Tutorial on **Advances in Simulation Technology for Web Applications**

– The Web Conf 2023

Presenters:
Da Xu (LinkedIn)
Shuyuan Xu (Rutgers University)
Bo Yang (Amazon)

Tutorial Webpage: https://foundation4recsys.github.io/Tutorial-WWW23/
Contact: daxu5180@gmail.com

Organizers

Da Xu

Staff AI Engineer @ LinkedIn
ML Manager @ WalmartLabs
Stats @ UC Berkeley

ML in Production, Theory and Foundation of IR & Recsys, Causal Inference, GenAI

Bo Yang

Applied Scientist @ Amazon
ML Engineer @ Linkedin
Stats @UVA

Advertising System, Causal Inference

Shuyuan Xu

CS PhD Candidate @ Rutgers University

ML, Recommender Systems, Causal Inference

Yongfeng Zhang

Assistant Professor @ Rutgers University

ML, Machine Reasoning, Information Retrieval, Recommender Systems, Explainable AI, Fairness in AI, and AI Economics.

TABLE OF CONTENTS

Q&A will be held at the end of the tutorial

01

Opening Remark

Overview, Motivation, Introduction & Scope

Milestones

90's

Simulation with simple physical models (with differential equations) discovered the Internet Service Provider (ISP) marketing model – no monthly fee + advertising – the foundation of modern IT companies' business model

00's

Use more fine-grained stochastic and physical models that reflect particular user behaviors to evaluate and understand simpler phenomena of the system

10's

Employ ideas from causal structure modeling to account for prior knowledge and other more complex aspects of real-world mechanisms for evaluating and studying system behaviors

2015+

Incorporate the counterfactual and sequential interaction nature of Web systems and combine them with the tools and ideas developed from reinforcement learning to drive realistic long-term simulations

2020+

Generative-AI-driven simulation where system components and transactions will be dominated by generative AI agents

Rough Numbers on the Percentage of Major Data Mining and Recsys Conference Publications that Have Simulation Study*

<2% — 10 YEARS AGO

~4% — PAST 10 - 3 YEARS

~7% — PAST 3 YEARS

* We use the term ‘simulation study’ as the indicator for doing our web-crawler-based survey, which is subject to missingness and bias

Motivation

Aim to provide a systematic review for the growth and advancement of a technology that is becoming progressingly crucial in the research and development of algorithms and systems for Web applications, with high relevance to both academia and industry

Introduction

Many simulation technologies that have emerged in the past 20 years can be encapsulated by the concept of agent-based simulation, which holds significant implications as we venture into the era of generative AI. We present and contextualize this concept with concrete design and realization patterns in the fields of information retrieval, recommender system, and advertising.

Scope

Since this is a 1.5 hr tutorial, we focus more on the conceptual aspects and point interested readers to the references for detail. Our discussion revolves around supporting the development, refinement, analysis and comparisons of algorithms and systems. Our discussions are not supposed to reflect the full extend of real-world behaviors.

02

Overview

Agent-based Simulation for Modern Web Applications

Agent-based Simulation

A Mindset More than Technology

- TL;DR: Creating system of agents and the relationships between them

Benefits

- Provide a natural description of the system
- Provide a visible pathway from **Hypothesis to Action,** bring algorithms out of the lab
- Fine-grained control over **environment complexity**, **reality faithfulness**, and **manipulability** for studying phenomena

Agent Decision

- Can exhibit **complex memory** and **path dependencies**
- Hysteresis, temporal correlations
- Can have complex behaviors such as **learning** and **adaptation**

Agent interactions

- Interactions are often **complex**, **non-linear**, **discontinuous**, and the state and action spaces are not fixed
- Emergent and interesting phenomena arises from agent interactions
- Often **heterogeneous** and exhibit such as network effects and delayed feedback

Beyond Average Analysis

- The population is rarely homogeneous, and can often characterized as mixture of clusters or hidden factors
- Aggregation smooths out fluctuations and masks heterogeneous behaviors

Agent-based Simulation for Web Applications

Information Retrieval (Example: SimIIR)

(A). User type and topic model, search context pool, query and document generator, search engine
(B,C).CSM (Complex Searcher Model, memory-aware)
(D). Markovian environment and stopping criteria

Domain knowledge
Real-world constraints

Description of the System

A. **Agent model & feature (state)** for such as user, item, document, …
B. **Agent choice model** (e.g. preference, interaction, stopping)
C. **Transition model**
D. **Environment** (reward, constraints, properties,...)

Recsys (Examples)

(A). Real-world user and item (e.g. **MarsGym**)
(B, C). Task- and User-type specific models (e.g. **PyRecGym**)
(D). Organic + Bandit E-commerce sessions (e.g. **RecoGYM**)
SlateQ environment (e.g. **RecSim**)

Marketing & Ads

(A). Open box auction simulation (**Genie**)
(D). Bandit learning for bidding strategies(**AuctionGym**)

Major Scenarios for Using Agent-based Simulation in Web Applications

System Evaluation

Take advantage of the full observability and unlimited horizon of simulation environments to examine algorithms, off-policy evaluation and learning methods, long-term properties of the system, etc

Opportunity Identification

Take advantage of the manipulability of simulation to design purposeful and controlled environments to verify hypothesis, detect patterns, and explore interesting phenomena

Stylized Analysis

Take advantage of the flexibility to create stylized models and environments that are more targeted and simpler to analysis than real-world systems, to investigate, identify and understand the process and mechanisms underlying the phenomena of interest

On Designing Agent-based Simulation

Reality

- Designed to identify, understand or address real-world problems
- Simplified, abstract, reusable, and externally valid components
- Complete ‘sim-to-real’ is often unnecessary

Assumption

- On agent models, states, interactions, transitions, and environment
- On initial conditions and behavior trajectory
- On objective and visible pathway between simulation outcome and hypothesis

Intervention

- Purposefully taking actions on the simulation environment
- Explore consequences, understand causal pathways
- Draw connection and test hypothesis

Faithfulness

Faithfulness to the real-world counterpart of the problem

Complexity

Complexity of the assumptions and hypothesis

Manipulability

Degree of freedom and level of constraints for the acting ability

The tradeoffs are delicate and should be carefully weighing in for different use cases in order to ensure feasibility and draw plausible conclusion.

How to get buy-in from reviewers, clients, colleagues, leadership

Communicate priority,
Find gaps

Define goals and questions to explore

Bring transparency,
Avoid unrealistic expectation

Big picture on capacity, limits, constraints, tradeoffs of the design

Build system / user story, present viable pathway between the design and hypothesis

Conduct prototyping and exploratory analysis to identify and validate the hypothesis and story

Test - Create - Modeling - Investigate

Use data-driven evidence and the underlying mechanism to prove the feasibility and plausibility

Summarize the outcome backed by concrete data points under the hypothesis, system design, and various constraints

03

For Information Retrieval

Search Engine
Conversational System

From Complex Searcher Model to Search-Engine Simulation

[Maxwell, Leif 2016]

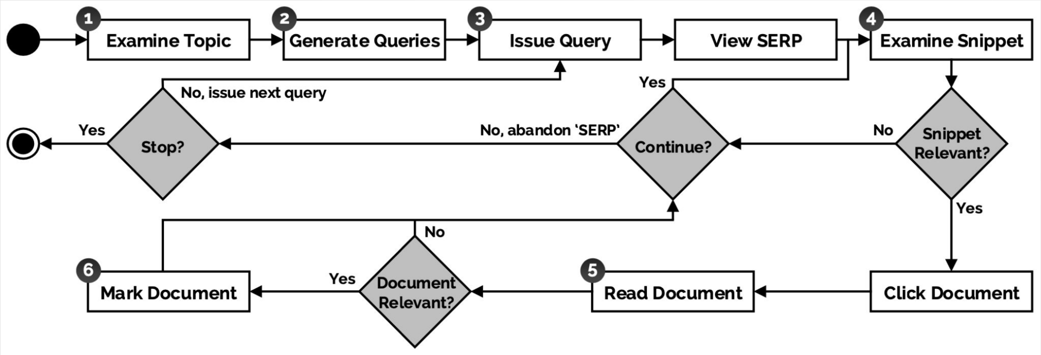

Gery components can be characterized by agent decisions

White components can be characterized by agent models

Basic setup:

- Query agent model (feature model)
- Document agent model (feature model)
- User agent model (feature model, choice model)
- Environment (stopping criteria, search context)
- ……

Goal: 'sim-to-real' scenario for system evaluation

Extend Basic CSM Setup to Accommodate More Scenarios and Design Strategies

[Zerhoudi, 2022]

Controlled user behavior by indexing user agent model with <u>user type</u>

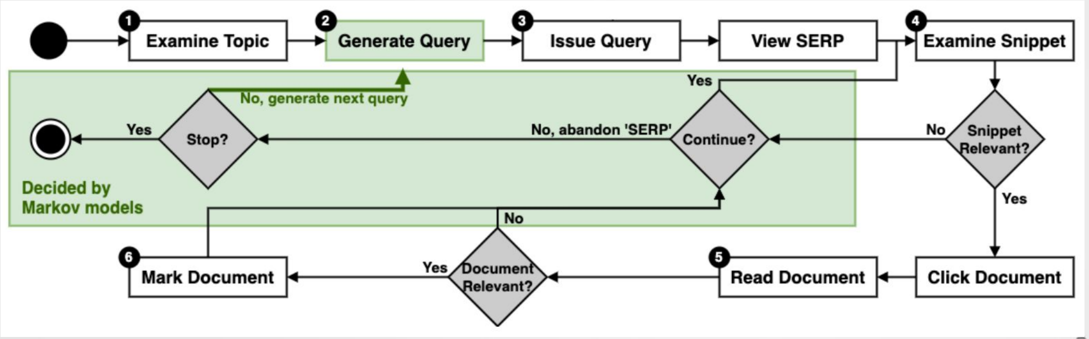

Markovian query generation by adding a <u>transition model</u> to query agent

Adding controlled behavior for the user agent:

- Stylized system evaluation
- User-centric and beyond-average analysis
- Opportunity discovery for different user types

Adding memory and path dependencies to query agent:

- Fine-grained control over environment complexity
- Reflect complex agent decision and interaction
- Faithfulness vs. manipulability

Agenda-based User Simulation for Dialogue System

[Schatzmann 2007]

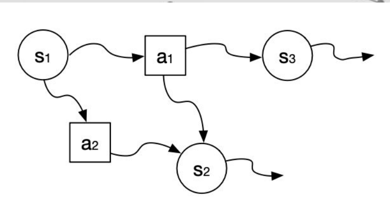

- **User agent** (goal model, agenda model, agenda act model, transition model)
- **Environment** (+ partially observed MDP)

User simulation at a semantic level	***Goal- and Agenda-Based State Representation***	***User act selection***	***State transition model***	***Agent update model***	***Goal update model***
Current state | User action | Intermediate state | System action | New state	A. User behave in a consistent and goal-driven fashion B. Agenda as a stack-like structure containing pending user acts (inform, request)	Pop items from the stack Can be made stochastic	● (Stochastic) push operations where dialogue acts are added to the agenda ● The hidden user constraints and requests changes with a given machine action ● MDP update		

Conversational System Simulation

[Zhang, Krisztian 2020]

- **User agent** (goal model, agenda model, agenda act model, transition model)
- **+ Interaction model, Preference model, Natural language agents**
- **Environment** (+ partially observed MDP)
- **+ conversational agent**

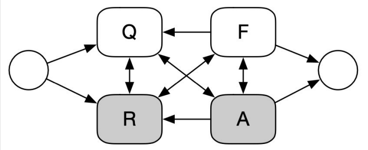

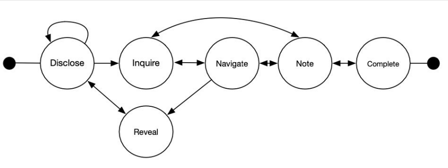

QRFA model
(Query, Request, Feedback, Accept)

CIR6 model
(Conversational item rec model)

Interaction model

Preference model

Single Item
Preference model

Personal knowledge graph

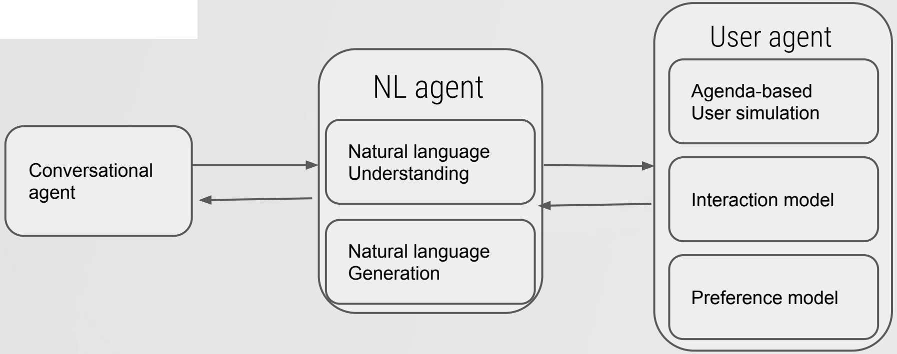

Template-based NL models

Natural language understanding

Natural language generation

04

For Recommender System

Personalized Recommendation

Data Generation Process in RS

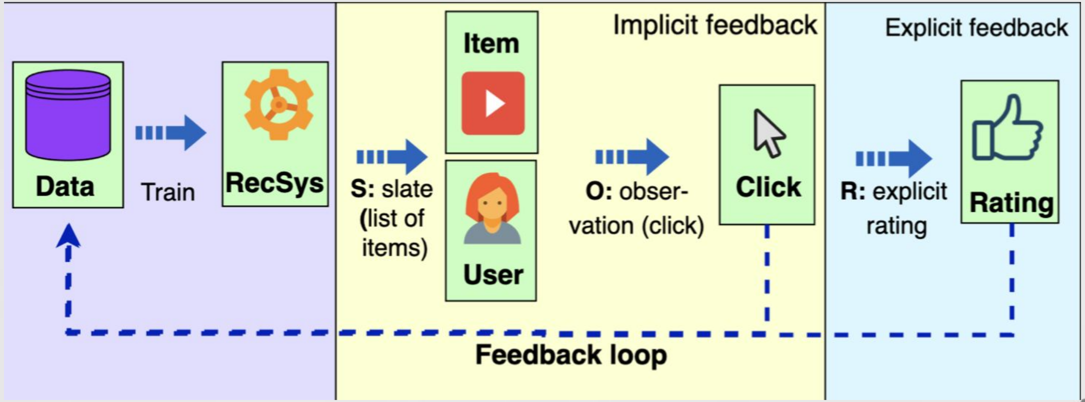

[Stavinova et al. 2022]

- **Users and Items Models**: Models for generating synthetic profiles of users and items.
- **Recommender Systems**: Models for recommending items to users.
- **User Response Models**: Models for providing feedback.

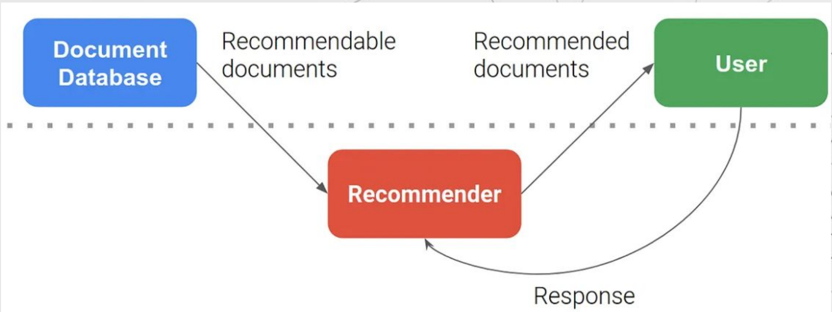

RecSim [Ie et al. 2019]

Simulator Design

Recommender Systems

- The construction of RecSys depends on scenarios and assumptions.
- Could be RL agent or pre-defined models.

Goal

- Scenarios: specific task and scenarios that could take place in the interaction between users and items
- Assumptions: a set of assumptions about the mechanism to satisfy the scenarios.

Users/Items Profile

- Realism: The profile can be generated with or without real data.
- Uncertainty: The profile may include noise and uncertainty.
- Dynamics: The profile can be dynamic or static.

User Response Models

- Related on the scenarios and assumptions..
- Based on user/item features, history, context, etc.
- Generate implicit/explicit feedback, time, etc.

RL-based Simulator Pipeline

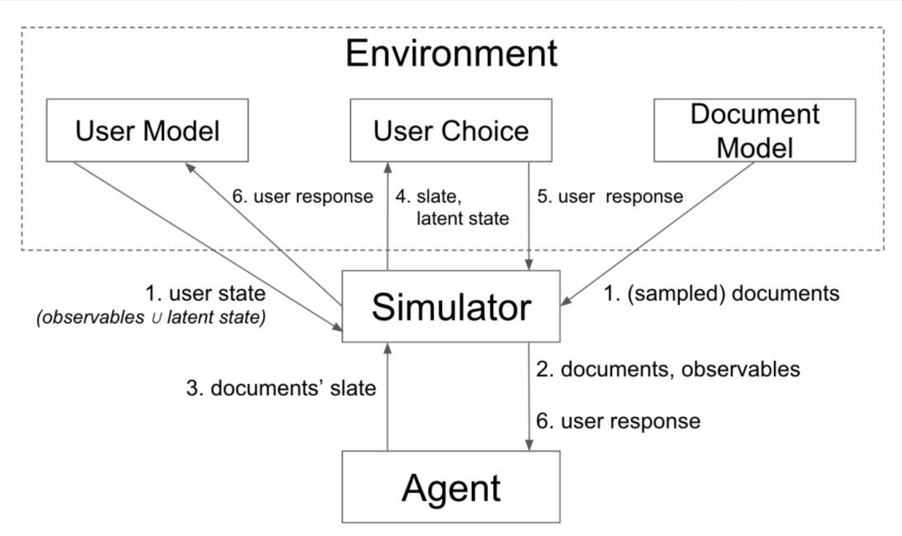

RecSim [Ie et al. 2019]

- The **environment** consists of a user model, a document (item) model and a user-choice model.
- The **agent** serves as a recommender system.
- The **action** is defined as recommending item(s).
- The **reward** will generally be a function of a user's responses

Case Study: RecSim

[Ie et al. 2019]

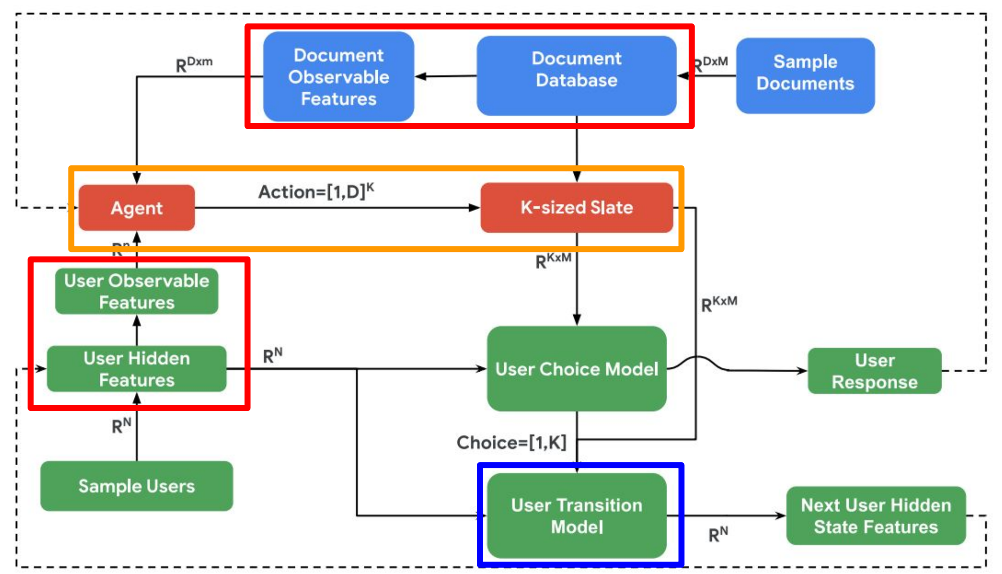

N - number of features that describe the user's hidden state
n - number of features that describe user's observed state
M - number of features describing document hidden state
m - number of features describing document observed state
D - total number of documents in the corpus
K - size of slate

- **Scenario**: providing environments that facilitate the development of new RL algorithm for recommender applications (Collaborative Interactive Recommenders). 'Sim-to-real' is not the priority concern.
- **User/Item profile**: sampled from a prior (based on real data or not) over user/item features, including both latent and observable features
- **Dynamics**: User profile will be updating along with interactions by User Transition Model
- **RecSys**: act as an agent to recommend slates of documents (items) based on observed features.
- **User Response Model**: generate user response depending on observable item features and all user features (latent and observable)

Manipulate the Simulator Generation

[Huang et al. 2020]

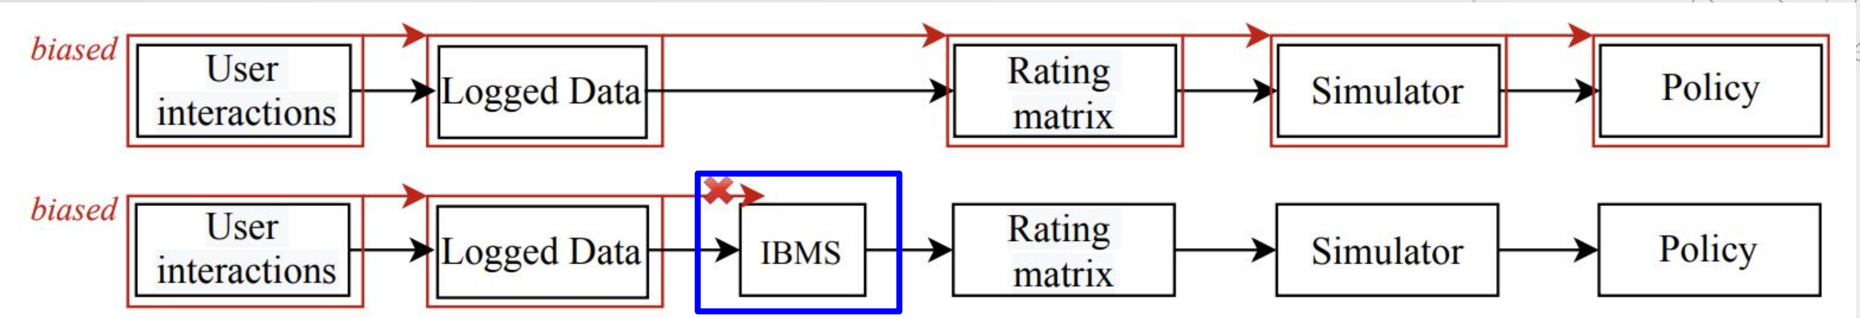

- **Scenario**: interactive recommender.
- **Concerns**: Using real data to build simulator may suffer from bias of real data.
- **User Response Model:** Predicted rating matrix
- **Intermediate Bias Mitigation Step (IBMS):** mitigating the effect of bias before the prediction model is learned by IPS

Case Study: Virtual-Taobao

[Shi et al. 2018]

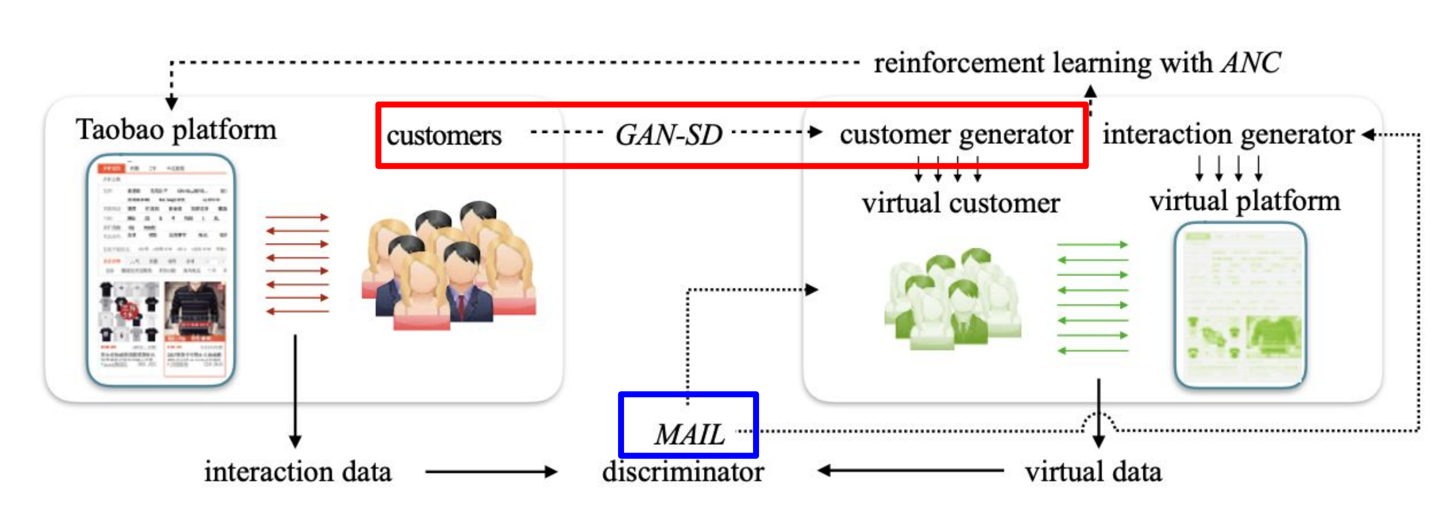

- **Scenario**: Real-world Online Retail Environment (Taobao)
- **User profile**: focus on realism, GAN-SD (GAN Simulation Distribution) to generate consumers similar to real data.
- **Interactions**: Generated by MAIL (multi-agent adversarial imitation learning), training the **customer policy** as well as the **engine policy**.

Case Study: Accordion

[McInerney et al. 2021]

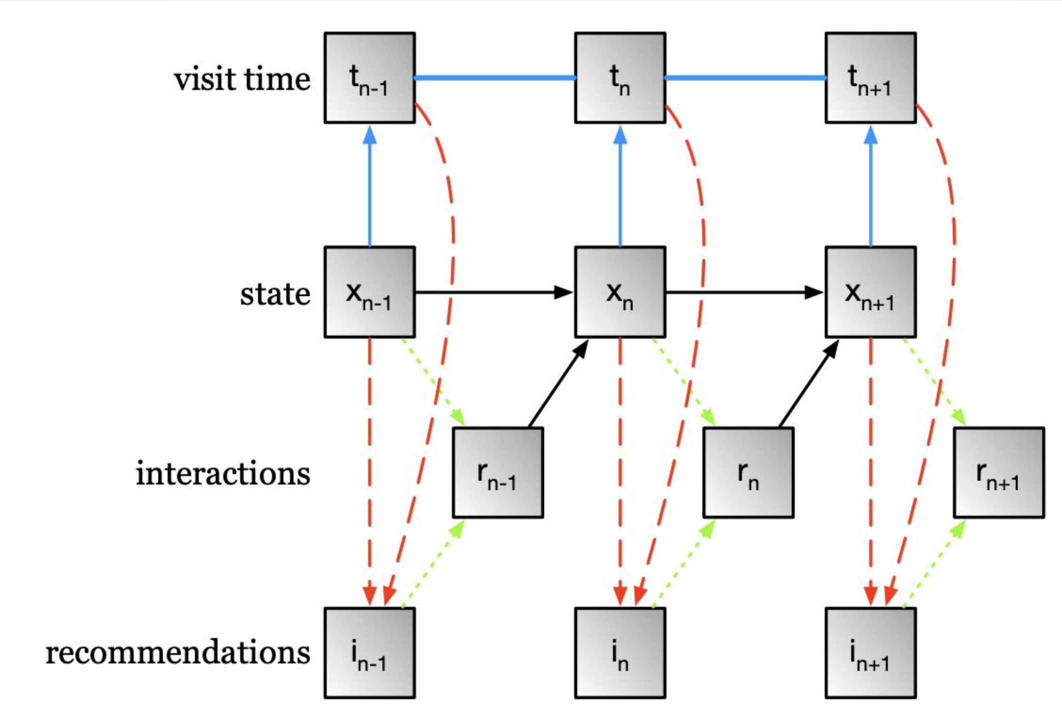

- **Scenario**: simulating Long-term interactive systems. Time-aware recommendation scenario
- **RecSys**: Recommendation models that considering visit time and user state for making recommendations.
- **User Response Models**: Consisted of visit model and selection model.
- **Visit Model**: A Poisson process based method for simulating user visit, time of the visits and the number of interactions in each visit.
- **Selection Model**: simulating outcome of the interaction (e.g., click, purchase, stream)

Simulation for Observation

[Kalimeris et al. 2021]

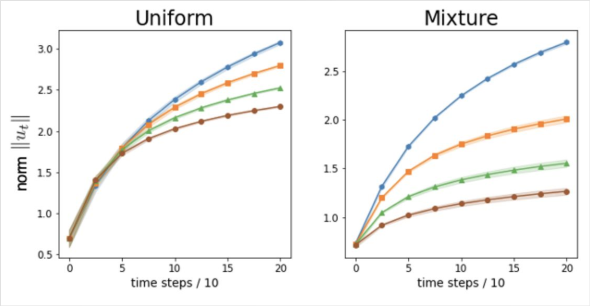

user preference

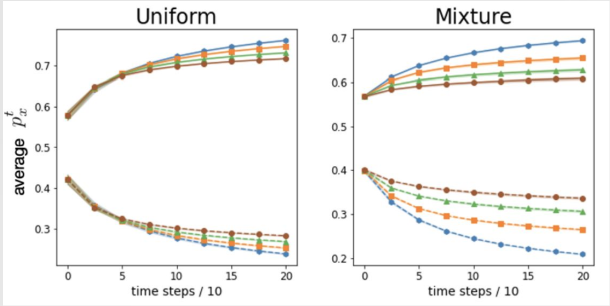

probability of likable vs non-likable item

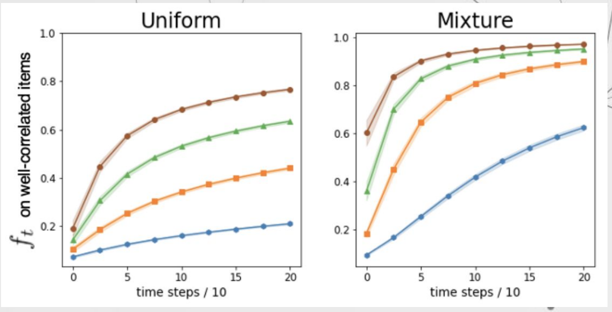

Probability mass on items correlating well with the initial user preference

- **Scenario**: study the preference amplification caused by MF models
- **User profile**: sampled from uniform distribution. Updated after every interactions.
- **Item profile**: sampled from pre-defined distributions (uniform or mixture of two uniform)
- **RecSys**: MF-based models.
- **Conclusion**: preference amplification, echo chambers, filter bubbles.

Manipulate User Response Models for Observation

[Yao et al. 2021]

	Name	Formula
Choice	Lazy	$\mathcal{M}_s^{lazy}(v) = \mathbb{1}[\text{Rank}(v) = 1]$
	Uniform	$\mathcal{M}_s^{uniform}(v) = \frac{1}{k}$
	Ranked	$\mathcal{M}_s^{ranked}(v) \propto \frac{1}{\log(1+\text{Rank}(v))}$
	α-preference	$\mathcal{M}_s^{\alpha-prefer}(v) \propto e^{\alpha\rho(v)}$
Feedback	Positive	$\mathcal{M}_f^{positive}(v) = +1$
	β-preference	$\mathcal{M}_f^{\beta-prefer}(v) = \begin{cases} \beta & \rho(v) \geq \rho_0 \\ -\beta & \text{else} \end{cases}$

- **Scenario**: measuring the impact of a recommender system (popularity bias) under different types of user behavior.
- **User profile**: trajectory obtained from the real data
- **Item profile**: item features obtained from the real data (popularity denoted as \rho(v))
- **User Response Models**: Consisted of choice model (Implicit) and feedback model (binary explicit).
- **RecSys**: Pre-trained MF and RNN.

05

For Marketing and Advertising

Bidding, Pricing, Ads Allocation

How RTB(Real-Time Bidding) Works

[Weinan, 2014]

- **Advertisers** create advertisement campaigns and place bids that describe how much they are willing to pay to see their ads displayed or clicked.
- **Publishers** provide web services. They provide the infrastructure to collect advertiser bids and display selected ads and expect to receive payments from the advertisers.
- **Users** reveal information about their current interests. They are offered web pages that contain a selection of ads. Users may view/click/place an order on an advertisement.

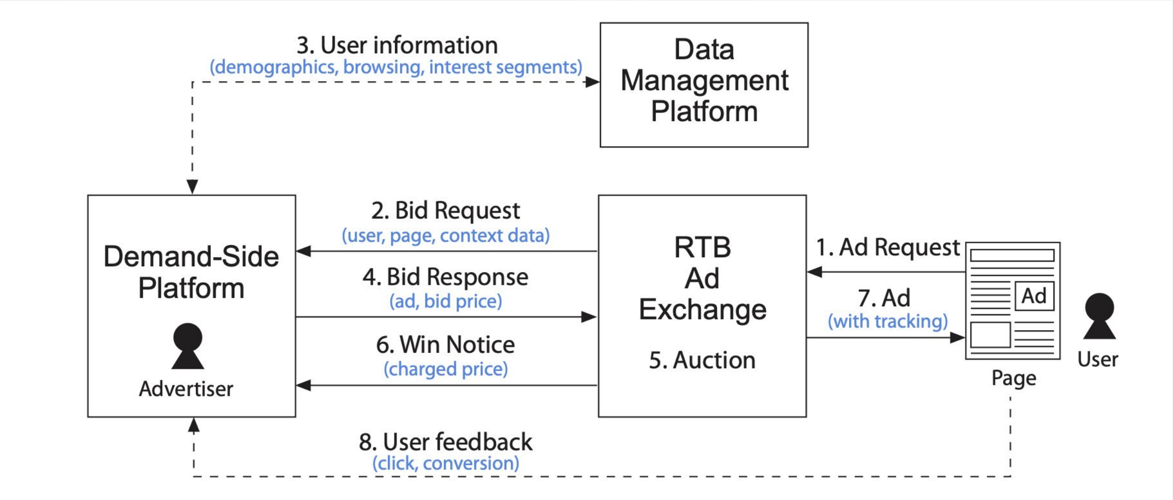

Auction Mechanism

Myerson [1981] Auction is generally regarded as a fair and transparent way for advertisers and publishers to agree with a price quickly, whilst enabling the best possible sales outcome.

- ***Publishers*** have access to partial information about the market demand from historic transactions. However, they do not have knowledge about how much an individual ad impression is worth on the market.
- ***Advertisers*** may have different (private) valuations of a given ad impression.

Second-price Auction.

- Truthful bidding is a dominant strategy in second-price auctions under several assumptions:
  - Bidder knows their expected valuation given a context
  - Placed bids do not influence the value of the good
  - Competitors all have access to the same information
  - Repeated rounds of auctions are statistically independent

First-price Auction.

- Bidders should optimally shade their bids to balance the trade-off between paying lower prices and decreasing their chances of winning.

Causation Issues in Computational Advertising

With the development of new ads marketplace algorithms there are always 'what if' questions with any policy, parameter or model change in the system that yields a different ad allocation.

Controlled Experiments (Kohava [2008])

- **Expensive** because they demand a complete implementation of the proposed modifications.
- **Slow** because each experiment typically demands a couple months.
- Splitting advertisers into treatment and control groups demands special attention because each auction involves multiple advertisers. **Simultaneously controlling for both users and advertisers is probably impossible**.

Statistical Methods (Simpson [1951])

- **Cheaper and faster** statistical methods are needed to drive essential aspects of the development of RTB engine. However, interpreting cheap and fast data can be very deceiving.
- **Confounding Data**: Assessing the consequence of an intervention is generally challenging because of difficulty to determine whether the observed effect is a consequence of the intervention or has uncontrolled causes.

Modeling Causal Systems

Bayir [2019] proposes counterfactual policy estimation framework called Genie to optimize Sponsored Search Marketplace. Genie employs an open box simulation engine with click calibration model to compute the KPI impact of any modification to the system.

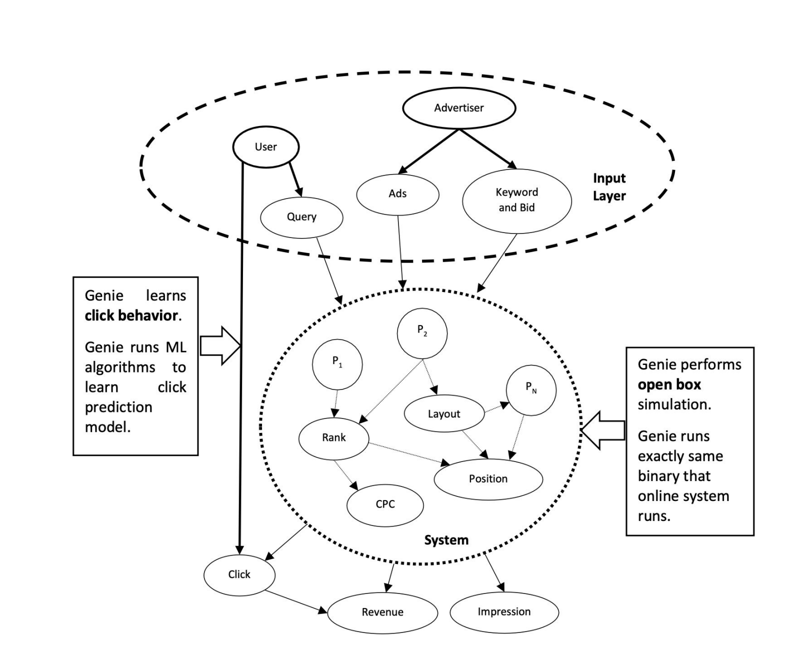

- **KPI impact** of any policy can be estimated by replaying training data with the modified policy and using user click behavior model that has tolerable noise.
- **Explore much wider parameter space** since it does not require real traffic with modification/exploration cost.
- Be leveraged to **tune completely new policies** where creating initial experiment is very costly due to cold start problem.

COUNTERFACTUAL REASONING AND LEARNING (Bottou [2013])

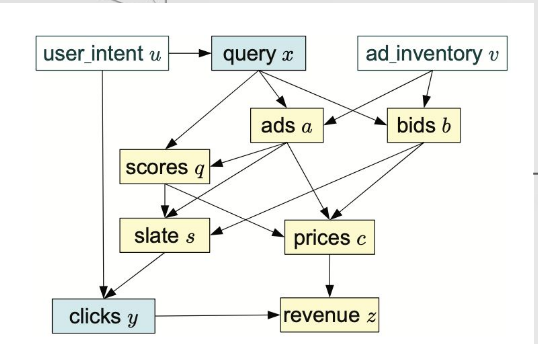

Casual Graph

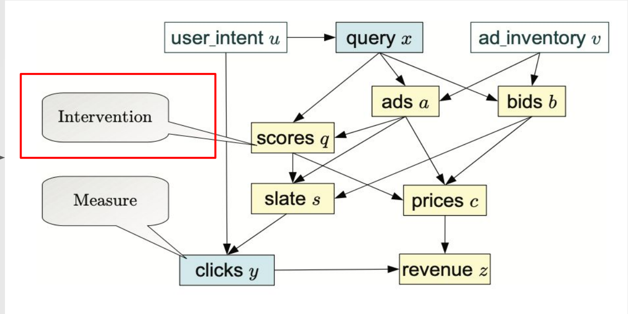

Intervention

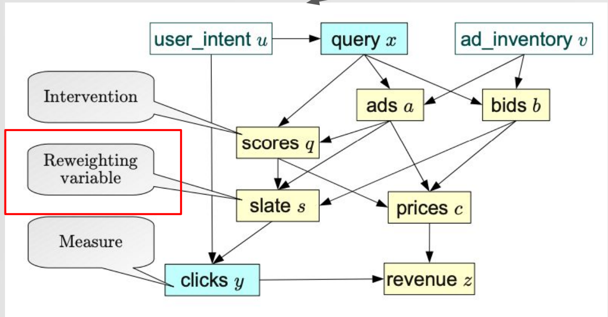

Displace reweighting point

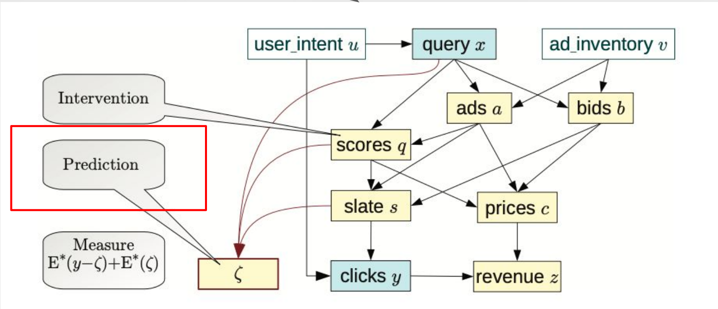

Use prediction point

Bandit learning for bidding strategies

Olivier [2022] introduce AuctionGym: a simulation environment that enables the use of bandit learning for bidding strategies in online advertising auctions.

Simulating auctions end-to-end

- an impression opportunity arises with features x~P(X)
- auction presents this opportunity to bidders
- bidders decide on an ad to show and a bid to place
- auctioneer decides on the auction winner and price
- the winning ad is shown and conversion/click/impression is observable by the winning bidder

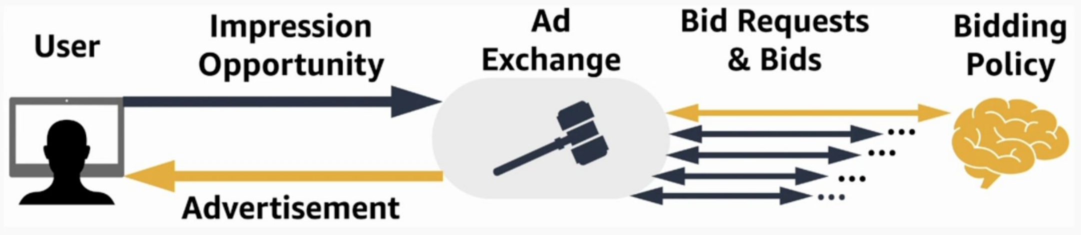

Auction Gym

- Policy-based and doubly robust formulation of bidding problem
  - Interactive and reactive nature of the repeated auction mechanism.
- Bandit-based “learning to bid”
  - Ad allocation problem
  - Bidding problem.

Auction Gym

- **Simulating Auctions (First/Second price auctions) to decide:**
  - who wins the auction
  - how much they will be charged
- **Simulating Bidders**
  - Every bidder has a private ad catalogue, private valuation for a given ads on conversion event. The ad-specific parameters are configurable.
- **Simulating Advertising Outcomes**
  - Simulate whether an allocation decision leads to a conversion event for the advertiser.

Off-policy estimation

- Choose a counterfactual (off-policy) estimator:
  - Given samples from π_0, what utility would I get from π?
- Learn the policy that maximizes this estimator: optimize π through gradient descent

- Value-based estimation (Direct Method)
  - Model winning probability function
  - high bias
- Policy-based estimation (IPS)
  - High variance
- Doubly robust estimation
  - Unbiased, lower variance

Case Study: Auction Gym

- **Model-based approach** stabilises quickly but suboptimally
  - Biased low-variance estimator.
- **Model-free importance sampling estimator** has high variance, and is able to improve upon the model-based estimator when sufficient learning steps are allowed.
  - The instability can lead to significant reductions in attainable welfare as it impacts training data collection for subsequent updates to the allocation model.
- **Doubly robust estimator** leads to improved surplus over all bidders participating in the auction
  - Lower variance than IPS.

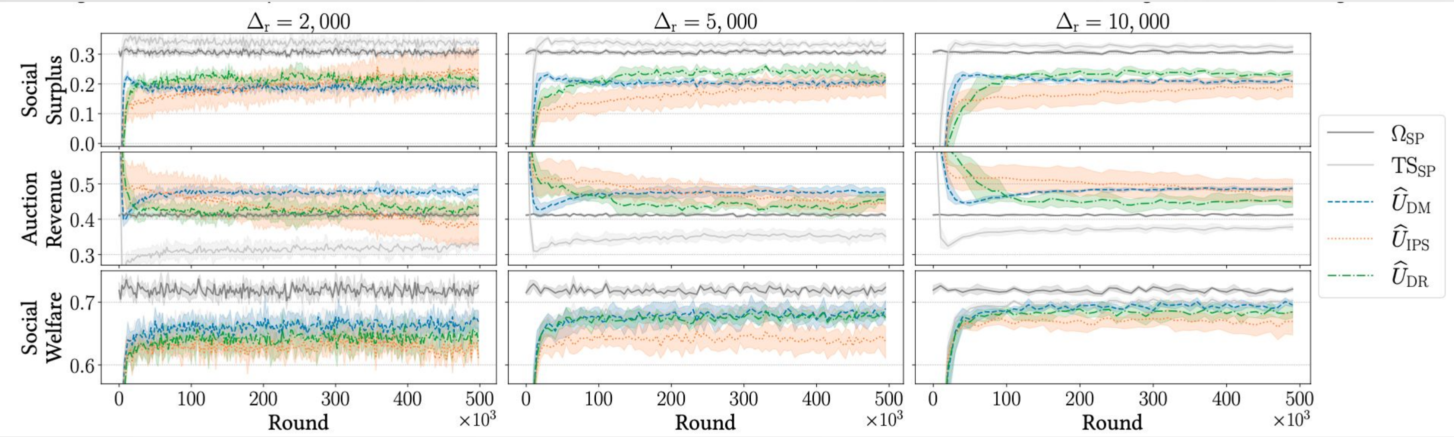

Challenges and Future Directions

Challenges

- Online vs offline parity
  - Tuning setup has significant deviations from existing policies/models in real traffic, yield large change in feature distributions.
- Increasingly large data size and search space
  - Data size grows aggressively including traffic volume, ads data and contextual data. Increasing complexity of the problem space need calibration on critical steps.

Future Directions

- Full reinforcement learning
  - Full reinforcement learning instantiations of the bidding problem, where current actions influence future states and a notion of planning can further improve bidder surplus
- Extend the simulation environment
  - Support advertiser budgets, multi-item and learnt auction mechanisms.

06

Summary & Future directions

Complete the loop of **pattern**, **prediction & inference**, **action**
Interpretation and **evaluation**
Generative AI and **LLM**

Summary of Agent-based Simulation for Modern Web Applications

Complete the loop

Pattern Recognition, Prediction & Inference, and Action are the three pillars of modern web applications and perhaps artificial intelligence (AI).

Agent-based simulation is able to methodically integrating and summarizing concepts and solutions developed throughout history for each component across various problems.

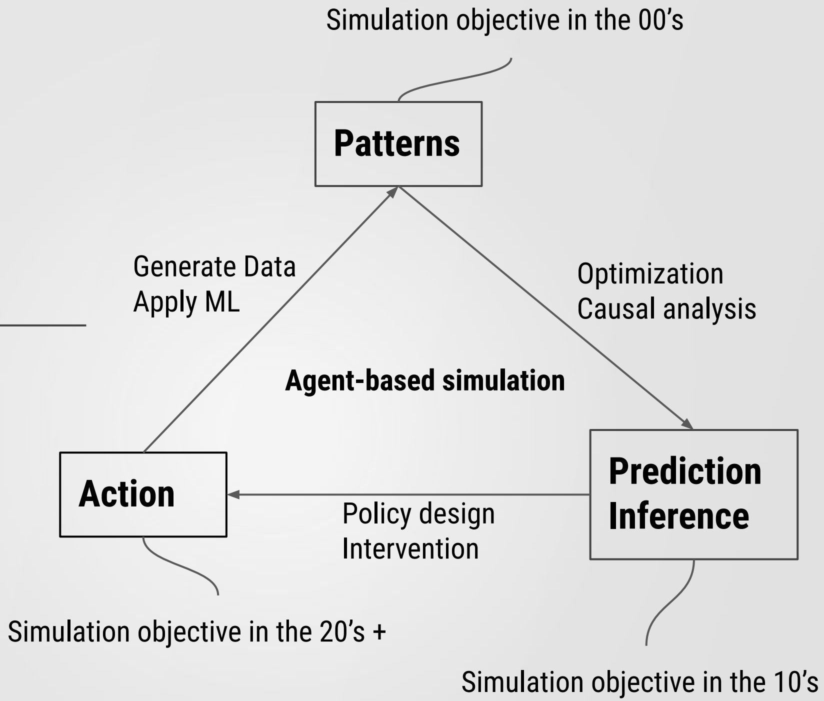

Result Interpretation and evaluation for Agent-based Simulation

An agent has to serve a purpose, how to build an agent with the right amount of description or detail remains an **art** than **science**.

Natural systems have **soft factors** that are difficult to quantify, calibrate, or even justify.

One must not make decisions on the basis of the **quantitative outcome** of an agent-based simulation that are designed purely at the **qualitative level**

Outcome correlation with real-world experiments and testings

Reality faithfulness

Complexity of the assumption and hypothesis supported by visible mechanisms from the interventions

Assumption complexity
Intervention manipulability

What's Next for Agent-based Simulation?

Employ such as generative adversarial network (**GAN**) as a feature (or initial states) agent to produce **close-to-real** distributions

Remark: considered as an alternative to the causal-structure-based simulation

Large language models pre-trained on whole-web data exhibit remarkable capacity for in-context learning, enabling them to effectively tackle a diverse array of complex tasks with human-like proficiency. With appropriate set up and prompt engineering, LLM can serve as both **task-specific and generic agents**.

Can we imagine a simulation consists entirely of GenAI agents?

Remark: although the directions presented in this slide have been explored to some extent, there remains much opportunities

Real-world observations could be an agent itself (e.g. reward-model agent), leading to the potentially more **powerful** & **targeted** semi-synthetic analysis

Remark: naturally connects to learning from partial feedback with bandit or RL

THANK YOU

Does anyone has any questions?

01 **Opening remark**
Motivation, Introduction & Scope

02 **Overview**
Agent-based Simulation for Modern Web Applications

03 **For Information Retrieval**
Web search engine
Conversation System

04 **For Recommender System**
Personalized Recommendation

05 **For Marketing and Advertising**
Bidding, Pricing, Ads Allocation

06 **Summary and future directions**
Landscape
Interpretation, evaluation
Generative AI, LLM

Slides and Recordings will be available on our webpage:
https://foundation4recsys.github.io/Tutorial-WWW23/

THANK YOU

Please feel free to contact us for follow-up and opportunities

daxu5180@gmail.com
boyang.emma@gmail.com
shuyuan.xu@rutgers.edu

Da Xu   Shuyuan Xu   Bo Yang

References

[1] Stavinova, E., Grigorievskiy, A., Volodkevich, A., Chunaev, P., Bochenina, K. and Bugaychenko, D., 2022. Synthetic Data-Based Simulators for Recommender Systems: A Survey. arXiv preprint arXiv:2206.11338.
[2] Ie, E., Hsu, C.W., Mladenov, M., Jain, V., Narvekar, S., Wang, J., Wu, R. and Boutilier, C., 2019. Recsim: A configurable simulation platform for recommender systems. arXiv preprint arXiv:1909.04847.
[3] Huang, J., Oosterhuis, H., De Rijke, M. and Van Hoof, H., 2020, September. Keeping dataset biases out of the simulation: A debiased simulator for reinforcement learning based recommender systems. In Proceedings of the 14th ACM Conference on Recommender Systems (pp. 190-199).
[4] Shi, J.C., Yu, Y., Da, Q., Chen, S.Y. and Zeng, A.X., 2019, July. Virtual-taobao: Virtualizing real-world online retail environment for reinforcement learning. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 33, No. 01, pp. 4902-4909).
[5] McInerney, J., Elahi, E., Basilico, J., Raimond, Y. and Jebara, T., 2021, September. Accordion: a trainable simulator for long-term interactive systems. In Proceedings of the 15th ACM Conference on Recommender Systems (pp. 102-113).
[6] Kalimeris, D., Bhagat, S., Kalyanaraman, S. and Weinsberg, U., 2021, August. Preference amplification in recommender systems. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (pp. 805-815).
[7] Yao, S., Halpern, Y., Thain, N., Wang, X., Lee, K., Prost, F., Chi, E.H., Chen, J. and Beutel, A., 2021. Measuring recommender system effects with simulated users. arXiv preprint arXiv:2101.04526.

References

[8] Zhang, Shuo, and Krisztian Balog. "Evaluating conversational recommender systems via user simulation." *Proceedings of the 26th acm sigkdd international conference on knowledge discovery & data mining*. 2020.
[9] Maxwell, David, and Leif Azzopardi. "Simulating interactive information retrieval: Simiir: A framework for the simulation of interaction." *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 2016.
[10] Schatzmann, Jost, et al. "Agenda-based user simulation for bootstrapping a POMDP dialogue system." *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*. 2007.
[11] Zerhoudi, Saber, et al. "The SimIIR 2.0 Framework: User Types, Markov Model-Based Interaction Simulation, and Advanced Query Generation." *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2022.
[12]Jeunen, Olivier, et al. "A probabilistic framework to learn auction mechanisms via gradient descent." (2023).

References

[13] Jeunen, Olivier, Sean Murphy, and Ben Allison. "Learning to bid with AuctionGym." (2022).
[14] Bayir, Murat Ali, et al. "Genie: An open box counterfactual policy estimator for optimizing sponsored search marketplace." Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining. 2019.
[15] Weinan Zhang, Shuai Yuan, and Jun Wang. Optimal real-time bidding for display advertising. In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, pages 1077–1086. ACM, 2014.
[16] Wang, Jun, Weinan Zhang, and Shuai Yuan. "Display advertising with real-time bidding (RTB) and behavioural targeting." Foundations and Trends in Information Retrieval 11.4-5 (2017): 297-435.