis typical of the bias-variance branch, positioning it at an intersection of algorithmic and theoretical contributions.

Among eight candidates examined across three contributions, none were found to clearly refute the proposed work. The binomial expansion method itself was assessed against one candidate with no refutation. Theoretical error bounds for BinomMAML examined three candidates, finding none that provide overlapping guarantees. The dynamic computational graph management contribution reviewed four candidates without identifying prior work offering the same memory-efficient implementation. This limited search scope—eight papers from semantic retrieval—suggests the analysis captures closely related work but cannot claim exhaustive coverage of all potential prior art in meta-gradient estimation or MAML variants.

Given the sparse population of the target leaf and the absence of refutations among examined candidates, the work appears to occupy a relatively unexplored niche within meta-gradient estimation. However, the small search scale (eight candidates) and the broader taxonomy context (50 papers total) indicate that while no direct overlap was detected, the novelty assessment remains contingent on this limited retrieval scope rather than a comprehensive field survey.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Binomial gradient-based meta-learning (BinomGBML) method

Description: The authors propose BinomGBML, a novel meta-gradient estimation method that uses truncated binomial expansion to incorporate more information than existing approaches while enabling efficient parallel computation. This method reformulates the meta-gradient as a cascade of vector operators that can be computed in parallel.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Learning to Coordinate: Distributed Meta-Trajectory Optimization Via Differentiable ADMM-DDP

URL: View paper

Brief Assessment

Distributed Meta-Trajectory[58] focuses on distributed trajectory optimization for multi-agent systems using ADMM-DDP, not meta-gradient estimation via truncated binomial expansion for general meta-learning tasks.

Contribution 2: Theoretical error bounds for BinomMAML

Description: The authors establish theoretical error bounds for BinomMAML under three different assumptions (Lipschitz gradient, convexity, and local strong convexity). They prove that BinomMAML achieves smaller estimation errors than existing methods, with super-exponential decay rates under certain conditions.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Fast-Rate PAC-Bayesian Generalization Bounds for Meta-Learning

URL: View paper

Brief Assessment

Fast-Rate PAC-Bayesian[52] focuses on PAC-Bayesian generalization bounds for meta-learning, which is a different theoretical framework than the gradient estimation error bounds analyzed in the original paper. The candidate addresses statistical generalization, not computational approximation errors in meta-gradient estimation.

2. Meta-D2AG: Causal Graph Learning with Interventional Dynamic Data

URL: View paper

Brief Assessment

Meta-D2AG[53] focuses on causal graph learning for dynamic time series data using meta-learning, not on meta-gradient estimation error bounds for gradient-based meta-learning algorithms like BinomMAML. The theoretical contributions are entirely different domains.

3. Nonlinear meta-learning can guarantee faster rates

URL: View paper

Brief Assessment

Nonlinear Faster Rates[51] focuses on meta-learning with nonlinear representations in RKHS for regression tasks, establishing convergence rates that scale with the number of tasks. BinomMAML addresses gradient-based meta-learning with binomial expansion for meta-gradient estimation in MAML, targeting computational efficiency and error bounds in a different technical framework.

Contribution 3: Dynamic computational graph management for memory efficiency

Description: The authors show that BinomMAML creates and releases computational graphs dynamically during execution, which significantly reduces memory consumption compared to vanilla MAML that stores all computation graphs. This addresses a key scalability limitation of MAML.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Learning to Propagate for Graph Meta-Learning

URL: View paper

Brief Assessment

Learning to Propagate[57] focuses on graph-based meta-learning for few-shot classification using message passing between class prototypes on a knowledge graph. It does not address computational graph management or memory efficiency in gradient-based meta-learning algorithms like MAML.

2. Research on traffic flow forecasting based on interactive dynamic meta-graph learning

URL: View paper

Brief Assessment

Interactive Dynamic Meta-Graph[54] focuses on traffic flow forecasting using spatial-temporal transformers and graph convolution for capturing dynamic spatial correlations. It does not address memory efficiency through dynamic computational graph management in meta-learning contexts.

3. Self-Distillation with Meta Learning for Knowledge Graph Completion

URL: View paper

Brief Assessment

Self-Distillation Knowledge Graph[55] focuses on knowledge graph completion using pruning and meta-learning for model compression, not on computational graph management for memory efficiency in meta-learning algorithms like MAML.

4. PyTorch Foundations and Applications: Definitive Reference for Developers and Engineers

URL: View paper

Brief Assessment

PyTorch Foundations[56] discusses PyTorch's dynamic computation graph capabilities as a general framework feature, not specifically addressing memory efficiency in meta-learning contexts like BinomMAML does.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Binomial Gradient-Based Meta-Learning for Enhanced Meta-Gradient Estimation View paper
- [1] A stochastic approach to Bi-Level optimization for hyperparameter optimization and meta learning View paper
- [2] Meta-learning with implicit gradients View paper
- [3] A contrastive rule for meta-learning View paper
- [4] Meta-learning with warped gradient descent View paper
- [5] A theoretical understanding of gradient bias in meta-reinforcement learning View paper
- [6] Meta-Reinforcement Learning with Evolving Gradient Regularization View paper
- [7] LaANIL: ANIL with Look-Ahead Meta-Optimization and Data Parallelism View paper
- [8] Settling the Bias and Variance of Meta-Gradient Estimation for Meta-Reinforcement Learning View paper
- [9] Efficient Curvature-Aware Hypergradient Approximation for Bilevel Optimization View paper
- [10] One step at a time: Pros and cons of multi-step meta-gradient reinforcement learning View paper

- [11] Promp: Proximal meta-policy search View paper
- [12] No DICE: An investigation of the bias-variance tradeoff in meta-gradients View paper
- [13] On the Iteration Complexity of Hypergradient Computation View paper
- [14] Debiasing Meta-Gradient Reinforcement Learning by Learning the Outer Value Function View paper
- [15] Learning locomotion skills via model-based proximal meta-reinforcement learning View paper
- [16] Meta-learning spiking neural networks with surrogate gradient descent View paper
- [17] Meta-Gradient Reinforcement Learning View paper
- [18] Scalable bayesian meta-learning through generalized implicit gradients View paper
- [19] Improving Generalization in Meta-Learning via Meta-Gradient Augmentation View paper
- [20] A Lazy Approach to Long-Horizon Gradient-Based Meta-Learning View paper
- [21] Adaptive Incentive Design with Multi-Agent Meta-Gradient Reinforcement Learning View paper
- [22] Collaborative Optimisation of Greyâ□□Markov Modelling and Adversarial Metaâ□□Learning for Resilient Intelligent Transportation Systems View paper
- [23] A History of Meta-gradient: Gradient Methods for Meta-learning View paper
- [24] Meta-learning & Compositional Generalization in Neural Networks View paper
- [25] Efficient Variance Reduction for Meta-Learning View paper
- [26] Topological Adversarial Attacks on Graph Neural Networks Via Projected Meta Learning View paper
- [27] Towards Adaptive Meta-Gradient Adversarial Examples for Visual Tracking View paper
- [28] Scalable Meta-Learning via Mixed-Mode Differentiation View paper
- [29] Robust meta gradient learning for high-dimensional data with noisy-label ignorance View paper
- [30] Edge Sparsification for Graphs via Meta-Learning View paper
- [31] Theory and Application of Meta Learning Techniques View paper
- [32] Reevaluating Meta-Learning Optimization Algorithms Through Contextual Self-Modulation View paper
- [33] Optimistic Meta-Gradients View paper
- [34] Self-supervised Meta-Prompt Learning with Meta-Gradient Regularization for Few-shot Generalization View paper
- [35] Offline Meta-Reinforcement Learning with Evolving Gradient Agreement View paper
- [36] Meta Gradient Boosting Neural Networks View paper
- [37] Meta-Gradient Search Control: A Method for Improving the Efficiency of Dyna-style Planning View paper
- [38] Meta Gradient Adversarial Attack View paper
- [39] A New First-Order Meta-Learning Algorithm with Convergence Guarantees View paper
- [40] A Meta-Gradient Approach to Learning Cooperative Multi-Agent Communication Topology View paper
- [41] Personalized Image Aesthetics Assessment via Meta-Learning With Bilevel Gradient Optimization. View paper
- [42] An Investigation of the Bias-Variance Tradeoff in Meta-Gradients View paper
- [43] Semi-Supervised Learning with Meta-Gradient View paper
- [44] Continuous-Time Meta-Learning with Forward Mode Differentiation View paper
- [45] Stateless neural meta-learning using second-order gradients View paper
- [46] Meta-Gradient Reinforcement Learning with an Objective Discovered Online View paper
- [47] Graph Meta Learning via Local Subgraphs View paper
- [48] Theoretical Convergence of Multi-Step Model-Agnostic Meta-Learning View paper
- [49] VIABLE: Fast Adaptation via Backpropagating Learned Loss View paper
- [50] Continual Meta Learning View paper
- [51] Nonlinear meta-learning can guarantee faster rates View paper
- [52] Fast-Rate PAC-Bayesian Generalization Bounds for Meta-Learning View paper
- [53] Meta-D2AG: Causal Graph Learning with Interventional Dynamic Data View paper
- [54] Research on traffic flow forecasting based on interactive dynamic meta-graph learning View paper
- [55] Self-Distillation with Meta Learning for Knowledge Graph Completion View paper
- [56] PyTorch Foundations and Applications: Definitive Reference for Developers and Engineers View paper
- [57] Learning to Propagate for Graph Meta-Learning View paper
- [58] Learning to Coordinate: Distributed Meta-Trajectory Optimization Via Differentiable ADMM-DDP View paper