

Novelty Assessment Report

Paper: Online Minimization of Polarization and Disagreement via Low-Rank Matrix Bandits

PDF URL: <https://openreview.net/pdf?id=nwkiK8vNd1>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

We study the problem of minimizing polarization and disagreement in the Friedkin-Johnsen opinion dynamics model under incomplete information. Unlike prior work that assumes a static setting with full knowledge of users' innate opinions, we address the more realistic online setting where innate opinions are unknown and must be learned through sequential observations. This novel setting, which naturally mirrors periodic interventions on social media platforms, is formulated as a regret minimization problem, establishing a key connection between algorithmic interventions on social media platforms and theory of multi-armed bandits. In our formulation, a learner observes only a scalar feedback of the overall polarization and disagreement after an intervention. For this novel bandit problem, we propose a two-stage algorithm based on low-rank matrix bandits. The algorithm first performs subspace estimation to identify an underlying low-dimensional structure, and then employs a linear bandit algorithm within the compact dimensional representation derived from the estimated subspace. We prove that our algorithm achieves an $\tilde{O}(\sqrt{T})$ cumulative regret over any time horizon T . Empirical results validate that our algorithm significantly outperforms a linear bandit baseline in terms of both cumulative regret and running time.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Online Minimization of Polarization and Disagreement in Opinion Dynamics**

A total of **33 papers** were analyzed and organized into a taxonomy with **10 categories**.

Taxonomy Overview

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

- **Algorithmic Intervention Strategies for Polarization Mitigation**
- **Polarization Mechanisms and Modeling**
- **Domain-Specific Applications and Extensions**

Complete Taxonomy Tree

- Online Minimization of Polarization and Disagreement in Opinion Dynamics Survey Taxonomy
- Algorithmic Intervention Strategies for Polarization Mitigation
 - Network Structure Modification Approaches
 - Link Recommendation and Edge Perturbation (5 papers)
 - [2] Towards consensus: Reducing polarization by perturbing social networks (Mikláš Z. Rájcz, 2023) [View paper](#)
 - [7] Minimizing polarization and disagreement in social networks via link recommendation (Liwang Zhu, 2021) [View paper](#)
 - [18] Mitigating opinion polarization in social networks using adversarial attacks (Michinori Ninomiya, 2024) [View paper](#)
 - [24] Reducing polarization in social networks with adversarial opinion perturbations (Lan Zhang, 2025) [View paper](#)
 - [25] Steering polarization toward consensus in signed majority-vote opinion models (Masaki Chujyo, 2025) [View paper](#)
 - Optimization Under Incomplete Information ★ (2 papers)
 - [0] Online Minimization of Polarization and Disagreement via Low-Rank Matrix Bandits (Anon et al., 2026) [View paper](#)
 - [4] Minimizing Polarization and Disagreement in the Friedkin-Johnsen Model with Unknown Innate Opinions (Cinus, 2025) [View paper](#)
 - Content and Recommendation System Interventions
 - Timeline and Feed Rebalancing (2 papers)
 - [8] Rebalancing social feed to minimize polarization and disagreement (Federico Cinus, 2023) [View paper](#)
 - [14] Modeling the Impact of Timeline Algorithms on Opinion Dynamics Using Low-rank Updates (Tianyi Zhou, 2024) [View paper](#)
 - Adaptive Recommender System Design (3 papers)
 - [10] Network-aware recommender system via online feedback optimization (Sanjay Chandrasekaran, 2025) [View paper](#)
 - [12] Optimizing Social Network Interventions via Hypergradient-Based Recommender System Design (Kāhne, 2025) [View paper](#)
 - [29] Recommender System Design via Online Feedback Optimization (S Chandrasekaran, 2024) [View paper](#)
 - Opinion-Based Intervention Methods (4 papers)
 - [5] How to Mitigate Disagreement and Polarization in Opinion Formation Processes on Social Networks (V Auletta, 2024) [View paper](#)
 - [9] Agent-based modelling of polarized news and opinion dynamics in social networks: a guidance-oriented approach (Shan Liu, 2024) [View paper](#)
 - [15] Control of kinetic opinion dynamics in popularity-adaptive social networks (Giacomo Albi, 2025) [View paper](#)
 - [30] Finite-Time Elimination of Disagreement of Opinion Dynamics via Covert Noise (Wei Su, 2022) [View paper](#)
- Polarization Mechanisms and Modeling
 - Algorithmic and Platform-Driven Polarization (6 papers)

- [1] Link recommendation algorithms and dynamics of polarization in online social networks (Fernando P. Santos, 2021) [View paper](#)
- [3] The algorithmic Management of Polarization and Violence on social media (Jonathan Stray, 2023) [View paper](#)
- [13] Analyzing the impact of filter bubbles on social network polarization (Uthsav Chitra, 2020) [View paper](#)
- [23] Understanding filter bubbles and polarization in social networks (Chitra, 2019) [View paper](#)
- [28] Harmonizing vs Polarizing Platform Influence Functions (Mallmann-Trenn, 2024) [View paper](#)
- [33] Modeling how social network algorithms can influence opinion polarization (Henrique Ferraz de Arruda, 2022) [View paper](#)
- Social Influence and Opinion Evolution Dynamics (5 papers)
- [16] Polarization on social media: Micro-level evidence and macro-level implications (Marijn A. Keijzer, 2024) [View paper](#)
- [17] Opinion amplification causes extreme polarization in social networks (Soo Ling Lim, 2022) [View paper](#)
- [22] MTOS: A LLM-Driven Multi-topic Opinion Simulation Framework for Exploring Echo Chamber Dynamics (Zuo, 2025) [View paper](#)
- [26] Time of influence as a decision variable in game-theoretic models of opinion dynamics in social groups (Jingjing, 2024) [View paper](#)
- [27] Opinion spreading and evolution in complex social networks: a simulation study (Yujie, 2024) [View paper](#)
- Agent-Based and Computational Simulation Studies (4 papers)
- [11] Modeling Disinformation Spread in Social Networks: Phase Transitions and Mean-Field Analysis (Spyridon Evangelatos, 2025) [View paper](#)
- [20] Agents of Discord: Modeling the Impact of Political Bots on Opinion Polarization in Social Networks (Hsuan-Wei Lee, 2024) [View paper](#)
- [21] Decoding Echo Chambers: LLM-Powered Simulations Revealing Polarization in Social Networks (Wang Chen-xi, 2024) [View paper](#)
- [32] Mitigating Influence of Disinformation Propagation Using Uncertainty-Based Opinion Interactions (Zhen Guo, 2022) [View paper](#)
- Domain-Specific Applications and Extensions
 - Public Health and Vaccine Hesitancy (1 papers)
 - [6] Modelling the mitigation of anti-vaccine opinion propagation to suppress epidemic spread: A computational approach (Sarah Alahmadi, 2024) [View paper](#)
 - Two-Party and Structured Opinion Systems (2 papers)
 - [19] Opinion de-polarization of social networks with GNNs (Mylonas Konstantinos, 2024) [View paper](#)
 - [31] Disagreement and Polarization in Two-Party Social Networks (Yuhao Yi, 2022) [View paper](#)

Narrative

Core task: online minimization of polarization and disagreement in opinion dynamics. The field addresses how to algorithmically reduce polarization and foster consensus in networked populations where opinions evolve over time. The taxonomy organizes research into three main branches: Algorithmic Intervention Strategies for Polarization Mitigation, which explores how platforms can actively modify network structures or content exposure to steer opinions toward agreement; Polarization Mechanisms and Modeling, which investigates the underlying dynamics that drive opinion fragmentation and echo chambers; and Domain-Specific Applications and Extensions, which applies these ideas to concrete settings such as social media feeds, political discourse, and misinformation spread. Within the intervention branch, some works focus on link recommendation and network rewiring (e.g., Link Recommendation Polarization[1], Consensus via Network Perturbation[2]), while others tackle content curation and feed design (e.g., Rebalancing Social Feed[8], Timeline Algorithms Low Rank[14]). The modeling branch examines phenomena like filter bubbles (Filter Bubbles Impact[13]) and adversarial manipulation (Adversarial Opinion Perturbations[24]), and the applications branch includes studies on bot-driven polarization (Political Bots Polarization[20]) and domain-specific interventions (Anti Vaccine Mitigation[6]).

A particularly active line of work centers on optimization under incomplete information, where the platform must learn which interventions reduce polarization without full knowledge of user opinions or network structure. Online Polarization Bandits[0] exemplifies this direction by framing the problem as a bandit task in which the platform sequentially selects edges to add while observing noisy feedback about polarization levels. This approach contrasts with works like Friedkin Johnsen Unknown Opinions[4], which also addresses unknown opinions but focuses on different inference or control mechanisms, and Mitigate Disagreement Networks[5], which may emphasize batch or offline optimization. A recurring theme across these studies is the trade-off between exploration (learning the network state) and exploitation (acting on current estimates), as well as the challenge of defining and measuring polarization in dynamic, partially observable environments. The original paper sits squarely within the network structure modification cluster, distinguished by its online learning perspective and bandit formulation.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Minimizing Polarization and Disagreement in the Friedkin-Johnsen Model with Unknown Innate Opinions

Authors: Cinus, Federico, Miyauchi, Atsushi, Federico Cinus, et al. (11 authors total) | **Year/Venue:** 2025 • International Joint Conference on Artificial Intelligence | **URL:** [View paper](#)

Abstract

The bulk of the literature on opinion optimization in social networks adopts the Friedkin-Johnsen (FJ) opinion dynamics model, in which the innate opinions of all nodes are known: this is an unrealistic assumption.

In this paper, we study opinion optimization under the FJ model without the full knowledge of innate opinions. Specifically, we borrow from the literature a series of objective functions, aimed at minimizing polarization and/or disagreement, and we tackle the budgeted optimization ...

Relationship Analysis

Both papers address polarization minimization in the Friedkin-Johnsen model under incomplete information about innate opinions, sharing the same taxonomy category of optimization under uncertainty. The original paper formulates an online learning problem using multi-armed bandits where innate opinions are learned through sequential noisy observations of scalar feedback, achieving regret bounds via low-rank matrix bandit techniques. The candidate paper takes a different approach by allowing limited queries of innate opinions, then reconstructing the full opinion vector to perform optimization, focusing on budgeted query selection and error propagation analysis rather than sequential online learning.

Contributions Analysis

Overall novelty summary. The paper contributes an online learning framework for polarization minimization in the Friedkin-Johnsen model, formulating the problem as regret minimization under incomplete information. It resides in the 'Optimization Under Incomplete

Information' leaf, which contains only two papers total (including this one). This leaf sits within the broader 'Network Structure Modification Approaches' branch, which encompasses five papers on link recommendation and edge perturbation. The sparse population of this specific leaf suggests the online bandit formulation for polarization reduction represents a relatively underexplored research direction within the broader intervention literature.

The taxonomy reveals neighboring work in 'Link Recommendation and Edge Perturbation' (five papers) and 'Content and Recommendation System Interventions' (five papers across two sub-leaves). The sibling paper in the same leaf addresses unknown opinions but appears to focus on different inference or control mechanisms rather than sequential bandit optimization. The broader 'Network Structure Modification Approaches' branch excludes content-based interventions and opinion-based methods, positioning this work firmly within topology-modification strategies. The taxonomy structure indicates that while network intervention is well-studied, the online learning perspective with incomplete information remains a niche area.

Among 29 candidates examined, the online regret formulation (Contribution 1) shows no clear refutation across 10 candidates, suggesting novelty in framing polarization reduction as a bandit problem. The two-stage algorithm with subspace estimation (Contribution 2) examined 9 candidates and found 5 potentially refutable, indicating substantial prior work on dimensionality reduction techniques in related bandit settings. The theoretical regret bound (Contribution 3) examined 10 candidates with no refutations, though this may reflect the specific combination of problem structure and analysis rather than entirely new proof techniques. The limited search scope (29 papers) means these assessments capture top semantic matches rather than exhaustive coverage.

Based on the top-29 semantic matches and taxonomy structure, the work appears to occupy a sparsely populated research direction (one of two papers in its leaf). The online formulation and regret analysis seem relatively novel, while the algorithmic approach draws on established bandit techniques. The analysis does not cover the full breadth of opinion dynamics or online learning literature, so conclusions about novelty remain provisional pending broader review.

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

Contribution 1: Online formulation of polarization and disagreement minimization as regret minimization problem

Description: The authors introduce a novel online learning framework for minimizing polarization and disagreement in the Friedkin-Johnsen opinion dynamics model under incomplete information. Unlike prior work assuming full knowledge of innate opinions, this formulation casts the problem as a stochastic low-rank matrix bandit problem where the learner observes only scalar feedback after each intervention.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Coevolution of opinion dynamics and recommendation system: Modeling analysis and reinforcement learning based manipulation

URL: [View paper](#)

Brief Assessment

Opinion Recommendation Coevolution[55] focuses on co-evolution of opinion dynamics with recommendation systems in social networks, not on online regret minimization for polarization reduction under incomplete information in the Friedkin-Johnsen model.

2. Minimizing Polarization and Disagreement in the Friedkin-Johnsen Model with Unknown Innate Opinions

URL: [View paper](#)

Brief Assessment

Friedkin Johnsen Unknown Opinions[4] addresses a budgeted query setting where innate opinions are partially observable through selective queries, not an online regret minimization framework with sequential scalar feedback.

3. Toward a social conflict evolution model: Examining the adverse power of conflictual social interaction in online learning

URL: [View paper](#)

Brief Assessment

Social Conflict Evolution[58] examines social conflict in online learning contexts through qualitative case study methods, not regret minimization or bandit algorithms for opinion dynamics interventions.

4. Harmony amidst division: leveraging genetic algorithms to counteract polarisation in online platforms

URL: [View paper](#)

Brief Assessment

Genetic Algorithms Polarization[56] focuses on using genetic algorithms to reduce polarization in social networks, while the original paper formulates the problem as a stochastic low-rank matrix bandit with regret minimization under incomplete information in the Friedkin-Johnsen model.

5. How to Mitigate Disagreement and Polarization in Opinion Formation Processes on Social Networks

URL: [View paper](#)

Brief Assessment

Mitigate Disagreement Networks[5] focuses on seeding-based interventions in social networks using greedy heuristics and online learning to estimate unknown edge weights, not on formulating the problem as a stochastic low-rank matrix bandit with regret minimization. The candidate does not address the bandit framework or the specific online learning formulation presented in the original paper.

6. Knowledge Co-Construction in online learning: applying social learning analytic methods and artificial intelligence

URL: [View paper](#)

Brief Assessment

Knowledge Co Construction[54] focuses on social learning analytics and knowledge co-construction in online learning environments, not on regret minimization frameworks for opinion dynamics or bandit algorithms.

7. Rebalancing social feed to minimize polarization and disagreement

URL: [View paper](#)

Brief Assessment

Rebalancing Social Feed[8] addresses polarization minimization in a static, offline setting using projected gradient descent on a fixed graph structure. It does not formulate the problem as online regret minimization or employ bandit feedback.

8. Centrality-Weighted Opinion Dynamics: Disagreement and Social Network Partition

URL: [View paper](#)

Brief Assessment

Centrality Weighted Disagreement[59] focuses on centrality-weighted opinion dynamics and network partitioning algorithms, not on online learning or regret minimization frameworks for sequential interventions.

9. Constructive online disagreement

URL: [View paper](#)

Brief Assessment

Constructive Online Disagreement[57] focuses on Wikipedia dispute resolution and detecting constructive disagreements through NLP methods, not on opinion dynamics models or regret minimization frameworks for polarization reduction.

10. Opinion Dynamics of Learning Agents: Does Seeking Consensus Lead to Disagreement?

URL: [View paper](#)

Brief Assessment

Learning Agents Disagreement[60] studies opinion dynamics in adaptive agent populations using boolean perceptrons, focusing on faction emergence rather than regret minimization frameworks. It does not address online learning with bandit feedback or low-rank matrix formulations.

Contribution 2: Two-stage algorithm with subspace estimation and dimensionality reduction

Description: The authors develop a two-stage algorithm that first estimates the latent subspace containing the unknown parameter matrix using nuclear-norm regularized least-squares, then runs a linear bandit method in a reduced $(2|V|-1)$ -dimensional space. This approach significantly reduces the problem dimension from $|V|^2$ to $O(|V|)$.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Efficient Generalized Low-Rank Tensor Contextual Bandits

URL: [View paper](#)

Prior Art Analysis

Efficient Tensor Bandits[40] demonstrates that a two-stage approach combining subspace estimation via nuclear-norm regularization followed by dimensionality reduction was already established in prior work. The candidate paper explicitly describes a nearly identical algorithmic framework: first estimating a low-rank subspace using nuclear-norm penalized optimization, then transforming the problem into a lower-dimensional linear bandit problem. Both papers employ nuclear-norm regularization for subspace estimation in stage 1, followed by dimensional reduction and linear bandit methods in stage 2, indicating this contribution was not novel to the original paper.

Evidence

Evidence 1 - **Rationale:** Both papers describe a two-stage algorithmic framework where the first stage explores/estimates a low-rank subspace and the second stage refines the solution using this subspace information. - **Original:** we propose a two-stage algorithm. the algorithm first performs subspace estimation to identify an underlying low-dimensional structure of the problem, and then employs a linear bandit algorithm within the compact $2|v| - 1$ dimensional representation, which is a significant reduction from the original... - **Candidate:** we introduce a new algorithm named "generalized low-rank tensor exploration subspace then refine" (g-lowtestr). the goal is to achieve an efficient and effective balance between exploration and exploitation when dealing with generalized low-rank tensor contextual bandits problems. such exploration a...

Evidence 2 - **Rationale:** Both papers explicitly describe the same two-stage process: nuclear-norm based subspace estimation followed by transformation to a lower-dimensional problem. - **Original:** the algorithm first performs subspace estimation to identify an underlying low-dimensional structure, and then employs a linear bandit algorithm within the compact $2|v| - 1$ dimensional representation derived from the estimated subspace. - **Candidate:** initially, g-lowtestr estimates w^* by solving a tensor nuclear norm penalized minimization problem using raw data (action features x_t and reward y_t). this provides an approximate estimation of the low-rank subspace of w^* . subsequently, based on this estimated low-rank subspace, g-lowtestr further tr...

Evidence 3 - **Rationale:** Both papers use nuclear-norm regularized optimization to estimate the low-rank subspace in the first stage, demonstrating the same methodological approach. - **Original:** to solve this bandit problem, we propose a two-stage algorithm. the algorithm first performs subspace estimation to identify an underlying low-dimensional structure of the problem, and then employs a linear bandit algorithm within the compact $2|v| - 1$ dimensional representation - **Candidate:** we first collect data to find the accurate row and column subspace estimation of w^* . especially, at each time t , the learner selects an action x_t and receives the corresponding reward y_t . given the horizon t_1 , a series of action and reward pairs $\{x_t, y_t\}_{t=1}^{t_1}$ are accumulated. considering the low-ra...

2. Low-rank bandits via tight two-to-infinity singular subspace recovery

URL: [View paper](#)

Prior Art Analysis

Tight Singular Subspace[38] demonstrates that two-stage algorithms combining subspace estimation followed by dimensionality reduction for low-rank matrix bandits were already established prior to the original paper. The candidate explicitly describes a two-phase approach where the first phase estimates singular subspaces of the reward matrix using spectral methods, then reduces the problem to a lower-dimensional linear bandit in the second phase. This matches the core algorithmic contribution claimed by the original paper, showing the approach was not novel.

Evidence

Evidence 1 - **Rationale:** Both papers describe the same two-stage algorithmic structure: first estimating subspaces, then solving a reduced-dimensional bandit problem. - **Original:** we propose a two-stage algorithm. the algorithm first performs subspace estimation to identify an underlying low-dimensional structure of the problem, and then employs a linear bandit algorithm within the compact $2|v| - 1$ dimensional representation - **Candidate:** we propose a generic method to design two-phase algorithms for low-rank bandit problems. in the first phase, we estimate the singular subspaces of the reward matrix m using simple spectral methods... in the second phase, the algorithm solves the resulting misspecified contextual linear bandit problem...

Evidence 2 - **Rationale:** Both describe reducing from a high-dimensional space ($|v|^2$ or mn) to a lower-dimensional representation ($2|v|-1$ or $r(m+n)-r^2$) through subspace estimation. - **Original:** the algorithm first performs subspace estimation to identify an underlying low-dimensional structure, and then employs a linear bandit algorithm within the compact $2|v| - 1$ dimensional representation, which is a significant reduction from the original $|v|^2$ -dimensional space - **Candidate:** in the first phase, we estimate the singular subspaces of the

reward matrix m using simple spectral methods... these new error bounds allow us to reformulate our low-rank bandits as a misspecified contextual linear bandit problem with dimension $(m+n)-r$ and controlled misspecification

Evidence 3 - Rationale: Both papers use spectral methods in the first phase to estimate the low-rank subspace structure of the parameter matrix. - **Original:** our main algorithm, `opd-min-est`, is summarized in algorithm 1. 4.1 stage 1: explore opinion subspace the initial t_1 rounds are dedicated to an exploration phase designed to learn the low-dimensional subspace containing the true parameter matrix $\theta^* = \text{ssT}$. to achieve this, we employ an estimator \hat{b}_a . - **Candidate:** to estimate the singular subspaces of m , we construct \hat{f}_m as: for all $(i, j) \in [m] \times [n]$, $\hat{f}_{m,i,j} = \frac{1}{t} \sum_{t=1}^t \omega_{i,j} \text{tr}(\xi_t \{(\text{it}, \text{jt}) = (i, j)\})$. then, we let \hat{c}_m be the best r -rank approximation of \hat{f}_m

Evidence 4 - Rationale: Both explicitly describe a second stage that reduces the bandit problem to a lower-dimensional linear bandit using the subspace estimates from the first stage. - **Original:** 4.2 stage 2: dimensionality reduction and subspace linear bandit after the subspace estimation phase, we reduce the original matrix bandit problem into a lower-dimensional linear bandit problem using the nuclear-norm penalized estimator \hat{b}_θ . - **Candidate:** in this section, we assume that we use t_1 reward observations to estimate the singular subspaces u and v of m as described in the previous section... based on the estimates of u and v , one could... reduce the problem faced by the learner in the remaining $(t - t_1)$ rounds to an almost low-dimensional ...

3. Low-rank contextual reinforcement learning from heterogeneous human feedback

URL: [View paper](#)

Brief Assessment

Low Rank Human Feedback[41] addresses contextual RLHF with heterogeneous human preferences using low-rank matrix structures, while the original paper focuses on opinion dynamics in social networks via low-rank matrix bandits. The candidate's two-stage approach (subspace estimation via rank-constrained MLE, then policy optimization in reduced space) serves a fundamentally different application domain—personalized language model alignment rather than polarization minimization in graph-based opinion systems.

4. On high-dimensional and low-rank tensor bandits

URL: [View paper](#)

Prior Art Analysis

Low Rank Tensor Bandits[35] demonstrates prior work on a two-stage algorithm that performs subspace estimation followed by dimensionality reduction for low-rank problems. The paper explicitly describes estimating low-dimensional subspaces in Phase A using tensor regression techniques, then converting the problem to a lower-dimensional space in Phase B. This structure directly parallels the ORIGINAL paper's claimed novel contribution of first estimating the latent subspace and then running a linear bandit method in reduced dimensions.

Evidence

Evidence 1 - Rationale: Both papers describe a two-stage approach where the first stage estimates low-dimensional subspaces and the second stage operates in a reduced parameter space. - **Original:** we propose a two-stage algorithm. the algorithm first performs subspace estimation to identify an underlying low-dimensional structure of the problem, and then employs a linear bandit algorithm within the compact $2|v| - 1$ dimensional representation. - **Candidate:** the `tofu` algorithm (presented in alg. 1) has two phases: a and b. phase a aims at estimating the unknown system tensor x up to a certain precision, especially its low-dimensional subspaces. with this estimate, the original bandit problem can be reformulated, such that the new problem has (approximate...)

Evidence 2 - Rationale: Both papers use the first phase to estimate the low-rank structure through subspace estimation techniques. - **Original:** the algorithm first performs subspace estimation to identify an underlying low-dimensional structure, and then employs a linear bandit algorithm within the compact $2|v| - 1$ dimensional representation, which is a significant reduction from the original $|v|$ 2-dimensional space. - **Candidate:** phase a adopts techniques in low-rank tensor regression (also known as low-rank tensor factorization or completion from linear measurements). especially, it considers the problem of estimating a low-rank tensor x by a collection of data $\{(a_t, r_t) : t \in [t_1]\}$ that are associated with x through eqn. (1), wh...

5. Generalized low-rank matrix contextual bandits with graph information

URL: [View paper](#)

Brief Assessment

[Final Audit Failure] The model insisted on a refutation claim but failed to provide verifiable evidence after multiple retries. Marked as `cannot_refute` for safety. Please manually verify the candidate text.

6. Multiagent low-dimensional linear bandits

URL: [View paper](#)

Brief Assessment

Multiagent Linear Bandits[37] addresses a multi-agent collaborative bandit problem where agents share subspace recommendations, not a single-agent opinion dynamics problem with low-rank matrix structure.

7. Effective generalized low-rank tensor contextual bandits

URL: [View paper](#)

Brief Assessment

Generalized Tensor Bandits[36] addresses tensor-structured contextual bandits with generalized linear rewards, while the original paper focuses on low-rank matrix bandits for opinion dynamics with quadratic objectives. The tensor formulation and generalized linear model represent fundamentally different problem structures.

8. High-dimensional gaussian process bandits

URL: [View paper](#)

Prior Art Analysis

Gaussian Process Bandits[42] demonstrates prior work on two-stage algorithms that perform subspace estimation followed by optimization in reduced dimensions. The paper explicitly describes a two-stage approach: first estimating a low-dimensional subspace using nuclear-norm regularized methods, then applying Gaussian process optimization (GP-UCB) in the reduced k -dimensional space. This directly parallels the ORIGINAL paper's claimed novel contribution of a two-stage algorithm with subspace estimation and dimensionality reduction from $|V|^2$ to $O(|V|)$ dimensions.

Evidence

Evidence 1 - Rationale: Both papers describe a two-stage approach where the first stage estimates a low-dimensional subspace and the second stage performs optimization in that reduced space. - **Original:** we propose a two-stage algorithm. the algorithm first performs subspace estimation to identify an underlying low-dimensional structure of the problem, and then employs a linear bandit algorithm within the compact $2|v| - 1$ dimensional representation, which is a significant reduction from the original... - **Candidate:** the `si-bo` algorithm performs two separate exploration and exploitation stages: (1) subspace identification (si), i.e. estimating the subspace on which the function is supported, and then (2) bayesian optimization (bo), in order to optimize the function on the learned subspace

Evidence 2 - **Rationale:** Both papers use low-rank recovery techniques to estimate the subspace in the first stage, then optimize in the reduced representation. - **Original:** the algorithm first performs subspace estimation to identify an underlying low-dimensional structure, and then employs a linear bandit algorithm within the compact $2\sqrt{v} - 1$ dimensional representation derived from the estimated subspace - **Candidate:** we first evaluate the function at several suitably chosen points and then use a low-rank recovery algorithm to compute a matrix \hat{a} that spans a subspace well aligned with the one generated by the true matrix a . once we have computed \hat{a} , similarly to [22, 7], we define the function which we optimize as ...

Evidence 3 - **Rationale:** Both papers employ low-rank matrix recovery for subspace estimation as the foundation of their two-stage approach. - **Original:** we propose a two-stage algorithm based on low-rank matrix bandits. the algorithm first performs subspace estimation to identify an underlying low-dimensional structure, and then employs a linear bandit algorithm within the compact dimensional representation derived from the estimated subspace - **Candidate:** we learn using the approach from [7], which reduces the learning problem to that of low rank matrix recovery. we construct a set of $m \times$ points $c = [\xi_1, \dots, \xi_{m \times}]$, which we call sampling centers, and consider the matrix x of gradients at those points $x = [\nabla f(\xi_1), \dots, \nabla f(\xi_{m \times})]$

Evidence 4 - **Rationale:** Both papers apply a bandit optimization algorithm (GP-UCB or linear bandit) in the reduced-dimensional space after subspace estimation. - **Original:** the algorithm first performs subspace estimation to identify an underlying low-dimensional structure of the problem, and then employs a linear bandit algorithm within the compact $2\sqrt{v} - 1$ dimensional representation - **Candidate:** once we have an approximate \hat{a} , we optimize the function $\hat{g}(z) = f(\hat{a}tz)$ on the low-dimensional domain $z = bk(1 + \epsilon)$. concretely, we use gp-ucb [12], because it exhibits state of the art empirical performance, and enjoys strong theoretical bounds for the cumulative regret

9. A simple unified framework for high dimensional bandit problems

URL: [View paper](#)

Prior Art Analysis

High Dimensional Bandit Framework[43] demonstrates that a two-stage approach with subspace estimation followed by dimensionality reduction was already established in the bandit literature. The candidate paper explicitly describes an 'explore-the-structure-then-commit (estc)' framework that first explores to learn low-dimensional structure, then commits to exploitation in a reduced-dimensional space. This directly parallels the original paper's claimed novelty of first estimating the latent subspace using nuclear-norm regularization, then running a linear bandit in reduced dimensions. The candidate's framework applies to multiple high-dimensional bandit problems including low-rank matrix bandits, achieving regret bounds that depend logarithmically on dimensions rather than polynomially, which predates the original paper's contribution.

Evidence

Evidence 1 - **Rationale:** Both papers describe a two-stage framework where the first stage explores to identify low-dimensional structure and the second stage exploits in a reduced-dimensional space. The candidate's ESTC framework explicitly establishes this paradigm for high-dimensional bandit problems. - **Original:** we propose a two-stage algorithm. the algorithm first performs subspace estimation to identify an underlying low-dimensional structure of the problem, and then employs a linear bandit algorithm within the compact $2\sqrt{v} - 1$ dimensional representation, which is a significant reduction from the original... - **Candidate:** we present a simple and unified algorithm framework named explore-the-structure-then-commit (estc) for high dimensional stochastic bandit problems and provide a problem-independent regret analysis framework for our algorithm. we show that to ensure a desirable regret, one simply needs to ensure that ...

Evidence 2 - **Rationale:** The candidate paper demonstrates that their framework achieves logarithmic dependence on dimensions for low-rank matrix bandits, which directly addresses the same problem class as the original paper. This shows the two-stage subspace estimation and dimensionality reduction approach was already established for low-rank matrix bandits. - **Original:** the algorithm first performs subspace estimation to identify an underlying low-dimensional structure, and then employs a linear bandit algorithm within the compact dimensional representation derived from the estimated subspace. - **Candidate:** we demonstrate the usefulness of our framework by applying it to different high dimensional bandit problems. we show that under a mild unified assumption, our algorithm can achieve desirable regret bounds. we first prove that our algorithm achieves comparable regret bounds in the lasso bandit problem ...

Contribution 3: Theoretical regret bound with sublinear dependence on time horizon

Description: The authors establish a cumulative regret bound of $eO(\sqrt{V}\sqrt{T})$ for their algorithm, demonstrating optimal \sqrt{T} dependence on the time horizon and linear rather than quadratic dependence on the number of users. This represents the first theoretical guarantee for sequential interventions in opinion dynamics without complete knowledge of innate opinions.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Quasi-safe bandit algorithms for the bid optimization problem in online advertising

URL: [View paper](#)

Brief Assessment

Bid Optimization Bandits[52] addresses budget and ROI constraints in online advertising with stochastic bandits, not sequential interventions in opinion dynamics. The technical setting, problem formulation, and application domain are fundamentally different from the original paper's focus on polarization minimization in social networks.

2. Would I regret being different? The influence of social norms on attitudes toward AI usage

URL: [View paper](#)

Brief Assessment

Social Norms AI[47] studies regret in human decision-making about AI usage under social norms, not regret bounds for sequential interventions in opinion dynamics or bandit algorithms.

3. Stochastic Top K-Subset Bandits with Linear Space and Non-Linear Feedback with Applications to Social Influence Maximization

URL: [View paper](#)

Brief Assessment

Top K Social Influence[53] addresses combinatorial bandits for influence maximization with non-linear feedback, achieving $\tilde{O}(k^{1/2}n^{1/3}T^{2/3})$ regret. The original paper focuses on opinion dynamics minimization with matrix bandits achieving $\tilde{O}(\sqrt{V}\sqrt{T})$ regret—different problem formulations and objectives.

4. Contextual bandit with herding effects: Algorithms and recommendation applications

URL: [View paper](#)

Brief Assessment

Herding Effects Bandits[44] addresses contextual bandits with biased feedback from herding effects in recommendation systems, not sequential interventions in opinion dynamics under incomplete information about innate opinions.

5. Combinatorial Rising Bandit

URL: [View paper](#)

Brief Assessment

Combinatorial Rising Bandit[49] addresses a fundamentally different problem structure—combinatorial bandits with rising rewards where base arms enhance over time. The original paper focuses on sequential interventions in opinion dynamics with low-rank matrix structure, not combinatorial action spaces with rising rewards.

6. $\hat{\alpha}$ in social influence by social distance in car-sharing decisions under uncertainty: A regret-minimizing hybrid choice model framework based on sequential stated $\hat{\alpha}$

URL: [View paper](#)

Brief Assessment

Social Distance Car Sharing[46] focuses on car-sharing decisions under uncertainty using hybrid choice models and stated adaptation experiments. This is a completely different domain from sequential interventions in opinion dynamics and does not address regret bounds for such problems.

7. Regret, Uncertainty, and Bounded Rationality in Norm-Driven Decisions

URL: [View paper](#)

Brief Assessment

Bounded Rationality Norms[48] studies regret in vaccination decisions within opinion dynamics, not sequential interventions in social networks. The candidate focuses on behavioral modeling of norm-driven choices rather than algorithmic regret bounds for bandit problems.

8. Incentive mechanism for spatial crowdsourcing with unknown social-aware workers: A three-stage stackelberg game approach

URL: [View paper](#)

Brief Assessment

Spatial Crowdsourcing Stackelberg[45] addresses worker recruitment in spatial crowdsourcing using multi-armed bandits and Stackelberg games, not sequential interventions in opinion dynamics. The regret bound applies to a fundamentally different problem domain (crowdsourcing task allocation vs. opinion network interventions).

9. Private owners' propensity to engage in shared parking schemes under uncertainty: comparison of alternate hybrid expected utility-regret-rejoice choice models

URL: [View paper](#)

Brief Assessment

Shared Parking Regret[51] focuses on modeling regret and rejoice emotions in parking owners' decision-making using stated choice experiments, not on establishing theoretical regret bounds for sequential interventions in opinion dynamics or bandit algorithms.

10. Family planning decision-making in relation to psychiatric disorders in women: a qualitative focus group study

URL: [View paper](#)

Brief Assessment

Family Planning Psychiatric[50] focuses on qualitative analysis of family planning decisions in women with psychiatric disorders, not regret bounds for sequential interventions in opinion dynamics or bandit algorithms.

Appendix: Text Similarity Detection

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

References

-
- [0] Online Minimization of Polarization and Disagreement via Low-Rank Matrix Bandits [View paper](#)
 - [1] Link recommendation algorithms and dynamics of polarization in online social networks [View paper](#)
 - [2] Towards consensus: Reducing polarization by perturbing social networks [View paper](#)
 - [3] The algorithmic Management of Polarization and Violence on social media [View paper](#)
 - [4] Minimizing Polarization and Disagreement in the Friedkin-Johnsen Model with Unknown Innate Opinions [View paper](#)
 - [5] How to Mitigate Disagreement and Polarization in Opinion Formation Processes on Social Networks [View paper](#)
 - [6] Modelling the mitigation of anti-vaccine opinion propagation to suppress epidemic spread: A computational approach [View paper](#)
 - [7] Minimizing polarization and disagreement in social networks via link recommendation [View paper](#)
 - [8] Rebalancing social feed to minimize polarization and disagreement [View paper](#)
 - [9] Agent-based modelling of polarized news and opinion dynamics in social networks: a guidance-oriented approach [View paper](#)
 - [10] Network-aware recommender system via online feedback optimization [View paper](#)
 - [11] Modeling Disinformation Spread in Social Networks: Phase Transitions and Mean-Field Analysis [View paper](#)
 - [12] Optimizing Social Network Interventions via Hypergradient-Based Recommender System Design [View paper](#)
 - [13] Analyzing the impact of filter bubbles on social network polarization [View paper](#)
 - [14] Modeling the Impact of Timeline Algorithms on Opinion Dynamics Using Low-rank Updates [View paper](#)
 - [15] Control of kinetic opinion dynamics in popularity-adaptive social networks [View paper](#)
 - [16] Polarization on social media: Micro-level evidence and macro-level implications [View paper](#)
 - [17] Opinion amplification causes extreme polarization in social networks [View paper](#)
 - [18] Mitigating opinion polarization in social networks using adversarial attacks [View paper](#)
 - [19] Opinion de-polarization of social networks with GNNs [View paper](#)
 - [20] Agents of Discord: Modeling the Impact of Political Bots on Opinion Polarization in Social Networks [View paper](#)
 - [21] Decoding Echo Chambers: LLM-Powered Simulations Revealing Polarization in Social Networks [View paper](#)
 - [22] MTOS: A LLM-Driven Multi-topic Opinion Simulation Framework for Exploring Echo Chamber Dynamics [View paper](#)
 - [23] Understanding filter bubbles and polarization in social networks [View paper](#)
 - [24] Reducing polarization in social networks with adversarial opinion perturbations [View paper](#)
 - [25] Steering polarization toward consensus in signed majority-vote opinion models [View paper](#)

- [26] Time of influence as a decision variable in game-theoretic models of opinion dynamics in social groups [View paper](#)
- [27] Opinion spreading and evolution in complex social networks: a simulation study [View paper](#)
- [28] Harmonizing vs Polarizing Platform Influence Functions [View paper](#)
- [29] Recommender System Design via Online Feedback Optimization [View paper](#)
- [30] Finite-Time Elimination of Disagreement of Opinion Dynamics via Covert Noise [View paper](#)
- [31] Disagreement and Polarization in Two-Party Social Networks [View paper](#)
- [32] Mitigating Influence of Disinformation Propagation Using Uncertainty-Based Opinion Interactions [View paper](#)
- [33] Modeling how social network algorithms can influence opinion polarization [View paper](#)
- [34] Generalized low-rank matrix contextual bandits with graph information [View paper](#)
- [35] On high-dimensional and low-rank tensor bandits [View paper](#)
- [36] Effective generalized low-rank tensor contextual bandits [View paper](#)
- [37] Multiagent low-dimensional linear bandits [View paper](#)
- [38] Low-rank bandits via tight two-to-infinity singular subspace recovery [View paper](#)
- [39] Bilinear bandits with low-rank structure [View paper](#)
- [40] Efficient Generalized Low-Rank Tensor Contextual Bandits [View paper](#)
- [41] Low-rank contextual reinforcement learning from heterogeneous human feedback [View paper](#)
- [42] High-dimensional gaussian process bandits [View paper](#)
- [43] A simple unified framework for high dimensional bandit problems [View paper](#)
- [44] Contextual bandit with herding effects: Algorithms and recommendation applications [View paper](#)
- [45] Incentive mechanism for spatial crowdsourcing with unknown social-aware workers: A three-stage stackelberg game approach [View paper](#)
- [46] $\hat{\pi}$ in social influence by social distance in car-sharing decisions under uncertainty: A regret-minimizing hybrid choice model framework based on sequential stated $\hat{\pi}$ [View paper](#)
- [47] Would I regret being different? The influence of social norms on attitudes toward AI usage [View paper](#)
- [48] Regret, Uncertainty, and Bounded Rationality in Norm-Driven Decisions [View paper](#)
- [49] Combinatorial Rising Bandit [View paper](#)
- [50] Family planning decision-making in relation to psychiatric disorders in women: a qualitative focus group study [View paper](#)
- [51] Private owners's propensity to engage in shared parking schemes under uncertainty: comparison of alternate hybrid expected utility-regret-rejoice choice models [View paper](#)
- [52] Quasi-safe bandit algorithms for the bid optimization problem in online advertising [View paper](#)
- [53] Stochastic Top K-Subset Bandits with Linear Space and Non-Linear Feedback with Applications to Social Influence Maximization [View paper](#)
- [54] Knowledge Co-Construction in online learning: applying social learning analytic methods and artificial intelligence [View paper](#)
- [55] Coevolution of opinion dynamics and recommendation system: Modeling analysis and reinforcement learning based manipulation [View paper](#)
- [56] Harmony amidst division: leveraging genetic algorithms to counteract polarisation in online platforms [View paper](#)
- [57] Constructive online disagreement [View paper](#)
- [58] Toward a social conflict evolution model: Examining the adverse power of conflictual social interaction in online learning [View paper](#)
- [59] Centrality-Weighted Opinion Dynamics: Disagreement and Social Network Partition [View paper](#)
- [60] Opinion Dynamics of Learning Agents: Does Seeking Consensus Lead to Disagreement? [View paper](#)