# CSC 84040: Advanced Data Mining *with AI Agents*

## Game Theory & Equilibrium in Data Mining: Coordination, Competition & Practical Control

By Saptarashmi Bandyopadhyay
Email: sbandyopadhyay@ccny.cuny.edu
Assistant Professor of Computer Science
City College of New York and Graduate Center at the City University of New York

Saptarashmi Bandyopadhyay

October 14, 2026 - CSC 84040

# From Static Datasets to Strategic Agent Loops

Traditional data mining assumes a passive, stationary data distribution, whereas real-world deployment operates in **continuous, agent-driven processing loops** where strategic actors adapt to the algorithm.

When multiple agents interact over shared graphs or data streams, learning becomes a game of **coordination and competition**: spammers manipulate link topologies, dominant groups skew empirical risk, and unconstrained clusters collapse.

Tonight we bridge **stationary spectral equilibria**, **min-max adversarial games**, **optimal transport polytopes**, and **closed-loop reinforcement control** to give data mining models practical control capabilities.

Subtarashmi Bandyopadhyay

# Lecture Agenda

**Adversarial Games & Web Equilibria**

01

Markov stationary distributions in PageRank, spider-trap taxation equilibria, spam-farm payoff amplification, TrustRank, and HITS mutual recursion.

**Min-Max Games & Bilevel Optimization**

02

Multi-agent feature coordination in BA-Cite, GroupDRO min-max worst-group games, adversarial venue invariance, and Stackelberg bilevel meta-learning.

**Optimal Transport & Team Dynamics**

03

Equipartition polytope equilibria in SwAV, Sinkhorn-Knopp scaling, SVGD attractive-repulsive particle games, and small-vs-large team disruption in SciSciNet.

**Practical Control & Embodied Agents**

04

Actor-Critic TD control loops, model-based world models, counterfactual intervention regularization, and hierarchical 3D scene-graph control.

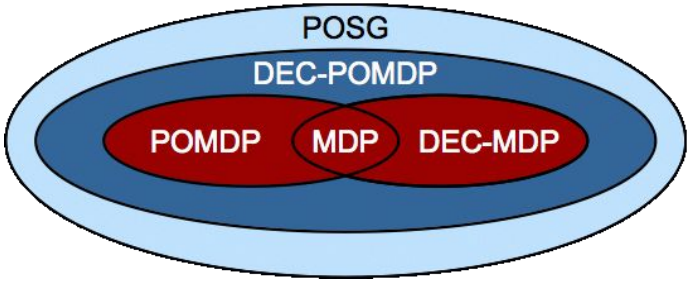
Saptarshmi Bandyopadhyay

# Part 0: Formalizing a “Game” & The Nash Equillibrium

Saptarashmi Bandyopadhyay

# What is the “Game” in “Game Theory”?

- **Single-Agent RL (MDPs):** The agent takes an action, the environment transitions to a new state, and gives a reward. The environment is stationary.
- **Multi-Agent RL (Games):** The "environment" now contains other thinking, adapting agents.
- **The Core Issue:** You cannot optimize your reward without predicting what the other agents are going to do. And they are simultaneously trying to predict what *you* are going to do!




Source: http://rbr.cs.umass.edu/camato/decpomdp/overview.html

Sayantani Bandyopadhyay

# Normal-Form Game

To solve these interactions, we formalize them mathematically as a **Normal-Form Game**. A game is defined by a tuple: $\langle N, A, u \rangle$

**Players ($N$):** A finite set of agents $N = \{1, 2, ..., n\}$ .

**Action Space ($A$):** $A = A_1 \times A_2 \times \ldots \times A_n$ . This represents all possible joint actions. A specific joint action is $a = (a_1, ..., a_n)$.

**Utility Function ($u$):** $u_i(a)$ gives the payoff for player $i$ when the joint action $a$ is played.

**Notation Trick:** We often write a joint action from the player $i$ perspective as $(a_i, a_{-i})$ , where $a_{-i}$ means "the actions of absolutely everyone else except me".

Soptarashmi Bandyopadhyay

# Payoff Matrices

For 2-player games, we visualize $\langle N, A, u \rangle$ using a **Payoff Matrix**.

**Row Player (Player 1)** chooses the row. **Column Player (Player 2)** chooses the column.

The cell contains the tuple of utilities: $(u_1, u_2)$ .

**Example:** If Row chooses "Top" and Column chooses "Left", we look at the top-left cell. If it says $(3, -1)$, Player 1 gets a reward of 3, and Player 2 gets a reward of -1.

| | **Left** | **Right** |
|---|---|---|
| **Top** | (3, -1) | (0, 0) |
| **Bottom** | (-2, 2) | (1, 1) |

Sabtarashmi Bandyopadhyay

# Pure Competition (Zero-Sum Games)

**The Paradigm:** What is good for me is strictly bad for you. Our goals are perfectly misaligned.

**The Math:** For any joint action $a$, the sum of all utilities is zero.

$$\sum_{i \in N} u_i(a) = 0$$

*(Or $u_1(a) = -u_2(a)$ for 2 players).*

Saptarashmi Bandyopadhyay

# Coordination and Mixed-Motives

**Pure Coordination (Common Payoff):** Agents share the exact same utility function $u_1(a) = u_2(a)$ .The only challenge is communicating or agreeing on the best joint action. (e.g., two robots carrying a heavy table).

**Mixed-Motive (General Sum):** The most complex and realistic paradigm. Agents have independent goals, but can benefit from temporary cooperation.

**The Classic Failure - Prisoner's Dilemma:** Even if mutual cooperation yields a high reward for both $(3, 3)$, individual rationality forces both agents to defect, leading to a terrible outcome $(1, 1)$ . How do we program AI to avoid this trap?

Saptarashmi Bandyopadhyay

# The Nash Equilibrium

Saptarashmi Bandyopadhyay

# Nash Equilibrium

How do we define a "solved" game? We use the **Nash Equilibrium**.

**Definition:** A joint strategy profile $\pi^*$ is a Nash Equilibrium if no player has an incentive to *unilaterally* deviate from it.

**The Mathematical Condition:** For all players $i$, and all possible alternative actions $a_i$:

$$u_i(\pi_i^*, \pi_{-i}^*) \geq u_i(a_i, \pi_{-i}^*)$$

**Intuition:** If you know exactly what your opponent is going to do, and you still wouldn't change your move, you are in a Nash Equilibrium.

Saptarashmi Bandyopadhyay

# Pure vs. Mixed Strategies

**Pure Strategy:** The agent chooses a single action deterministically (100% chance to play "Rock").

**The Problem:** Many games (like Rock-Paper-Scissors) *do not have* a Pure Strategy Nash Equilibrium. If you always play Rock, I will just play Paper.

**Mixed Strategy ( $\pi$ ):** The agent plays a probability distribution over its available actions.

**Nash's Theorem:** *Every* game with a finite number of players and actions has at least one Mixed Strategy Nash Equilibrium.

Saptarashmi Bandyopadhyay

# Part 1: Adversarial Competition & Equilibria in Web Mining

Saptarashmi Bandyopadhyay

# Web Search as a Non-Cooperative Repeated Game

Web search is an ongoing **strategic game** between search engines seeking to rank genuinely useful pages and commercial spammers seeking to manipulate rankings for profit.

In the first round of this game, early search engines ranked pages via **inverted index term frequencies**, prompting spammers to deploy **term spam**—stuffing thousands of invisible background-colored keywords into pages.

Google shifted the game's equilibrium via two mechanisms: simulating **random surfer congregation (PageRank)** and judging a page by **anchor text in incoming links** controlled by independent third parties.

Soutarashmi Bandyopadhyay

# PageRank as a Markov Chain Stationary Equilibrium

We model the Web as a directed graph of n pages with **transition matrix M**, where entry $m_{ij} = 1/k$ if page j has k outgoing arcs including an arc to page i, and $m_{ij} = 0$ otherwise.

Initializing a random surfer uniformly across pages via $v_0 = (1/n)e$ and iterating $v_{t+1} = M v_t$ converges to the **stationary equilibrium distribution v** satisfying:
$v = M v$
where v is the principal eigenvector of column-stochastic matrix M associated with eigenvalue $\lambda = 1$.

Convergence to a unique, non-zero equilibrium strictly requires two topological conditions: the graph must be **strongly connected** and contain **no dead ends** (nodes with zero out-links).

Sooptarashmi Bandyopadhyay

# Breaking Sinks: Taxation as Equilibrium Control

Real web graphs exhibit a **bowtie structure** (in-component, strongly connected core, out-component, tendrils, and tubes) that violates strong connectivity and destroys naive Markov equilibrium.

**Dead ends** make M substochastic so probability mass drains to 0, while **spider traps** (closed subsets with no outgoing arcs) act as absorbing sinks that monopolize 100% of the surfer distribution.

We restore a unique stationary equilibrium via **taxation (recycling surfers)** with damping control parameter $\beta \in [0.8, 0.9]$:
$v' = \beta M v + (1 - \beta) e / n$
where e is an n-dimensional vector of ones, taxing each surfer's rank by $(1 - \beta)$ at every step and redistributing it uniformly.

Saptarashmi Bandyopadhyay

# Adversarial Response: Spam Farm Equilibrium Payoff

To defeat PageRank, link spammers partition the Web into **inaccessible pages**, **accessible pages** (blogs/forums where spammers drop comment links), and **owned pages** arranged in a star-like **spam farm**.

A spammer surrounds target page t with m supporting pages that link exclusively back to t; each supporting page reaches equilibrium rank **βy/m + (1 - β)/n**, where y is the unknown PageRank of target t and n is total Web size.

Summing external inflow x from accessible pages and recycled support rank yields the **closed-form equilibrium payoff** for target page t:
$y = x / (1 - \beta^2) + c \cdot (m / n)$, where $c = \beta / (1 + \beta)$
At β = 0.85, the farm amplifies external rank x by **$1/(1 - \beta^2) = 360\%$** while capturing **$c = 46\%$** of its web share m/n.

Soutarashmi Bandyopadhyay

# Defender Counter-Strategies: TrustRank & Spam Mass

Spammers can even form **cooperative coalitions** by linking target pages across separate spam farms to mutually amplify y, making purely structural motif detection an endless arms race.

**TrustRank** shifts the game by restricting the teleport set S to verified **trustworthy seed pages** (human-vetted top nodes or controlled .edu / .gov domains):
$v_{trust} = \beta M v_{trust} + (1 - \beta) e_S / |S|$
where $e_S$ has 1s for trusted nodes in S and 0s elsewhere, starving supporting spam pages of their $(1 - \beta)/n$ tax subsidy.

The search engine then computes **Spam Mass** for every node p with PageRank r and TrustRank t:
$SpamMass(p) = (r - t) / r$
Values near 1.0 expose nodes whose authority is overwhelmingly manufactured by link spam, triggering automated demotion.

Saptarshi Bandyopadhyay

# HITS: Coupled Hub-Authority Fixed-Point Equilibrium

Rather than a single importance score, Kleinberg's **HITS algorithm** decomposes graph value into two coupled roles: **Authorities** (pages providing core information) and **Hubs** (pages curating links to authorities).

Given n × n link adjacency matrix L (where $L_{ij} = 1$ if i links to j), hub vector h and authority vector a evolve via **mutually recursive fixed-point equations**:
$h = \lambda L a, \quad a = \mu L^T h \quad \Rightarrow \quad h = \lambda \mu (L L^T) h, \quad a = \lambda \mu (L^T L) a$
where λ and μ are scaling constants that normalize the maximum vector component to 1 at each step.

Because $L L^T$ (bibliographic coupling) and $L^T L$ (co-citation) are symmetric matrices, alternating updates converge to a **coupled eigenvector equilibrium** without requiring taxation for dead ends or spider traps.

Saptarshi Bandyopadhyay

# Part 2: Min-Max Games, GroupDRO & Bilevel Optimization

Saptarashmi Bandyopadhyay

# Cooperative Multi-Agent Feature Extraction (BA-Cite)

In cold-start citation prediction for newborn papers, early cascade signals are absent and single models overfit **explicit shortcut proxies** like venue prestige (V).

**BA-Cite** coordinates **six specialized autonomous agents** that mine metadata, GitHub, Wikipedia, and LLMs to construct a unified node attribute vector **f_p = [A, V, R, C, H, Q]**.

Here, the coordinated agents extract Role-Aware Author Reputation ($A_1$, $A_2$, $A_3$), Venue Prestige (V), Reproducibility code/data ($R \in \{0, 1\}$), Collaboration Network (C), Topic Hotness (H), and LLM Text Quality (Q), injecting f_p into a **Heterogeneous Graph Transformer (HGT)**.

Soprarashmi Bandyopadhyay

# Causal Mediation & Adversarial Venue-Invariance Game

To prevent explicit venue prestige V from dominating predictions, BA-Cite enforces the **two-stage mediation pathway V → E → Y** separating early exposure E from future citations Y:
Stage A: $\hat{E} = g_{\phi}([z_p; f_p^{(+V)}])$, Stage B: $\hat{Y} = f_{\theta}([z_p^{(-V)}; f_p^{(-V)}; \hat{E}])$
strictly reducing shortcut mutual information **$I(V; Y \mid \hat{E}) < I(V; Y)$**.

To purge residual venue leakage from Stage-B representations $s_p$, the encoder $f_\theta$ plays a **two-player zero-sum min-max game** against a venue discriminator $D_\psi$:
$\min_\theta \max_\psi \mathbb{E}_{p \sim \mathcal{D}}[\log D_\psi(v_p \mid f_\theta(s_p))]$

While discriminator $D_\psi$ maximizes its accuracy in predicting venue tier $v_p$, encoder $f_\theta$ updates $\theta$ to fool $D_\psi$, converging to **venue-invariant latent representations**.

Saptarshi Bandyopadhyay

# GroupDRO: Min-Max Equilibrium Across Environments

Standard Empirical Risk Minimization (ERM) minimizes average loss, causing heavy-tailed citation models to sacrifice accuracy on **underrepresented, lowly cited papers**.

**Group Distributionally Robust Optimization (GroupDRO)** casts training as a zero-sum game between model parameters θ and an adversarial environment player weighting groups $\mathcal{E} = \{\text{low}, \text{high}\}$:

$$\min_\theta \max_{w \in \Delta^2} \sum_{e \in \mathcal{E}} w_e \mathcal{L}_e(\theta), \quad \text{where} \quad \Delta^2 = \{ w \geq 0 \mid w_{low} + w_{high} = 1 \}$$

At each step t, the adversary amplifies weights on high-loss environments via **exponentiated gradient ascent**:

$$\tilde{w}_e \propto w_e^{(t)} \exp\left( \alpha \, (\mathcal{L}_e - \bar{L}) / (\sigma_{\mathcal{L}} + \varepsilon) \right), \quad w_e^{(t+1)} = \text{clip}(\tilde{w}_e, w_{min}, w_{max}) / \sum_{e'} \hat{w}_{e'}$$

where $\bar{L}$ and $\sigma_{\mathcal{L}}$ are the mean and standard deviation of group risks and $\alpha = 0.1$ is the step size.

Saptarashmi Bandyopadhyay

# Bilevel Meta-Learning as a Stackelberg Game

In low-resource learning, **bilevel meta-optimization** structures training as a hierarchical **leader-follower (Stackelberg) game** between meta-parameters $\phi$ (leader) and task-specific parameters $\theta$ (follower).

Across M source tasks $D_{source} = \{(D_{source}^{train}, D_{source}^{val})^{(i)}\}$, the outer-level optimization is constrained by the inner-level optimal response $\theta^{*(i)}(\phi)$:

$$\phi^* = \text{argmin}_\phi \sum_{i=1}^{M} \mathcal{L}^{meta}(\theta^{*(i)}(\phi), \phi, D_{source}^{val}(i)) \quad \text{s.t.} \quad \theta^{*(i)}(\phi) = \text{argmin}_\theta \mathcal{L}^{task}(\theta, \phi, D_{source}^{train}(i))$$

During the inner phase, follower $\theta$ adapts to individual low-resource tasks holding $\phi$ fixed; during the outer phase, leader $\phi$ updates across validation tasks to extract **transferable equilibrium meta-knowledge $\phi^*$**.

Saptarashmi Bandyopadhyay

# Part 3: Optimal Transport, Particle Games & Team Dynamics

Saptarashmi Bandyopadhyay

# SwAV: Optimal Transport & Equipartition Equilibrium

In online self-supervised visual clustering, unconstrained optimization suffers from **mode collapse**, where the encoder maps every image in a batch of size B to the exact same prototype.

**SwAV** enforces a balanced equilibrium by constraining code matrix $Q = [q_1, ..., q_B] \in \mathbb{R}_+^{K \times B}$ to the **transportation polytope** $\mathcal{Q}$ over K prototypes $C = [c_1, ..., c_K]$:

$$\max_{Q \in \mathcal{Q}} \operatorname{Tr}(Q^T C^T Z) + \varepsilon H(Q), \quad \text{where} \quad \mathcal{Q} = \{ Q \in \mathbb{R}_+^{K \times B} \mid Q 1_B = (1/K) 1_K, \; Q^T 1_K = (1/B) 1_B \}$$

Entropy regularization $H(Q) = -\sum_{ij} Q_{ij} \log Q_{ij}$ (with small $\varepsilon = 0.05$) prevents uniform collapse while **3 Sinkhorn-Knopp iterations** yield the unique soft equilibrium $\mathbf{Q^* = Diag(u) \exp(C^T Z / \varepsilon) Diag(v)}$.

Soutarashmi Bandyopadhyay

# Stein Variational Gradient Descent (SVGD) Particle Game

To approximate an intractable target posterior p(x) without costly MCMC sampling, **SVGD** evolves a swarm of m interacting particles $\{x_i^0\}_{i=1}^m$ through deterministic incremental transforms T(x) = x + ε φ(x).

In the unit ball of a Reproducing Kernel Hilbert Space (RKHS) with kernel k(x, x'), the steepest descent direction minimizing KL(q || p) reduces to the **Kernelized Stein Discrepancy (KSD)** update:
$x_i^{(t+1)} \leftarrow x_i^t + (\varepsilon_t / m) \sum_{\square=1}^m [\, k(x_\square^t, x_i^t)\ \nabla_{x_\square^t} \log p(x_\square^t) + \nabla_{x_\square^t} k(x_\square^t, x_i^t) \,]$

Particle equilibrium emerges from two competing forces: the first term **drives particles toward high-probability modes** of p(x), while the repulsive gradient **$\nabla k(x_\square^t, x_i^t)$ pushes particles apart** to prevent mode collapse.

Soutarashmi Bandyopadhyay

# Harmonic Graph Equilibrium & Spin-Glass States (KnIT)

In automated hypothesis generation (KnIT), candidate entities form a weighted similarity network where known biological annotations $y \in \{0, 1\}^n$ diffuse globally to reach a **harmonic label equilibrium $f \in \mathbb{R}^n$**.

Drawing on risk minimization through **accessible spin-glass ground states**, f minimizes the sum of anchor loss and Laplacian quadratic energy:
$\min_f \{ (f - y)^T(f - y) + \mu f^T L f \} \Rightarrow f^* = (I + \mu L)^{-1} y$
where $L = D - A$ is the graph Laplacian and $\mu = 1 / \max_i(\sum_k |L_i k|)$ guarantees strict convexity.

At equilibrium $f^*$, every unlabeled node settles at a potential balancing its local neighbors' states against global conservation of initial evidence y, enabling **0.840 AUC retrospective kinase discovery**.

Sagnik Bandyopadhyay

# Team Coordination Trade-offs in Science (SciSciNet)

Mining **134,129,188 papers in SciSciNet** reveals a steady multi-decade expansion in average research **team size** across all 19 top-level scientific fields from 1950 to 2020.

However, evaluating the **Disruption Index D = ($n_i$ - $n_j$) / ($n_i$ + $n_j$ + $n_k$)**—where $n_i$ cites only the focal paper, $n_j$ cites both focal and its references, and $n_k$ cites only references—exposes a fundamental **coordination trade-off**.

As team size grows, **citation impact ($C_{10}$) rises** but **disruption percentile drops sharply**: **small teams disrupt** science by overturning paradigms (high D), whereas **large teams coordinate to develop** existing frontiers.

Sautarashmi Bandyopadhyay

# Part 4: Practical Control Capabilities & Embodied Agents

Saptarashmi Bandyopadhyay

# Closed-Loop Control: Actor-Critic & TD Error

Giving data mining models **practical control capabilities** requires transitioning from passive prediction to a closed-loop **Actor-Critic reinforcement learning architecture**.

The **Actor (policy π)** selects control actions from observed states, while the **Critic (value function V)** evaluates state or state-action expected returns from environmental rewards.

The **Temporal-Difference (TD) error** measures the discrepancy between predicted and actual returns, providing a unified gradient signal that updates both the Critic's value estimate and the Actor's control policy.

Saptarashmi Bandyopadhyay

# World Models & Distributionally Robust Control

**Model-Based World Models:** To minimize costly real-world interactions, agents learn latent environment dynamics (**Dyna, PlaNet, MBPO, PETS, Dreamer, EfficientZero**) to execute virtual planning rollouts.

**Prior-Guided & Demonstration Control:** Algorithms accelerate convergence in sparse-reward regimes using **Deep Q-Learning from Demonstrations (DQfD)**, causal MDP priors, and **LLM action-distribution priors**.

**Distributionally Robust Offline RL:** Under computational and dataset constraints, minimax optimal offline RL controllers and **hyperspherical normalization** guarantee stable policy updates under environmental shift.

Soutarashmi Bandyopadhyay

# Actionable Control via Counterfactual Regularization

Predictive models deployed for decision support must respond sensibly to **controllable interventions** rather than exploiting spurious correlations.

In BA-Cite, the regularization head performs **what-if interventions** on actionable factors v (such as toggling code reproducibility R: 0 → 1), recomputing exposure $\hat{E}^{(v\uparrow)}$ to measure counterfactual effect:
$\Delta_v(p) = \hat{Y}(s_p^{(v\uparrow)}) - \hat{Y}(s_p)$, where $s_p^{(v\uparrow)} = [z_p; f_p^{(-V)}$ with $v \leftarrow$ high; $\hat{E}^{(v\uparrow)}]$

Control stability is enforced by combining **directional monotonicity** and **smoothness penalties**:
$\mathcal{L}_{mono}^{(v)} = \mathbb{E}[\max(0, -t_v \Delta_v(p)) \cdot 1\{v(p) < \tau_v\}]$, $\mathcal{L}_{smooth}^{(v)} = \mathbb{E}[\Delta_v(p)^2]$
guaranteeing that improving reproducibility R or text quality Q never degrades predicted impact.

Sudarshini Bandyopadhyay

# Hierarchical RL & Spatial Scene-Graph Control

**Hierarchical Reinforcement Learning (HRL)** decomposes long-horizon control into multi-scale modules—**perception, planning, execution, and feedback**—where high-level policies emit sub-goals or prompts to low-level controllers.

In embodied robotics, **Kimera** structures physical environments into a **4-layer spatial hierarchical scene graph**: (1) Metric-Semantic Mesh, (2) Objects and Agents, (3) Places and Structures, and (4) Rooms.

In open-world tasks, **DECKARD** uses few-shot LLMs to hypothesize an **Abstract World Model** for sub-goal planning, while **SAM** executes 11-step interactive point-and-mask control in **~50 ms on CPU** via a 3-mask winner-take-all head.

Soutarashmi Bandyopadhyay

# Game Theory, Equilibrium & Control Across Paradigms

| Mining / AI System | Interacting Players / Forces | Equilibrium / Objective | Control / Stabilization Mechanism |
|---|---|---|---|
| **PageRank vs. Link Spam** | Search Engine vs. Spam Farm | $v = \beta M v + (1-\beta)e/n$; $y = x/(1-\beta^2)+c(m/n)$ | Taxation β, TrustRank teleport set S, Spam Mass |
| **BA-Cite Graph Learning** | Encoder $f_\theta$ vs. Group/Venue Adversary | $\min_\theta \max_{\{w \in \Delta^2\}} \sum w_e \mathcal{L}_e(\theta)$ | Two-stage V→E→Y, GroupDRO, Monotonicity $\mathcal{L}_{mono}$ |
| **SwAV Visual Clustering** | Similarity $Tr(Q^T C^T Z)$ vs. Collapse | $Q^* = Diag(u)\ \exp(C^T Z/\varepsilon)\ Diag(v)$ | Equipartition polytope $\mathcal{Q}$ & 3-step Sinkhorn-Knopp |
| **SVGD & KnIT Diffusion** | Attraction / Anchor vs. Repulsion / L | $f^* = (I + \mu L)^{-1} y$; KSD steepest descent | RKHS repulsive kernel $\nabla k$; Laplacian bound μ |
| **Embodied HRL & SAM** | High-Level Planner vs. Low-Level Actor | TD Value Equilibrium & Winner-Take-All | World models (Dreamer/DECKARD) & 50ms mask loop |

# Summary & Preparation for Next Week

**Equilibrium as a Unifying Principle:** Whether computing **PageRank/HITS eigenvectors**, solving **SwAV Sinkhorn-Knopp optimal transport**, or diffusing **KnIT Laplacian potentials**, stable data mining relies on well-conditioned fixed-point equilibria.

**Robustness via Min-Max Games:** Formulating learning as an adversarial game—via **TrustRank**, **GroupDRO worst-group reweighting**, **venue-invariance discriminators**, and **bilevel meta-learning**—prevents shortcut overfitting and long-tail collapse.

**Next Lecture (Oct 21, 2026): Evolutionary Learning, Mean-Field Learning, Self-play for Self-Improvement, Continual Learning & Decentralized AI**—exploring real-time multi-agent communication and scaling paradigms.

Sooptarashmi Bandyopadhyay

# Questions?

Saptarashmi Bandyopadhyay