Emergence of Complex Network Structures in LLM Social Simulations

Aleksandar Tomašević, Darja Cvetković, Sara Major, Slobodan Maletić, Miroslav Anđelković, Ana Vranić, Boris Stupovski, Dušan Vudragović, Aleksandar Bogojević. Marija Mitrović Dankulov

Belgrade, August, 2025

Motivation

- Post-API era of social media research (Mimizuka et al., 2025)
- Generative LLM simulations are promising alternative to ABMs (Kozlowski & Evans, 2025)
- Start with simple, small simulations of niche community

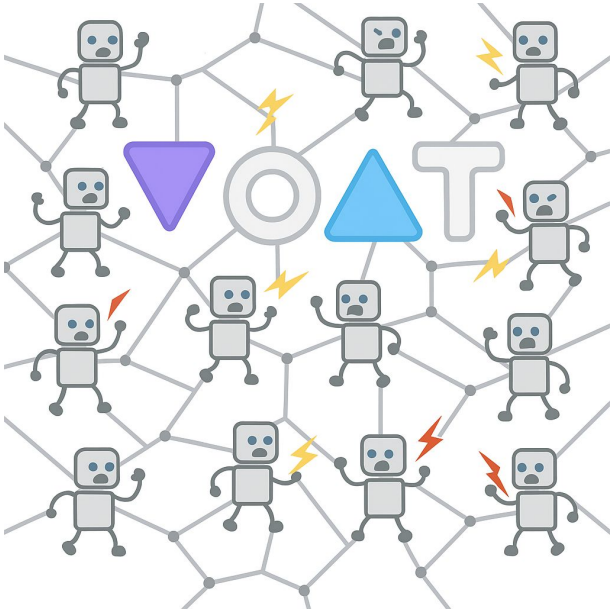

Mimizuka, K., Brown, M. A., Yang, K.-C., & Lukito, J. (2025). Post-post-API age: Studying digital platforms in scant data access times. http://arxiv.org/abs/2505.09877

Kozlowski, A. C., & Evans, J. (2025). Simulating subjects: The promise and peril of artificial intelligence stand-ins for social agents and interactions. https://doi.org/10.1177/00491241251337316

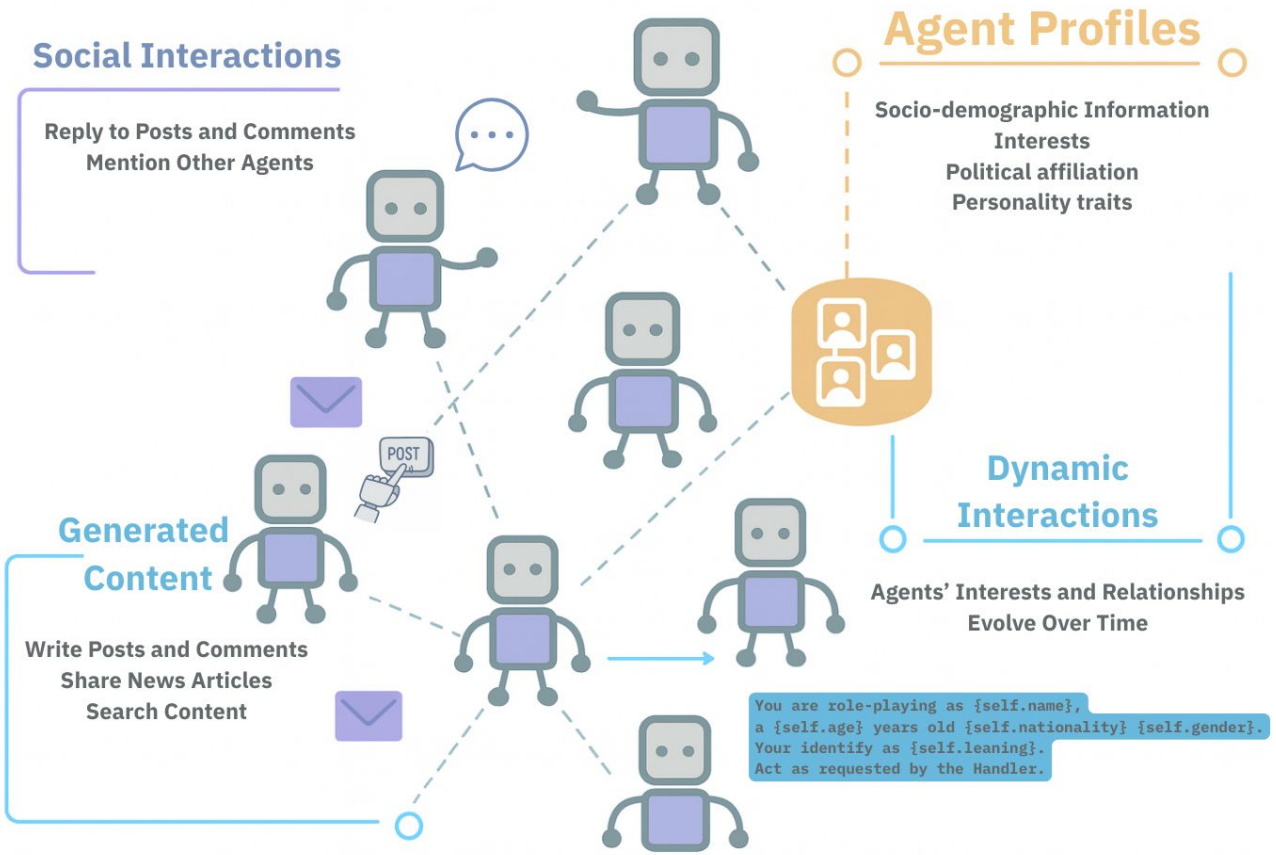
Social Interactions
Reply to Posts and Comments
Mention Other Agents
Agent Profiles
Socio-demographic Information
Interests
Political affiliation
Personality traits
POST
Generated Content
Write Posts and Comments
Share News Articles
Search Content
Dynamic Interactions
Agents' Interests and Relationships
Evolve Over Time
You are role-playing as {self.name},
a {self.age} years old {self.nationality} {self.gender}.
Your identify as {self.leaning}.
Act as requested by the Handler.

Outline

- We use **Voat** (alt-right clone of Reddit) for calibration
- Reddit-like: feed-focused instead of follower-network; **reverse chronological feed**
- Run simulation, map user–user **undirected weighted interaction network**
- Validation: **operational validity** against matched Voat windows; compare activity distributions and network structure

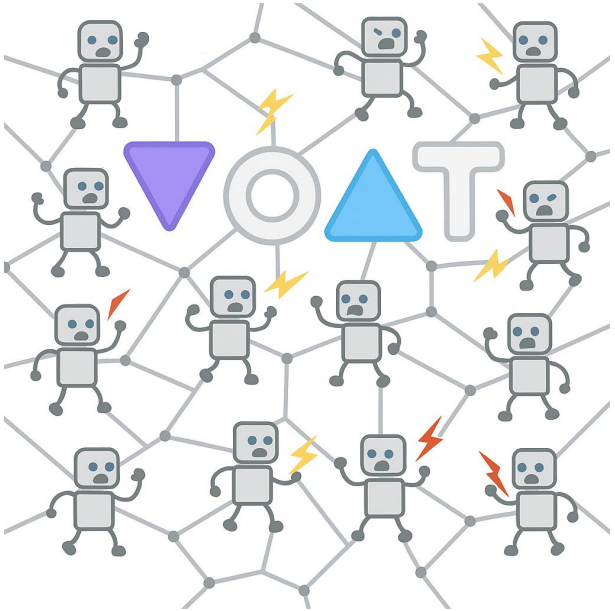

Research Question

Do LLM agents, under Voat-like rules and real content flows, generate interaction networks and activity patterns that match real Voat at distributional and structural levels?

Methods: Study design & calibration

- Target community: **v/technology**
- Calibration windows: 10 non-overlapping 30-day samples from **MADOC** dataset (Mitrović Dankulov et al., 2025)
- Derive daily activity, churn, thread depth.

Means and standard deviations are across windows; min and max are window extremes.

Metric	Mean	SD	Min–Max
Users per 30d sample (unique)	576.10	111.11	385–721
Active users per day	31.52	5.96	21.50–40.57
New users per day (%)	59.44	2.29	55.69–62.49
Churned users per day (%)	75.13	1.73	71.95–76.80
Comments per post (sample-level)	1.07	0.09	0.96–1.19
Posts per 30d sample	618.40	109.69	440–819
Comments per 30d sample	664.50	135.36	435–864
Active users on day 1	32.60	15.05	14–66

Mitrović Dankulov, M., Tomašević, A., Maletić, S., Anđelković, M., Vranić, A., Cvetković, D., Stupovski, B., Vudragović, D., Major, S., & Bogojević, A. (2025). Multi-Platform Aggregated Dataset of Online Communities (MADOC). Proceedings of the International AAAI Conference on Web and Social Media, 19, 2529–2538. https://doi.org/10.1609/icwsm.v19i1.35954

Methods: Engine & Agents

- **YSocial digital-twin architecture** (Rossetti et al., 2024)

  Server = full platform state and recommender. Client = clock, scheduling, queries. Ollama serves LLMs on request.
- Base LLM: **Dolphin 3.0** (based on Llama-3.1-8B) uncensored to allow disagreement and toxicity
- Personas: **uniform sampling** of demographics + activity & toxicity parameters
- Content seeding: 30+ tech **RSS feeds**.

Rossetti, G., Stella, M., Cazabet, R., Abramski, K., Cau, E., Citraro, S., Failla, A., Improta, R., Morini, V., & Pansanella, V. (2024). Y Social: an LLM-powered Social Media Digital Twin. arXiv. http://arxiv.org/abs/2408.00818

Agent Population Parameters

Attribute	Values / Range
Locale	English (American)
Education level	high school, bachelor, master, phd
Political leaning	Republican
Age	18–60
Actions per activation (round)	1–2
Number of interests	2–5
Toxicity propensity level	Absolutely No, No, Moderately, Extremely

Interests catalog

- Social Media & Online Platforms
- Internet Policy & Regulation
- Artificial Intelligence
- Electric Vehicles & Transportation
- Software Development
- Clean Energy & Sustainability
- Cybersecurity & Privacy
- Big Tech
- Space Technology
- Open Source Projects

A Day in the Life (of an Agent)

Morning activation

- **10:00 AM (Round 10).** The agent is activated; according to their profile, they will perform two actions in this round.
- **Round action 1:** the simulator offers `[COMMENT, SHARE_LINK, NONE]`; the LLM chooses `SHARE_LINK`.
  - Selects an article from the local news database matching interests; e.g., "New battery tech for grid storage."
  - Reads the article and generates commentary, posted as a root submission with a URL to the source.
- **Round action 2:** the simulator offers `[READ, POST, NONE]`; the LLM chooses `READ`.
  - Reads a recommended post (root + comments up to the configured depth).
  - Lurking behavior; no follow-up action.
- After two actions, the agent becomes inactive.

Evening activation

- **5:00 PM (Round 17).** The agent is activated again; they will perform two actions.
- **Round action 1:** the simulator offers `[READ, COMMENT, NONE]`; the LLM chooses `COMMENT`.
  - Chooses a candidate post, reviews context, and writes a reply.
- **Round action 2:** the simulator offers `[SEARCH, COMMENT, NONE]`; the LLM chooses `NONE`.
  - Observes the feed without acting.
- After two actions, the agent is deactivated and is not activated again for the rest of the day.

Methods: Networks (core-periphery)

- Simulation output: CSV file with all posts and comments: *interaction id, author id, parent post id*
- User–User **undirected weighted interaction network**
  - edges = replies: comment to post, comment to comment
  - weights = interaction counts, normalized

Stochastic Block Model (Gallagher et al., 2021)

- Hub-and-spoke with prior dense core, sparse periphery
- **Multi-run** core–periphery inference (Gibbs/MCMC), take posterior samples, and evaluate MDL plus quality metrics
  - core/periphery densities
  - core–periphery cross-density
- Select a best partition via a composite score (quality + MDL)

Gallagher, R. J., Young, J.-G., & Welles, B. F. (2021). A clarified typology of core-periphery structure in networks. Science Advances, 7(12), eabc9800. https://doi.org/10.1126/sciadv.abc9800

Results - Basic metrics

Metric	Simulation	Voat
Posts	754	704
Comments	802	793
Users (unique)	641	721
Avg thread length	2.06	2.09
Mean toxicity	0.15	0.10

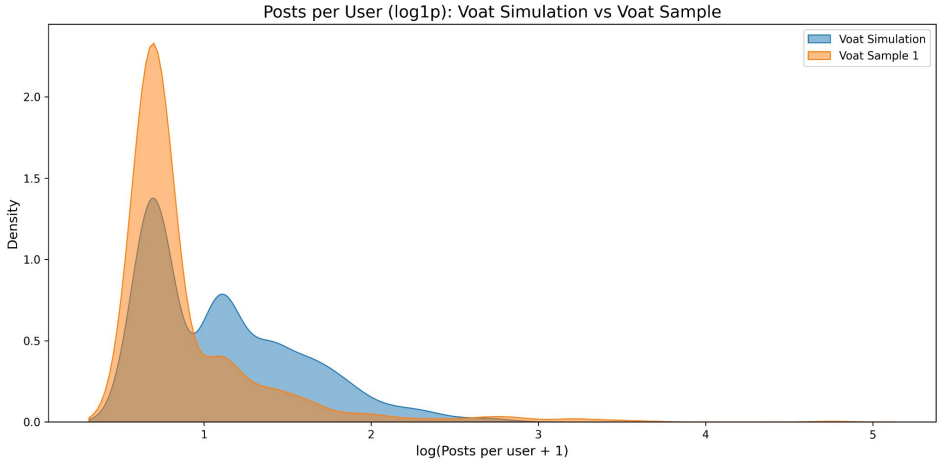

Network statistics

Metric	Simulation	Voat dataset
Nodes	641	554
Edges	711	623
Avg degree	2.218	2.249
Weighted avg degree	0.497	0.415
Clustering coefficient	0.0060	0.0027
Density	0.00347	0.00407
Largest component (nodes)	466 (72.7%)	435 (78.5%)

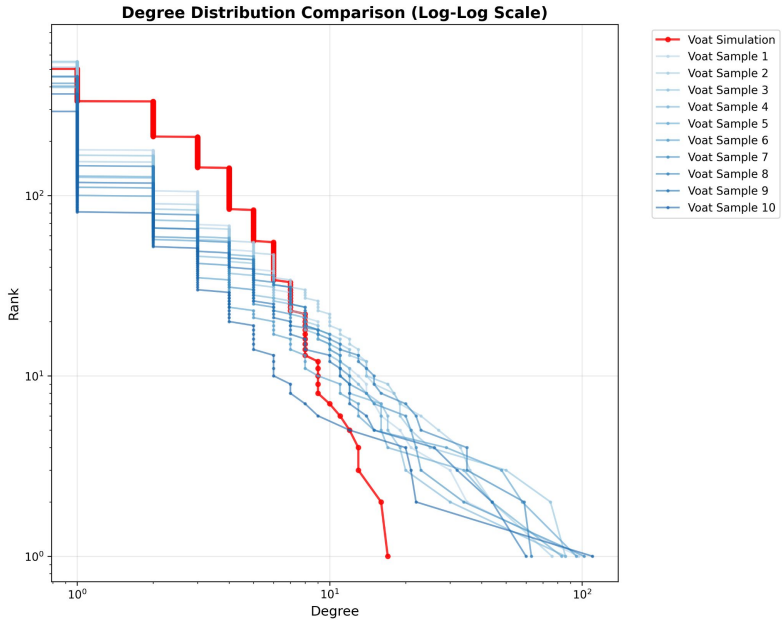

Network visualization

Network visualization

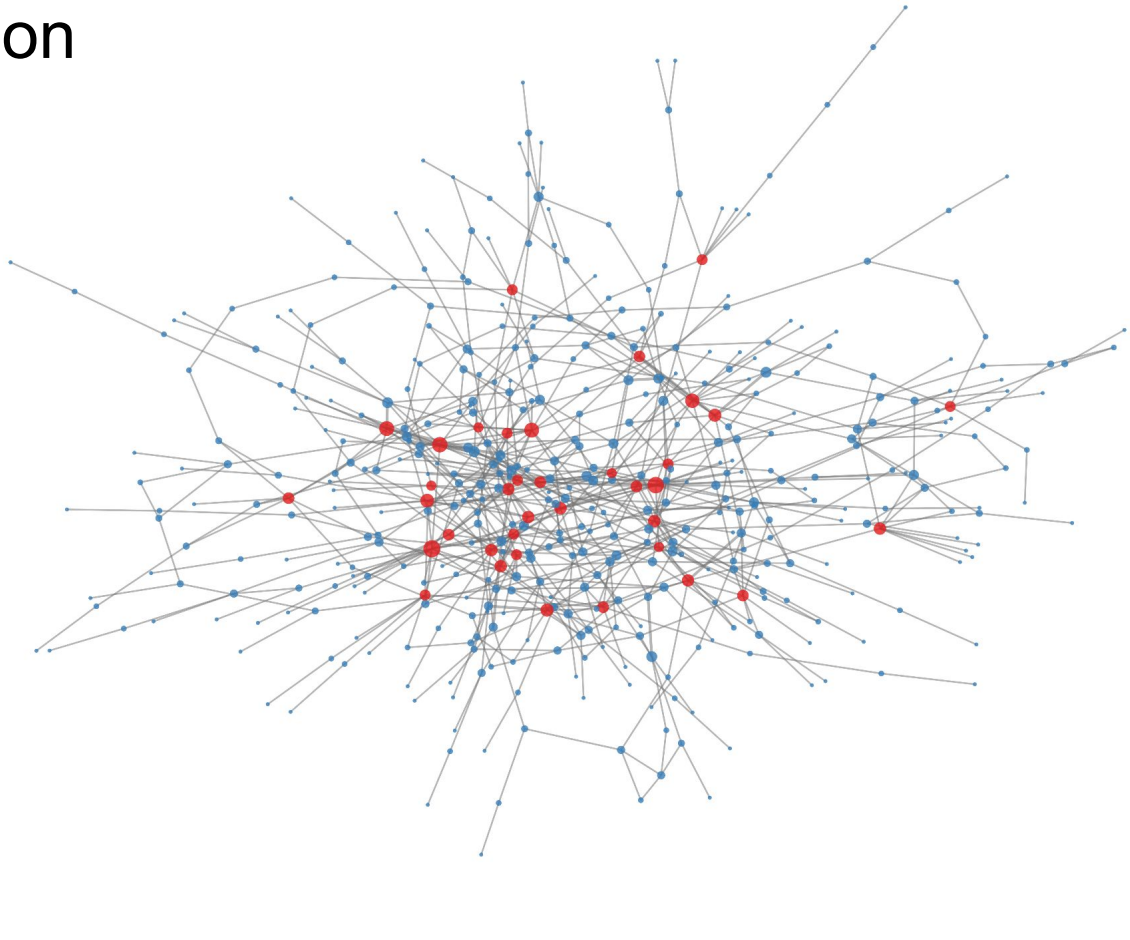

Network comparison

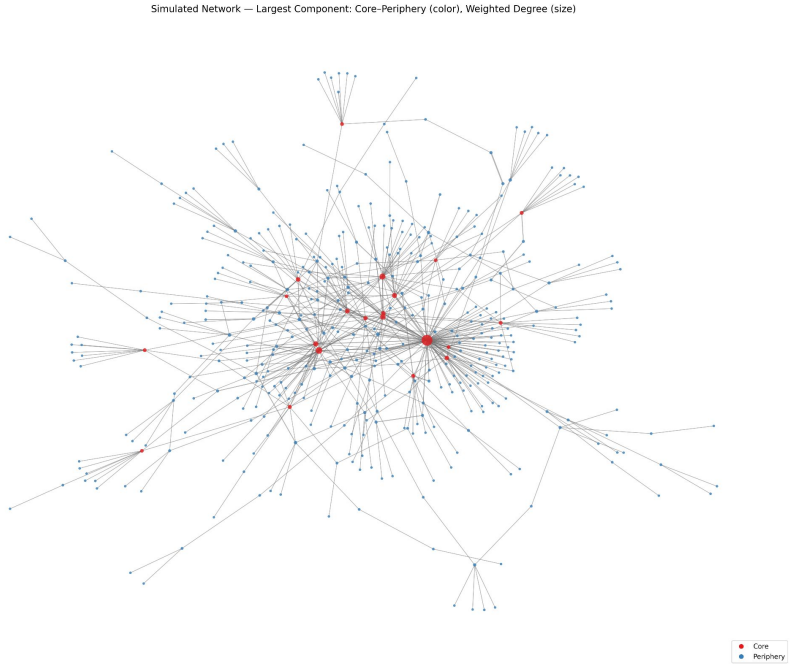

Core-periphery results

Metric	Simulation	Voat
LCC nodes, edges, density	466, 690, ~0.0064	435, 555, ~0.0059
Core size (range; mean)	33–49; ≈40.4	21–23; ≈22.5
Best core size (% LCC)	37 (7.94%)	21 (4.83%)
Core density	~0.057	~0.157
Core–periphery density	~0.015	~0.036
Core avg degree	~8.59	~18.14

Takeaways

- Core–periphery emerges in a simple feed-centric LLM simulation without explicit programming/prompting.
- Real core is smaller and denser with stronger core-periphery coupling; the simulated core is larger and diffuse.
- **Uniform** per-round activity and **simplified feed** damp hub consolidation.
- With realistic calibration and algorithmic feeds, YSocial can produce simulations as a sandbox to reproduce complex network patterns *in silico*.

Thank you!

Email: atomasevic@ipb.ac.rs

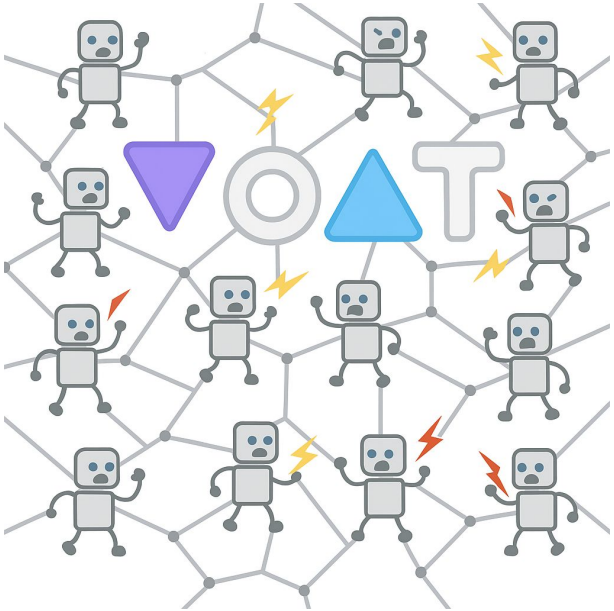

Additional Results

Simulation Configuration: Macro parameters

Parameter	Value
Duration (days)	30
Starting agents	50
New agents/iteration	0.30
Removal/iteration	0.90
Engagement likelihood: post	0.005
Engagement likelihood: link share	0.060
Engagement likelihood: comment	0.060
Engagement likelihood: read	0.40
Engagement likelihood: search	0.10
Max thread length reading	3

Agent Population Parameters

Attribute	Values / Range
Locale	English (American)
Education level	high school, bachelor, master, phd
Political leaning	Republican
Age	18–60
Actions per activation (round)	1–2
Number of interests	2–5
Toxicity propensity level	Absolutely No, No, Moderately, Extremely

Interests catalog

- Social Media & Online Platforms
- Internet Policy & Regulation
- Artificial Intelligence
- Electric Vehicles & Transportation
- Software Development
- Clean Energy & Sustainability
- Cybersecurity & Privacy
- Big Tech
- Space Technology
- Open Source Projects

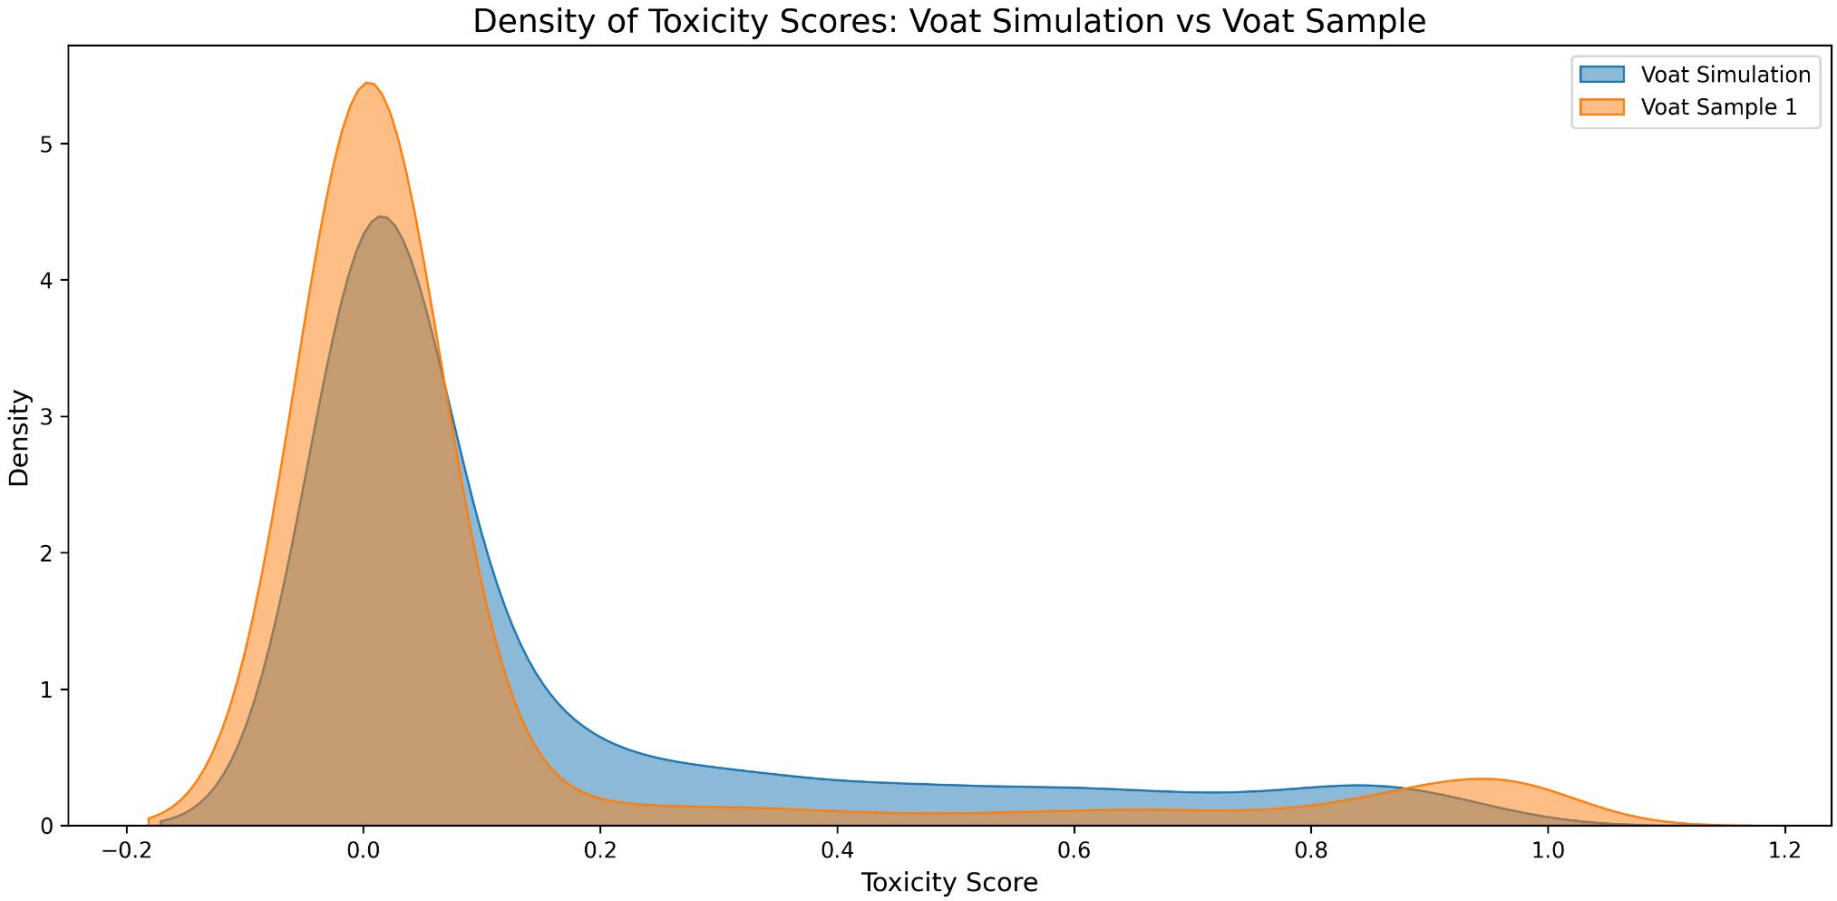
Density of Toxicity Scores: Voat Simulation vs Voat Sample
Voat Simulation
Voat Sample 1
Density
0
1
2
3
4
5
−0.2
0.0
0.2
0.4
0.6
0.8
1.0
1.2
Toxicity Score

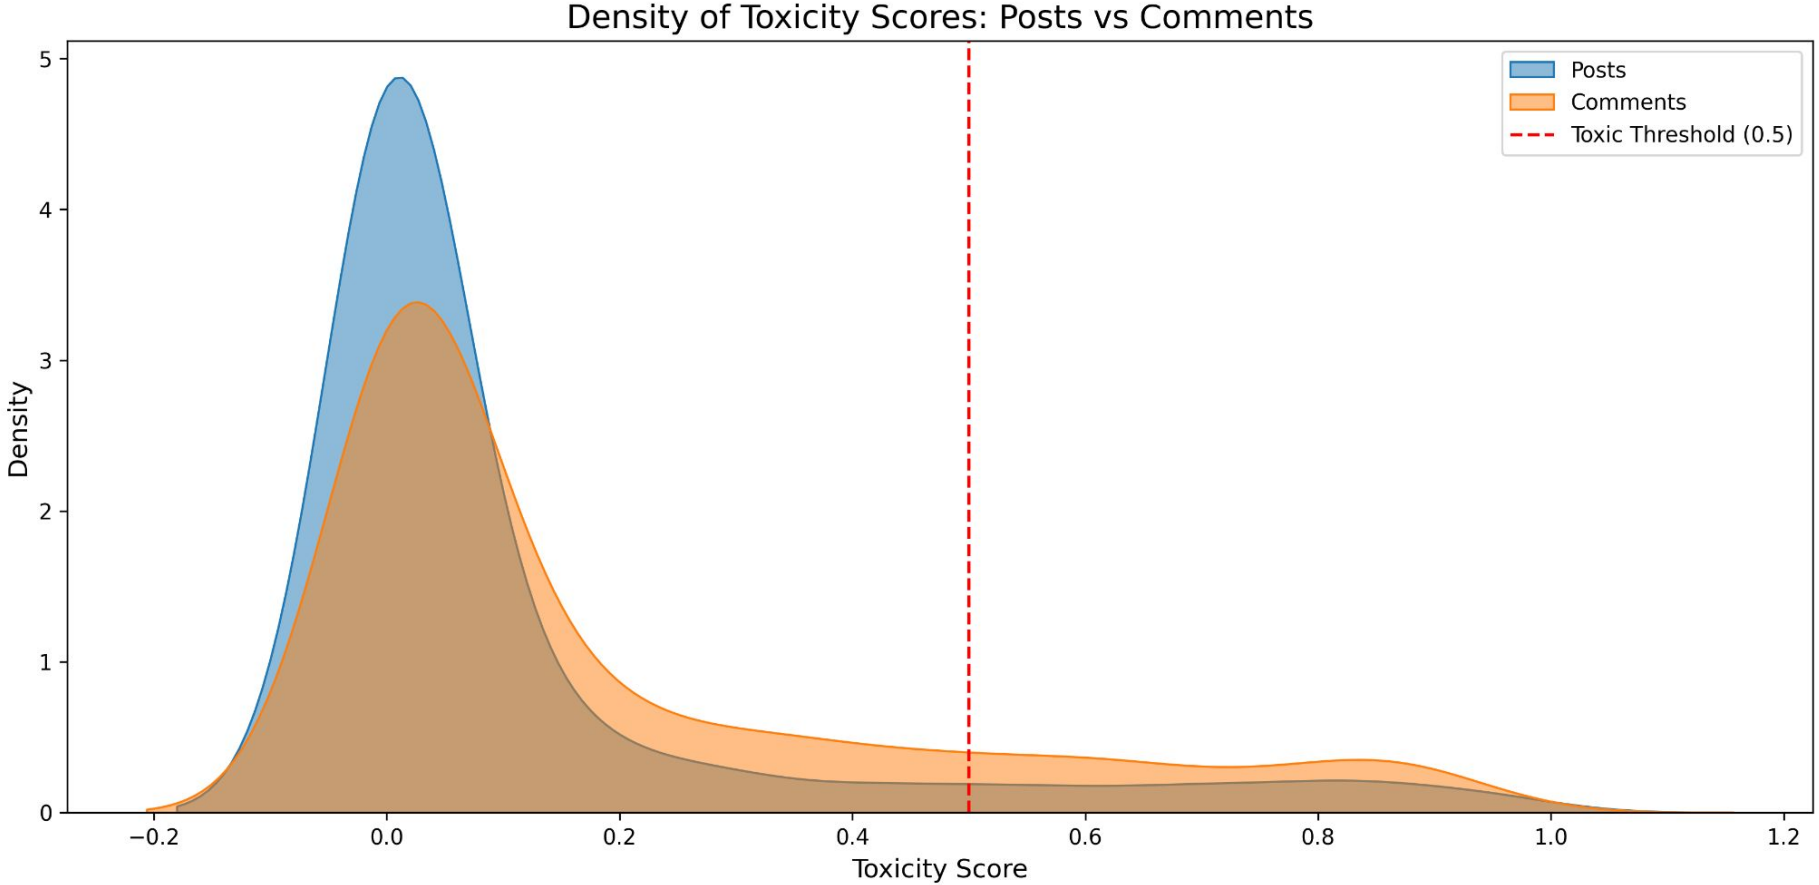
Density of Toxicity Scores: Posts vs Comments
Posts
Comments
Toxic Threshold (0.5)
Density
Toxicity Score
0
1
2
3
4
5
−0.2
0.0
0.2
0.4
0.6
0.8
1.0
1.2

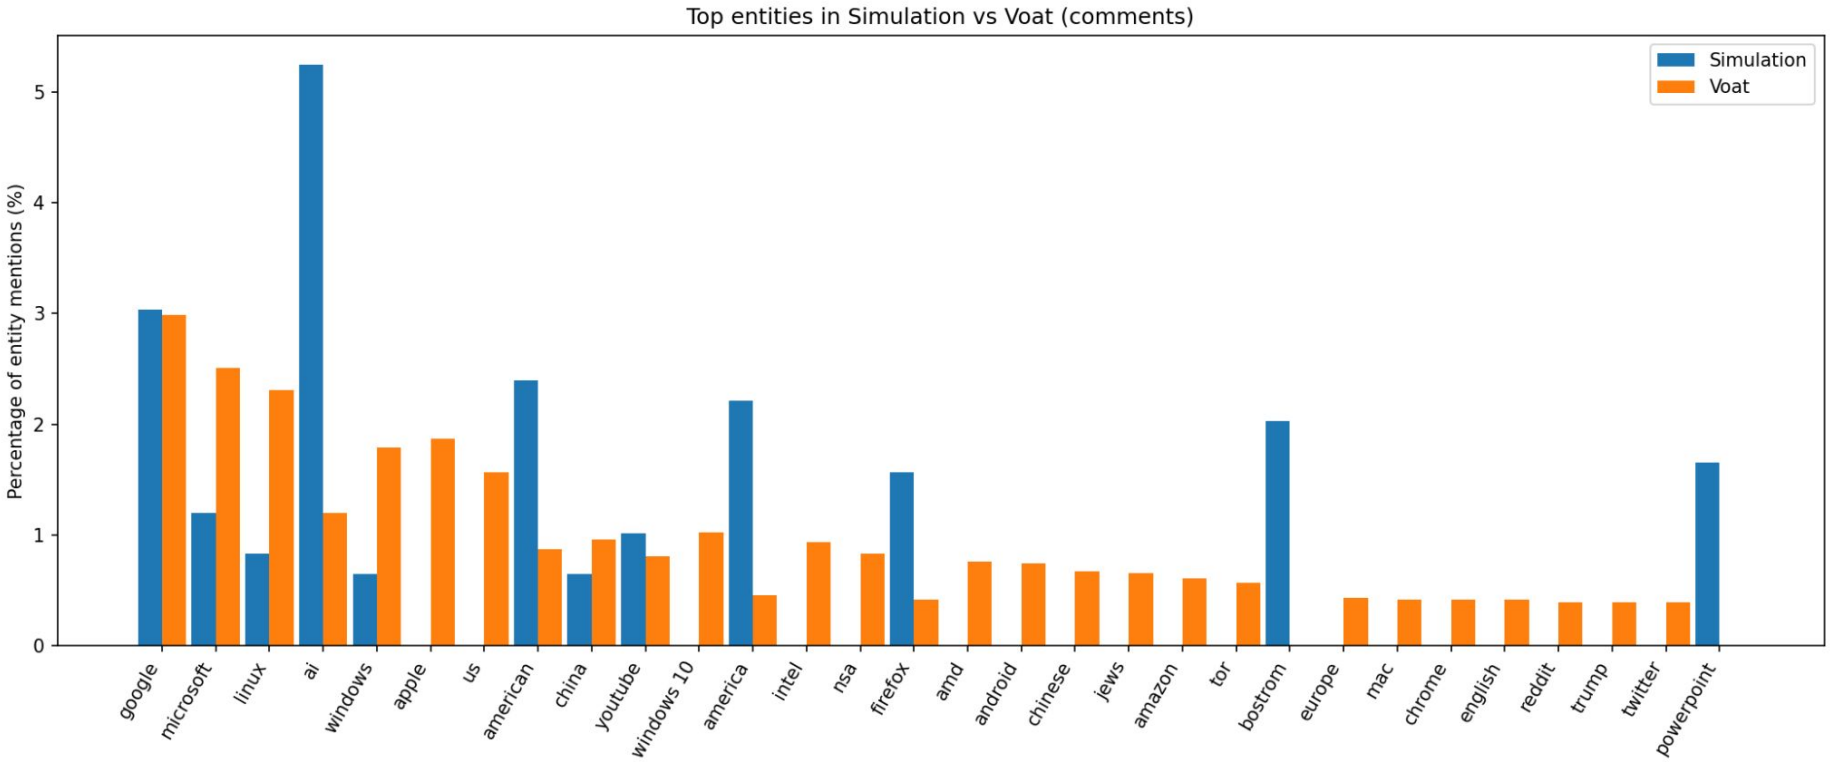
Top entities in Simulation vs Voat (comments)
Simulation
Voat
Percentage of entity mentions (%)
0
1
2
3
4
5
google
microsoft
linux
ai
windows
apple
us
american
china
youtube
windows 10
america
intel
nsa
firefox
amd
android
chinese
jews
amazon
tor
bostrom
europe
mac
chrome
english
reddit
trump
twitter
powerpoint

Simulation (posts+comments): Entity Co-occurrence

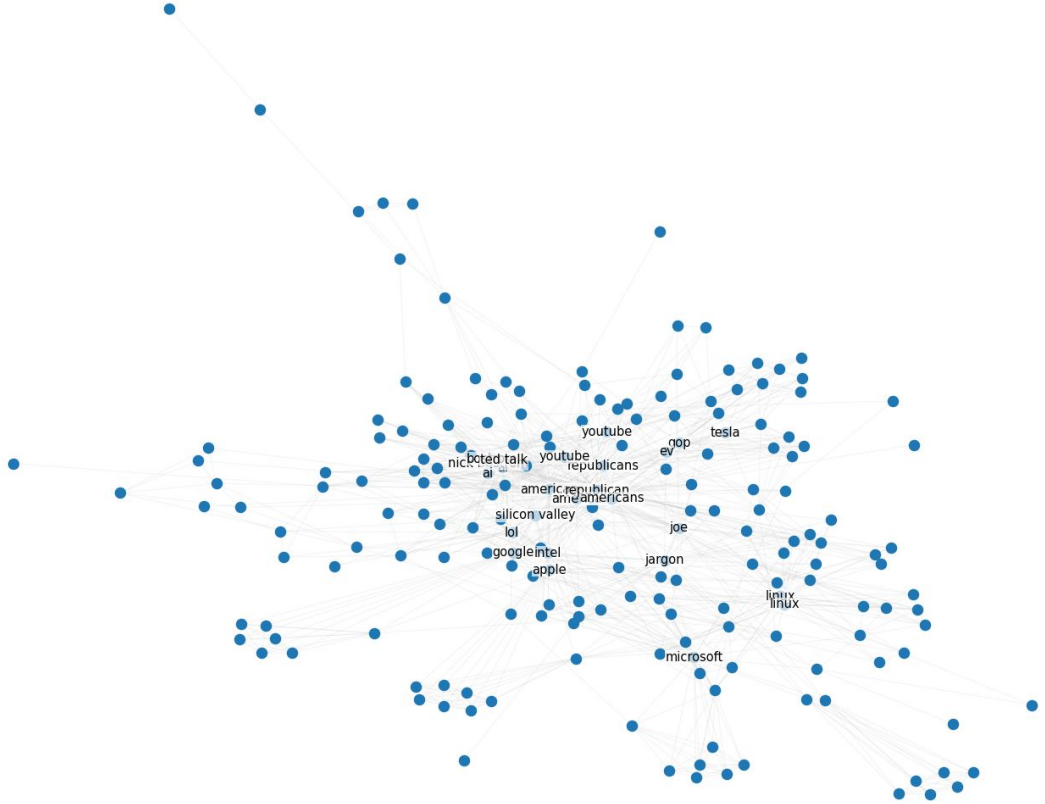

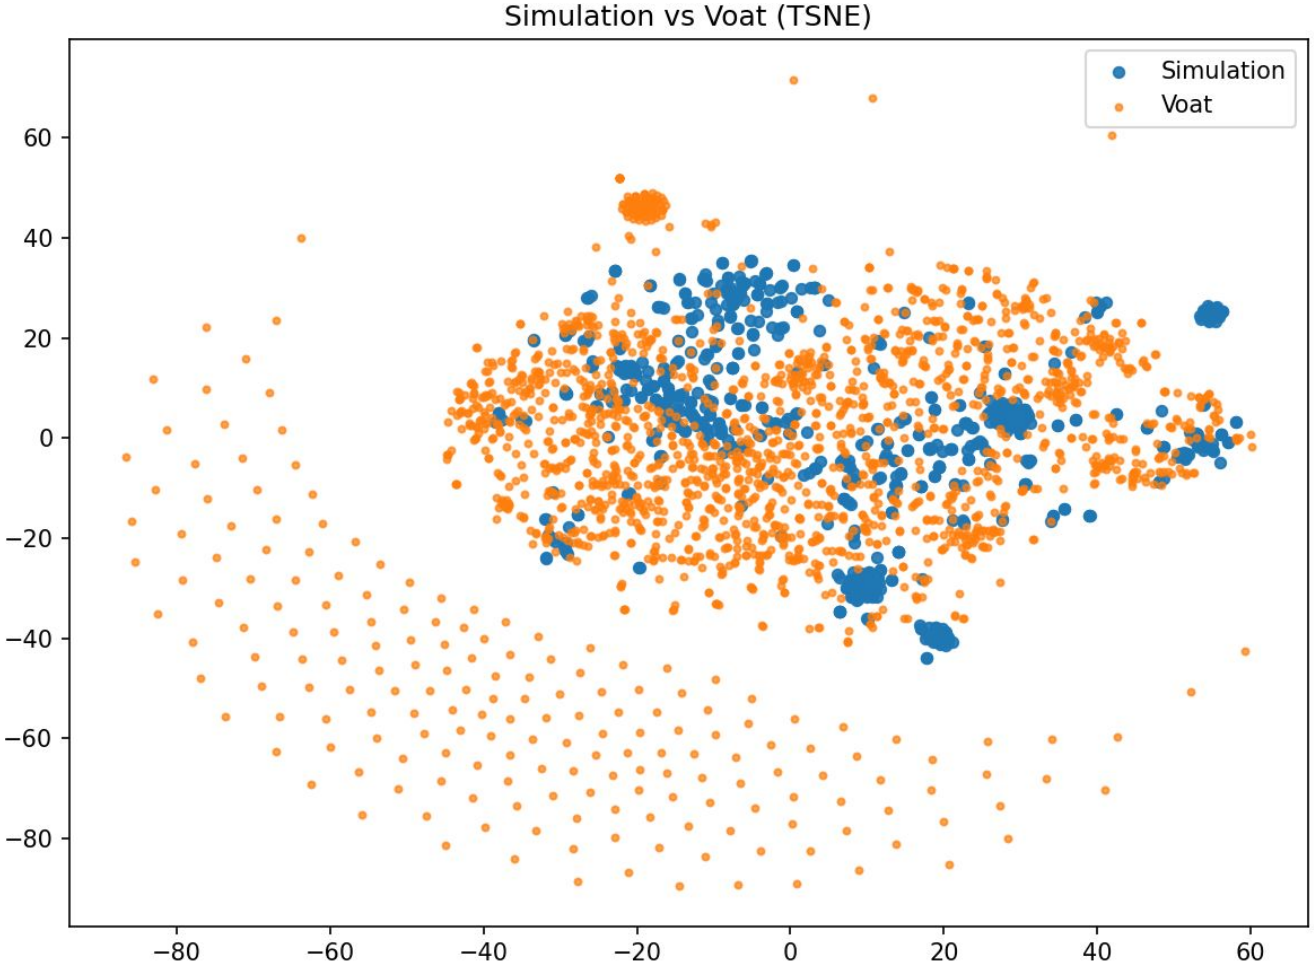
Simulation vs Voat (TSNE)
Simulation
Voat
60
40
20
0
−20
−40
−60
−80
−80
−60
−40
−20
0
20
40
60

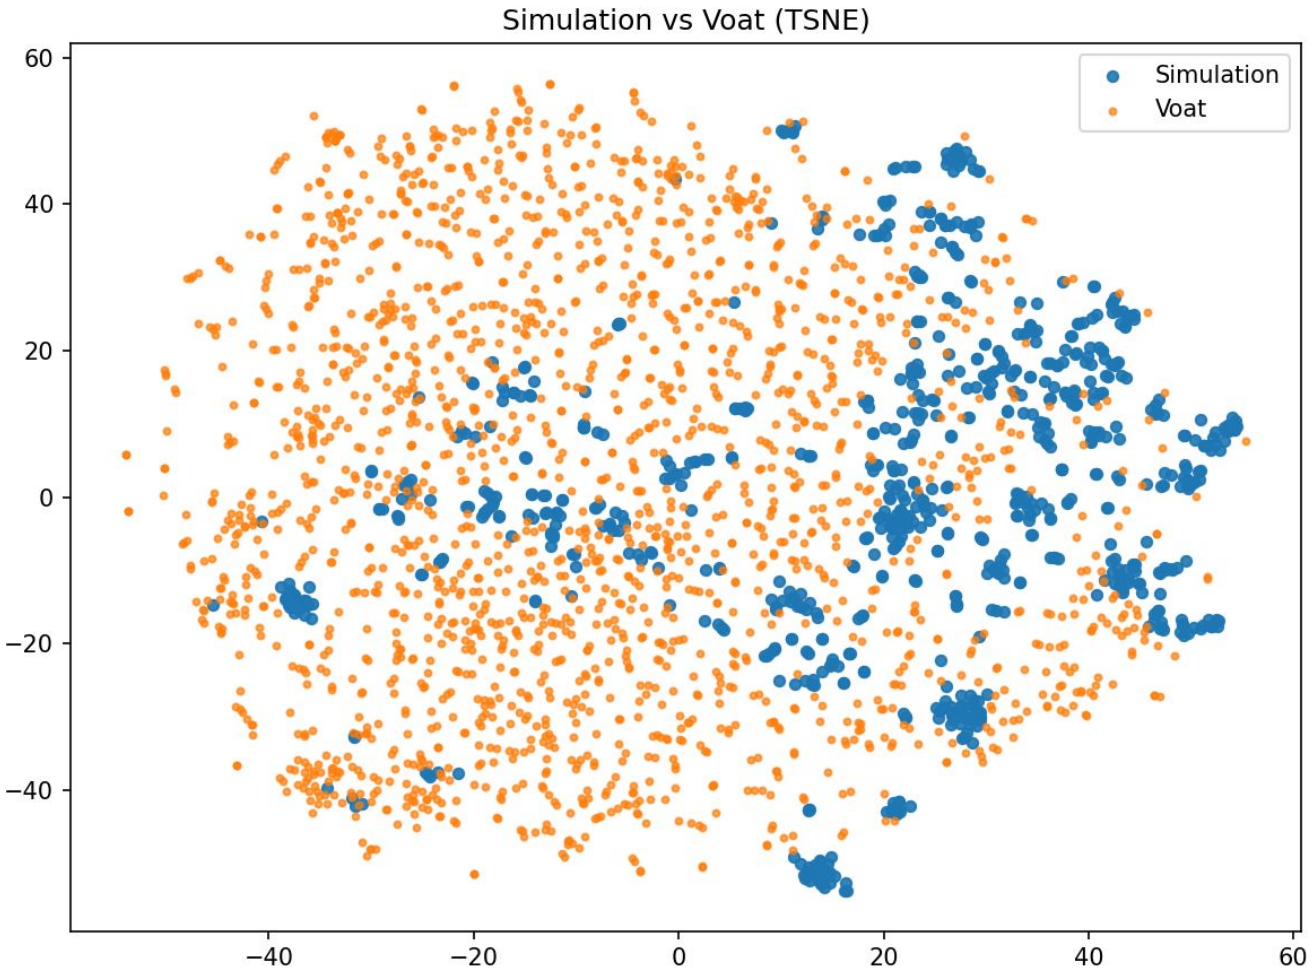
Simulation vs Voat (TSNE)
Simulation
Voat
60
40
20
0
−20
−40
−40
−20
0
20
40
60