香港中文大學（深圳）
The Chinese University of Hong Kong, Shenzhen

SCHOOL OF DATA SCIENCE
數據科學學院

DYP Research and Travel Award Poster Session

12:00-14:00, November 26, 2025
Lobby of Dao Yuan Building

R01. Enhanced Dual-sided Projection with Tensor t-decomposition for Efficient Model Reduction of Polynomial Systems

Presenter: **Wang, Ziming**

Abstract:

This paper presents a novel interpolation-based model order reduction technique designed specifically for high-order polynomial systems. By building upon the existing two-sided projection interpolation approach, a new subspace construction method is proposed. In this method, the basis of the matrix successfully retains the crucial property of matching the generalized moments of the polynomial system. To further exploit the characteristics of high-order tensors, a tensor decomposition method, namely t-eigen decomposition, is introduced based on tensor t-product and t-SVD. This allows for the derivation of low-rank approximations for the Hessian tensor and cross term tensor, effectively capitalizing on the high-dimensional linear algebraic structure of the polynomial system and significantly reducing the complexity and storage requirements during the offline stage. The feasibility and efficiency of the proposed method are verified through a series of numerical experiments.

R02. Prof. Robot: Differentiable Robot Rendering Without Static and Self-Collisions

Presenter: **Lei, Jiabao**

Abstract:

Differentiable rendering has gained significant attention in the field of robotics, with differentiable robot rendering emerging as an effective paradigm for learning robotic actions from image-space supervision. However, the lack of physical world perception in this approach may lead to potential collisions during action optimization.

In this work, we introduce a novel improvement on previous efforts by incorporating physical awareness of collisions through the learning of a neural robotic collision classifier. This enables the optimization of actions that avoid collisions with static, non-interactable environments as well as the robot itself. To facilitate effective gradient optimization with the classifier, we identify the underlying issue and propose leveraging Eikonal regularization to ensure consistent

gradients for optimization. Our solution can be seamlessly integrated into existing differentiable robot rendering frameworks, utilizing gradients for optimization and providing a foundation for future applications of differentiable rendering in robotics with improved reliability of interactions with the physical world.

Both qualitative and quantitative experiments demonstrate the necessity and effectiveness of our method compared to previous solutions.

R03. DriveX: Omni Scene Modeling for Learning Generalizable World Knowledge in Autonomous Driving

Presenter: **Shi, Chen**

Abstract:

Data-driven learning has advanced autonomous driving, yet task-specific models struggle with out-of-distribution scenarios due to their narrow optimization objectives and reliance on costly annotated data. We present DriveX, a self-supervised world model that learns generalizable scene dynamics and holistic representations (geometric, semantic, and motion) from large-scale driving videos. DriveX introduces Omni Scene Modeling (OSM), a module that unifies multimodal supervision-3D point cloud forecasting, 2D semantic representation, and image generation-to capture comprehensive scene evolution. To simplify learning complex dynamics, we propose a decoupled latent world modeling strategy that separates world representation learning from future state decoding, augmented by dynamic-aware ray sampling to enhance motion modeling. For downstream adaptation, we design Future Spatial Attention (FSA), a unified paradigm that dynamically aggregates spatiotemporal features from DriveX's predictions to enhance task-specific inference. Extensive experiments demonstrate DriveX's effectiveness: it achieves significant improvements in 3D future point cloud prediction over prior work, while attaining state-of-the-art results on diverse tasks including occupancy prediction, flow estimation, and end-to-end driving. These results validate DriveX's capability as a general-purpose world model, paving the way for robust and unified autonomous driving frameworks.

R04. BadVideo: Stealthy Backdoor Attack against Text-to-Video Generation

Presenter: **Wang, Ruotong**

Abstract:

Text-to-video (T2V) generative models have rapidly advanced and found widespread applications across fields like entertainment, education, and marketing. However, the adversarial vulnerabilities of these models remain rarely explored. We observe that in T2V generation tasks, the generated videos often contain substantial redundant information not explicitly specified in the text prompts, such as environmental elements, secondary objects, and additional details, providing opportunities for malicious attackers to embed hidden harmful content. Exploiting this

inherent redundancy, we introduce BadVideo, the first backdoor attack framework tailored for T2V generation. Our attack focuses on designing target adversarial outputs through two key strategies: (1) Spatio-Temporal Composition, which combines different spatiotemporal features to encode malicious information; (2) Dynamic Element Transformation, which introduces transformations in redundant elements over time to convey malicious information. Based on these strategies, the attacker's malicious target seamlessly integrates with the user's textual instructions, providing high stealthiness. Moreover, by exploiting the temporal dimension of videos, our attack successfully evades traditional content moderation systems that primarily analyze spatial information within individual frames. Extensive experiments demonstrate that BadVideo achieves high attack success rates while preserving original semantics and maintaining excellent performance on clean inputs. Overall, our work reveals the adversarial vulnerability of T2V models, calling attention to potential risks and misuse.

R05. A Distributional Approach to Uncertainty-Aware Preference Alignment Using Offline Demonstrations

Presenter: **Xu, Sheng**

Abstract:

Designing reward functions in Reinforcement Learning (RL) often demands significant task-specific expertise. Offline Preference-based Reinforcement Learning (PbRL) provides an effective alternative to address the complexity of reward design by learning policies from offline datasets that contain human preferences between trajectory pairs. Existing offline PbRL studies typically model a reward function by maximizing its likelihood of generating the observed human preferences. However, due to the varying number of samples within the limited dataset, less frequently compared trajectories exhibit greater uncertainty, which potentially leads to unreliable behaviors during reward and policy updates. To solve this issue, in this work, we introduce Uncertainty-Aware PbRL (UA-PbRL) to learn a distributional reward model and a risk-sensitive policy from an offline preference dataset. Our approach employs a Maximum A Posteriori (MAP) objective to update trajectory rewards and incorporates an informative prior to account for the uncertainties. Building upon this reward update, we propose a generative reward model to capture the reward distribution, utilizing the offline distributional Bellman operator and the Conditional Value-at-Risk (CVaR) metric to train a risk-sensitive policy. Experimental results demonstrate that UA-PbRL effectively identifies and avoids states with high uncertainty, facilitating risk-averse behaviors across various tasks, including robot control and language model alignment.

R06. Vevo: Controllable Zero-Shot Voice Imitation with Self-Supervised Disentanglement

Presenter: **Zhang, Xueyao**

Abstract:

The imitation of voice, targeted on specific speech attributes such as timbre and speaking style, is crucial in speech generation. However, existing methods rely heavily on annotated data, and struggle with effectively disentangling timbre and style, leading to challenges in achieving controllable generation, especially in zero-shot scenarios. To address these issues, we propose Vevo, a versatile zero-shot voice imitation framework with controllable timbre and style. Vevo operates in two core stages: (1) Content-Style Modeling: Given either text or speech's content tokens as input, we utilize an autoregressive transformer to generate the content-style tokens, which is prompted by a style reference; (2) Acoustic Modeling: Given the content-style tokens as input, we employ a flow-matching transformer to produce acoustic representations, which is prompted by a timbre reference. To obtain the content and content-style tokens of speech, we design a fully self-supervised approach that progressively decouples the timbre, style, and linguistic content of speech. Specifically, we adopt VQ-VAE as the tokenizer for the continuous hidden features of HuBERT. We treat the vocabulary size of the VQ-VAE codebook as the information bottleneck, and adjust it carefully to obtain the disentangled speech representations. Solely self-supervised trained on 60K hours of audiobook speech data, without any fine-tuning on style-specific corpora, Vevo matches or surpasses existing methods in accent and emotion conversion tasks. Additionally, Vevo's effectiveness in zero-shot voice conversion and text-to-speech tasks further demonstrates its strong generalization and versatility.

R07. MoDGS: Dynamic Gaussian Splatting from Casually-captured Monocular Videos with Depth Priors

Presenter: **Liu, Qingming**

Abstract:

In this paper, we propose MoDGS, a new pipeline to render novel views of dy namic scenes from a casually captured monocular video. Previous monocular dynamic NeRF or Gaussian Splatting methods strongly rely on the rapid move ment of input cameras to construct multiview consistency but struggle to recon struct dynamic scenes on casually captured input videos whose cameras are either static or move slowly. To address this challenging task, MoDGS adopts recent single-view depth estimation methods to guide the learning of the dynamic scene. Then, a novel 3D-aware initialization method is proposed to learn a reasonable deformation field and a new robust depth loss is proposed to guide the learning of dynamic scene geometry. Comprehensive experiments demonstrate that MoDGS is able to render high-quality novel view images of dynamic scenes from just a casually captured monocular video, which outperforms state-of-the-art meth ods by a significant margin. The code will be publicly available.

R08. Understanding Constraint Inference in Safety-Critical Inverse Reinforcement Learning

Presenter: **Yue, Bo**

Abstract:

In practical applications, the underlying constraint knowledge is often unknown and difficult to specify. To address this issue, recent advances in Inverse Constrained Reinforcement Learning (ICRL) have focused on inferring these constraints from expert demonstrations. However, the ICRL approach typically characterizes constraint learning as a tri-level optimization problem, which is inherently complex due to its interdependent variables and multiple layers of optimization. Considering these challenges, a critical question arises: Can we implicitly embed constraint signals into reward functions and effectively solve this problem using a classic reward inference algorithm? The resulting method, known as Inverse Reward Correction (IRC), merits investigation. In this work, we conduct a theoretical analysis comparing the sample complexities of both solvers. Our findings confirm that the IRC solver achieves lower sample complexity than its ICRL counterpart. Nevertheless, this reduction in complexity comes at the expense of generalizability. Specifically, in the target environment, the reward correction terms may fail to guarantee the safety of the resulting policy, whereas this issue can be effectively mitigated by transferring the constraints via the ICRL solver. Advancing our inquiry, we investigate conditions under which the ICRL solver ensures-optimality when transferring to new environments. Empirical results across various environments validate our theoretical findings, underscoring the nuanced trade-offs between complexity reduction and generalizability in safety-critical applications.

R09. Flow-Based Delayed Hawkes Process

Presenter: **Yang, Chao**

Abstract:

Multivariate Hawkes processes are classic temporal point process models for event data. These models are simple and parametric in nature, offering interpretability by capturing the triggering effects between event types. However, these parametric models often struggle with low model capacity, limiting their expressive power to capture heterogeneous data patterns influenced by latent variables. In this paper, we propose a simple yet powerful extension: the Flow-based Delayed Hawkes Process, which integrates Normalizing Flows as a generative model to parameterize the Hawkes process. By generating all model parameters through the flow-based network, our approach significantly improves flexibility and expressiveness while preserving interpretability. We provide theoretical guarantees by proving the identifiability of the model parameters and the consistency of the maximum likelihood estimator under mild assumptions. Extensive experiments on both synthetic and real-world datasets show that our model outperforms existing baselines in capturing intricate and heterogeneous event dynamics.

R10. Action-Priority Driven Policy Optimization for Flexible Job Shop Scheduling Problem

Presenter: **Zheng, Wenjun**

Abstract:

Efficient scheduling significantly impacts productivity and operational performance, particularly in flexible resource allocation scenarios. The Flexible Job Shop Scheduling Problem (FJSP) extends traditional scheduling by allowing each operation to be processed on multiple machines, enhancing practical applicability but increasing computational complexity. Existing methods, ranging from exact solvers to heuristic/metaheuristic algorithms, either incur high computational costs or yield suboptimal solutions. Recent reinforcement learning (RL)-based approaches attempt a balance between optimality and efficiency; however, intrinsic characteristics of FJSP remain underutilized for better solution quality. This paper introduces the Action-Priority Driven Policy Optimization (APDPO) method, explicitly designed to address three critical challenges: achieving superior solutions, improving interpretability via an intuitive policy representation, and enhancing computational efficiency through direct, gradient-free policy updates. A policy improvement theorem ensures performance enhancement under specified conditions. Comparative evaluations confirm that APDPO outperforms existing approaches by delivering better solution quality, reduced computational requirements, and greater decision transparency.

R11. A Unified Framework for Cell-type-specific eQTL Prioritization by Integrating Bulk and scRNA-seq Data

Presenter: **Yu, Xinyi**

Abstract:

Genome-wide association studies (GWASs) have identified numerous genetic variants associated with complex traits, yet the biological interpretation remains challenging, especially for variants in non-coding regions. Expression quantitative trait locus (eQTL) studies have linked these variations to gene expression, aiding in identifying genes involved in disease mechanisms. Traditional eQTL analyses using bulk RNA sequencing (bulk RNA-seq) provide tissue-level insights but suffer from signal loss and distortion due to unaddressed cellular heterogeneity. Recently, single-cell RNA-seq (scRNA-seq) has provided higher resolution, enabling cell-type-specific eQTL (ct-eQTL) analyses. However, these studies are limited by their smaller sample sizes and technical constraints. In this work, we present a statistical framework, IBSEP, which integrates bulk RNA-seq and scRNA-seq data for enhanced ct-eQTL prioritization. Our method employs a hierarchical linear model to combine summary statistics from both data types, overcoming the limitations while leveraging the advantages associated with each technique. Through extensive simulations and real data analyses, including peripheral blood mononuclear cells and brain cortex datasets, IBSEP demonstrated superior performance in identifying ct-eQTLs compared to existing methods. Our approach unveils transcriptional regulatory

mechanisms specific to cell types, offering deeper insights into the genetic basis of complex diseases at a cellular resolution.

R12. The Impact of Autonomous Vehicles on Ride-Hailing Platforms with Strategic Human Drivers

Presenter: **Wu, Xinyuan**

Abstract:

We consider a ride-hailing platform that operates a mixed fleet of autonomous vehicles (AVs) and conventional vehicles (CVs), where AVs are fully controlled by the platform and CVs are operated by self-interested human drivers. Each vehicle is modeled using a Markov Decision Process where the vehicle maximizes long-run average rewards by choosing its repositioning actions. The behavior of CVs corresponds to a large-scale game in which agents interact through resource constraints that result in fluid queues. To optimize the mixed AV-CV system for arbitrary networks, we formulate a bi-level optimization problem OPT in which the platform moves first by controlling the demand revealed to the CVs and subsequently assigning the optimal actions to the AVs, while the CVs react by forming an equilibrium characterized by the solution to a convex optimization problem. We prove several structural properties of the optimal solution and analyze simple heuristics, such as AV-first, where we solve for the optimal dispatch of AVs without taking into account the subsequent reaction of CVs. We also propose three numerical algorithms to solve OPT, which is a non-convex non-smooth problem, and evaluate their performance for large networks. Finally, we use our computational tools to show some interesting trends in the optimal AV-CV fleet dimensioning when vehicle supply is exogenous and endogenous, and apply these results to New York City using demand and trip-time data from real-world taxi service datasets. Our results suggest that our model can be used to predict traffic behavior and optimize mixed-fleet deployment given topology and cost/reward information.

T01. Prophet Inequalities with Stochastic Choices

Presenter: **Wang, Shaoyu**

Abstract:

In the classic single-choice prophet inequality problem, the decision-maker's reward is determined by the first admitted item. In this paper, we study a generalized variant with two key modifications: (1) after admitting each item, the decision-maker must also optimize an associated action, and (2) the final reward is derived from at most one of the admitted items, selected via a stochastic choice function. This generalization introduces significant technical challenges, including complex dependency structures among admitted items and non-submodular reward functions that prevent direct application of standard prophet inequality techniques. Despite these challenges, we establish that under mild conditions, a computationally efficient, order-oblivious, and non-adaptive threshold policy guarantees a constant-factor approximation to the optimal

objective. Furthermore, we extend our analysis to settings where only a limited number of items can be admitted, demonstrating that our results on item arrival information generalize existing constrained-order prophet inequalities for both single and multiple-unit cases without stochastic choices. Additionally, we explore scenarios in which a knapsack constraint is imposed on the admitted items and propose efficient heuristic approaches to address this constraint.

T02. Dynamic Patient Remote Monitoring under Alert Fatigue

Presenter: **Yang, Wanting**

Abstract:

We study the optimal appointment scheduling policy for chronic disease patients who use imperfect self-monitoring devices. We depart from existing literature by modeling patient no-show behavior as endogenous, driven by an ``alert fatigue'' state that accumulates after false-positive alerts. We formulate this problem as a discrete-time, finite-horizon Partially Observable Markov Decision Process (POMDP), where the state captures both the provider's belief of the patient's true health and the patient's accumulated fatigue. We analytically characterize the structure of the optimal policy. We prove it is a threshold policy based on the patient's posterior health belief, and we analyze the monotonicity of this threshold with respect to both behavioral uncertainty (fatigue) and diagnostic uncertainty (error rates). We generalize this structure, proving that for the ξNξ-patient competitive case, the belief space is partitioned into ξN+1ξ decision regions that all intersect at a unique indifference point. Our findings provide a theoretical foundation for designing scheduling policies that balance the immediate benefits of clinical intervention with the long-term erosion of patient engagement.

T03. X2-DFD: A framework for eXplainable and eXtendable Deepfake Detection

Presenter: **Chen, Yize**

Abstract:

Deepfake content is rapidly evolving in both scale and realism, making reliable and interpretable detection increasingly crucial. While recent multimodal large models demonstrate promising semantic understanding, they still fall short in identifying fine-grained visual artifacts that are essential for robust deepfake detection. Existing artifact-based detectors, on the other hand, excel at signal-level cues but lack generalization across diverse forgery types. This work addresses these complementary strengths and weaknesses through X2-DFD, the first framework designed to be both explainable and extendable for deepfake detection. Our framework begins with a comprehensive capability characterization of multimodal models, revealing that they possess strong semantic reasoning but uneven sensitivity to local forgery traces. Based on these findings, we design a modular system that strategically combines multimodal reasoning with specialized artifact detectors. X2-DFD adaptively selects the most informative signal pathways depending on the type of forgery, enabling both robust performance and transparent explanation of the model's

decision process. Empirically, X2-DFD achieves state-of-the-art performance across multiple benchmark datasets while providing human-interpretable visual and textual explanations. Notably, our framework excels at cross-dataset generalization, a major challenge in deepfake detection, and maintains stability against previously unseen forgery patterns. The extendable design also allows new detectors or modalities to be integrated with minimal cost, enhancing its practicality for real-world applications. Overall, X2-DFD offers a principled and scalable approach to trustworthy synthetic media detection, advancing both performance and interpretability in a unified framework.

T04. Searching and Detecting Structurally Similar Communities in Large Heterogeneous Information Networks

Presenter: **Wang, Shu**

Abstract:

Heterogeneous information networks (HINs) are prevalent in various domains, including bibliographic information networks, social media, and knowledge graphs. As a fundamental topic in HIN mining, community mining has found various real applications, such as recommendation, biological data analysis, and event organization. Most existing works often rely on meta-paths, relational constraints, spectral partitioning, label propagation, and network representation to define the communities. However, almost all these works do not explicitly consider the structural similarity between vertices, which plays a vital role in modeling communities and also ignore the specific roles of vertices. In this paper, we propose a novel community model, called structurally similar community (SSC), which models the HIN communities by explicitly considering the structural similarity between vertices. In particular, SSC can not only support various structural similarity measures, but also identify different roles of the vertices in the community, such as cores, noncores, hubs, and outliers. Based on the SSC, we develop fast online and index-based algorithms that support both efficient searching and detecting SSCs in large HINs, where the former one searches an SSC containing a specific query vertex while the latter one detects all the SSCs from the HIN. Extensive experiments on real-world datasets demonstrate the effectiveness of SSC model in revealing meaningful communities and the high efficiency of our proposed algorithms.

T05. UTCS: Effective Unsupervised Temporal Community Search with Pre-training of Temporal Dynamics and Subgraph Knowledge

Presenter: **Zhang, Yue**

Abstract:

In many real-world applications, the evolving relationships between entities can be modeled as temporal graphs, where each edge has a timestamp representing the interaction time. As a fundamental problem in graph analysis, {\it community search (CS)} in temporal graphs has

received growing attention but exhibits two major limitations: (1) Traditional methods typically require predefined subgraph structures, which are not always known in advance. (2) Learning-based methods struggle to capture temporal interaction information. To fill this research gap, in this paper, we propose an effective \textbf{U}nsupervised \textbf{T}emporal \textbf{C}ommunity \textbf{S}earch with pre-training of temporal dynamics and subgraph knowledge model (\textbf{UTCS}). UTCS contains two key stages: offline pre-training and online search. In the first stage, we introduce multiple learning objectives to facilitate the pre-training process in the unsupervised learning setting. In the second stage, we identify a candidate subgraph and compute community scores using the pre-trained node representations and a novel scoring mechanism to determine the final community members. Experiments on five real-world datasets demonstrate the effectiveness.

T06. Improving Task-Specific Multimodal Sentiment Analysis with General MLLMs via Prompting

Presenter: **Zhang, Haoyu**

Abstract:

The field of Multimodal Sentiment Analysis (MSA) has recently witnessed an emerging direction seeking to tackle the issue of data incompleteness. Recognizing that the language modality typically contains dense sentiment information, we consider it as the dominant modality and present an innovative Language-dominated Noise-resistant Learning Network (LNLN) to achieve robust MSA. The proposed LNLN features a dominant modality correction (DMC) module and dominant modality based multimodal learning (DMML) module, which enhances the model's robustness across various noise scenarios by ensuring the quality of dominant modality representations. Aside from the methodical design, we perform comprehensive experiments under random data missing scenarios, utilizing diverse and meaningful settings on several popular datasets (e.g., MOSI, MOSEI, and SIMS), providing additional uniformity, transparency, and fairness compared to existing evaluations in the literature. Empirically, LNLN consistently outperforms existing baselines, demonstrating superior performance across these challenging and extensive evaluation metrics.

T07. Toward Exploratory Inverse Constraint Inference with Generative Diffusion Verifiers

Presenter: **Zhao, Runyi**

Abstract:

An important prerequisite for safe control is aligning the policy with the underlying constraints in the environment. In many real-world applications, due to the difficulty of manually specifying

these constraints, existing works have proposed recovering constraints from expert demonstrations by solving the Inverse Constraint Learning (ICL) problem. However, ICL is inherently ill-posed, as multiple constraints can equivalently explain the experts' preferences, making the optimal solutions not uniquely identifiable. In this work, instead of focusing solely on a single constraint, we propose the novel approach of Exploratory ICL (ExICL). The goal of ExICL is to recover a diverse set of feasible constraints, thereby providing practitioners the flexibility to select the most appropriate constraint based on the practical needs of deployment. To achieve this goal, we design a generative diffusion verifier that guides the trajectory generation process using the probabilistic representation of an optimal constrained policy. By comparing these decisions with those made by expert agents, we can efficiently verify a candidate constraint. Driven by the verification feedback, ExICL implements an exploratory constraint update mechanism that strategically facilitates diversity within the collection of feasible constraints. Our empirical results demonstrate that ExICL can seamlessly and reliably generalize across different tasks and environments.