
Tracking the spread of climate change skepticism on X with simulations and deep learning

Uwaila Ekhatior

School of Computing
Boise State University
Boise, ID 83725

uwailaekhatior@u.boisestate.edu

Mason Youngblood

Institute for Advanced Computational Science
Stony Brook University
Stony Brook, NY 11794
masonyoungblood@gmail.com

Vicken Hillis

Human-Environment Systems
Boise State University
Boise, ID 83725
vickenhillis@boisestate.edu

Abstract

Climate change continues to be a global challenge that requires urgent action. However, the ongoing presence of climate skepticism undermines society’s ability to confront this important challenge. Understanding the mechanisms driving the spread of climate skepticism might give policymakers additional tools to combat climate change. Here, we propose a methodological approach that combines computational simulation (in the form of an agent-based model representing online X communication) with simulation-based inference using amortized deep neural networks. Our approach allows us to infer the relative importance of a variety of different learning strategies that can contribute to the spread of climate skepticism and support.

1 Introduction

Despite ongoing climate change and ever-increasing impacts[1], substantial skepticism continues to exist, particularly in countries like the US [2]. While social media platforms like X (formally Twitter) have been instrumental in raising public awareness on climate change, the opinions on such platforms have become increasingly polarized [3] and can contribute to the spread of misinformation. [4]. Climate skepticism erodes the sense of urgency required for action[5], while polarization makes consensus policy more difficult to achieve. Understanding the mechanisms that drive the spread of climate skepticism on X is therefore a priority.

We propose a methodological approach combining agent-based simulation with deep-learning based inference to identify which social learning strategies (SLS) drive the spread of climate skepticism on X. SLS or biases, are the different ways in which individuals choose to learn from others[6]. These biases can take a variety of forms. For example, a content bias exists when the inherent characteristics of a tweet, such as negative emotional valence, influence its probability of being retweeted.[7]). Frequency bias exists when the prevalence of a tweet disproportionately influences its probability of being retweeted. A demonstrator bias causes individuals to preferentially retweet a tweet based on certain characteristics of the original tweeter and an ideological bias leads individuals to preferentially retweet something that aligns with their own ideological perspective. Understanding the relative importance of these (and possibly other types of learning biases) is important because their interaction

can lead to various emergent population-level outcomes such as the spread of misinformation or polarization [6].

The spread of climate change beliefs is a continuous process, so relying on static data analysis might miss critical temporal dynamics. Longitudinal network data with a full sequence of retweets, quotes or followers, would be required to use empirical data analysis alone, but such data is difficult to obtain. In addition, inferring SLS on social media can be challenging because exposures are unobserved and the different biases can produce similar patterns (equifinality). Our proposed method overcomes this by inferring the relative importance of different learning biases from the distributional data.

Our research objective is to quantify the relative importance of a variety of different biases within the climate change community on X. The insights from our work would explain why skepticism persists on X and offer guidance for social media platforms and climate communicators seeking to improve the quality of climate dialogue online.

2 Proposed Methodology

Our proposed methodology in figure 1 has two main stages: the Agent-Based Model (ABM) and Simulation-Based Inference (SBI), which we will explain in details below.




Figure 1: Proposed methodology workflow.

The dataset used in this study was curated by Brady et al. [8]. The data includes a large sample of climate change X data ($n = 507,358$).

2.1 ABM

We propose an ABM adapted from Youngblood et al. [7], which borrows elements from Carrignon et al. [9], Lachlan et al. [10], and Youngblood and Lahti [11]. Our ABM differs in that we include an ideology bias motivated by research that shows that climate change beliefs and discourse are politically polarized[12]. Our ABM simulates interactions between a population of X users with follower counts (T), activity levels (r), ideologies and probabilities of tweeting original tweets (μ) from real users in the observed data. μ is the proportion of a user’s total tweets and retweets that are original tweets, while r is the total number of original tweets and retweets by a user.

The ABM is initialized with a fully connected population of N users and is run for 101 timesteps, each of which corresponds to a 6-h interval in the real dataset (the highest resolution possible given computational limits). Each simulated user is assigned T , r , and μ corresponding to a real user in the observed data. In this way, we retain the correlation structure of follower count, activity, and retweet probability in our simulation model. This reflects the real variation in users’ behavior on X (e.g., in both contexts, there is a substantial number of users with few followers who exclusively retweet). The ABM is also initialized with a set of tweets with retweet frequencies drawn randomly from the first timestep in the observed data. Each tweet is also assigned a negativity index (M), computed using the VADER sentiment analysis tool [13]. At the start of each timestep, a pseudo-random subset of users becomes active (weighted by their values of r) and tweets according to the observed overall level of activity in the same timestep. Each active user either tweets original tweets or retweets existing tweets based on their unique value of μ . New original tweets are assigned (M) randomly sampled from the tweets in the first timestep.

The probability of retweeting each tweet is computed in a series of steps. First, the "baseline" attractiveness of each tweet (x) for retweeting (b_x) is computed using eqn (1) below:

$$b_x = s(\log F_x)^a \cdot s(\log T_x)^d \cdot s(M_x)^c \cdot s(\text{age}_x)^{-g} \quad (1)$$

F is the number of times that a tweet has been previously retweeted, and is log-transformed, normalized, and raised by the level of frequency bias (a). a is the same across all agents, where values < 1 simulate conformity bias and values > 1 simulate anticonformity bias. T is log-transformed, normalized, and raised by the level of follower influence (d). d is the same across all agents, where values of 0 simulate neutrality by removing variation in follower count and values > 0 simulate increasing levels of follower influence. M is normalized and raised by the level of content bias (c). c is the same across all agents, where values of 0 simulate neutrality by removing variation in the attractiveness of content and values > 0 simulate increasing levels of content bias. The final term simulates the decreasing probability that a tweet is retweeted as it ages, where g controls the rate of decay. Then, to reduce the computational cost of simulation, a pseudo-random sample of 5,000 candidate tweets (r) is collected, weighted by their attractiveness (b_x). $q \subset \{1, 2, \dots, n\}$, $|q| = 5000$. The ideological distance between each user and each of these tweets ($G_{u,y}$) is computed by taking the absolute difference between the ideology of the user and the ideology of the tweet, both of which are gotten from the observed data, normalized and raised by the level of ideological bias (k),

$$G_{u,y} = |\text{ideology}_u - \text{ideology}_y|, \quad \text{for each tweet } y \in q \quad (2)$$

and multiplied by the tweet’s attractiveness (b_y) to produce a retweet probability:

$$P_{u,y} = \frac{b_y \cdot s(G_{u,y})^k}{\sum_{z \in q} b_z \cdot s(G_{u,z})^k} \quad (3)$$

k is the same across all agents, where values of 0 simulate neutrality by removing variation in ideological distances and values > 0 simulate increasing levels of ideological bias. Once the active users are done tweeting and retweeting, each tweet increases in age by 1 and the next timestep begins. We provide working implementations of the ABM model in both C++ (for computational efficiency) and R in an anonymous repository.

We will run 55,000 iterations of the ABM, with parameter values sampled from uniform priors. Prior predictive tests will be conducted to choose which priors best cover the observed frequency distribution. The ABM produces five distributions ((i) the log-transformed frequency distribution of retweets, (ii) the density of negativity scores, (iii) the density of log-transformed follower counts, (iv) the density of ages, and (v) the density of ideology scores), constructed using kernel density estimation, weighted by frequency from all tweets in the population.

2.2 SBI

We will conduct SBI using BayesFlow in Python [14], a library that uses amortized deep neural networks to approximate the posterior distributions of parameters in a generative model [15]. BayesFlow out-competes approximate Bayesian computation [14] and has already been used to infer social learning strategies in other studies [16, 17]. The results from the ABM will be split into a training set ($n = 50,000$) and a validation set ($n = 5,000$). Then, we will run BayesFlow for 200 epochs using the Mish activation function and Adam optimizer with a batch size of 1000 and a learning rate of 0.001. Informative summary statistics are learned from the five output distributions using a hybrid convolutional-recurrent network, which is then mapped to parameter values. Finally, we will apply the trained neural network to produce 10,000 posterior predictions for each parameter in the ABM, based on the five distributions computed from the real data. The posterior distributions reveal the combinations of biases that best reproduce the retweet dynamics in the real dataset.

3 Impact and Conclusion

Our proposed approach will provide insights vital to online climate change communication. As noted in the introduction, climate skepticism is widespread on X, and understanding why it persists is crucial. By quantifying the strengths of the five SLS, our analysis will help to explain what makes climate retweets spread in climate change dialogues. We believe that these insights can guide policy makers, climate communicators, and media platform managers on how to better recommend content, frame climate messaging, and regulate the kind of climate change content being shared online. As one brief example, if our analysis uncovers a bias toward highly negative tweets (a finding consistent in other domains[7]), then content managers might develop stronger controls on negative posts, especially those with misinformation.

References

- [1] IPCC. *Summary for Policymakers*. IPCC, Geneva, Switzerland, 2023.
- [2] A. Leiserowitz, E. Maibach, S. Rosenthal, J. Kotcher, E. Goddard, J. Carman, M. Verner, T. Myers, J. Ettinger, J. Fine, E. Richards, M. Goldberg, and J. Marlon. Climate change in the american mind: Beliefs & attitudes, spring 2025. Report, Yale Program on Climate Change Communication, Yale University and George Mason University, New Haven, CT, 2025.
- [3] Mo Jones Jang and P. Hart. Polarized frames on “climate change” and “global warming” across countries and states: Evidence from twitter big data. *Global Environmental Change*, 32:11–17, 03 2015.
- [4] Yanmengqian Zhou and Lijiang Shen. Confirmation bias and the persistence of misinformation on climate change. *Communication Research*, 49(4):500–523, 2022.
- [5] Matthew J. Hornsey, Cassandra M. Chapman, and Jacquelyn E. Humphrey. Climate skepticism decreases when the planet gets hotter and conservative support wanes. *Global Environmental Change*, 74:102492, 2022.
- [6] Rachel L. Kendal, Neeltje J. Boogert, Luke Rendell, Kevin N. Laland, Mike Webster, and Patricia L. Jones. Social learning strategies: Bridge-building between fields. *Trends in Cognitive Sciences*, 22(7):651–665, 2018.
- [7] Mason Youngblood, Joseph M Stubbersfield, Olivier Morin, Ryan Glassman, and Alberto Acerbi. Negativity bias in the spread of voter fraud conspiracy theory tweets during the 2020 us election. *Humanities and Social Sciences Communications*, 10(1):1–11, 2023.
- [8] William J. Brady, Julian A. Wills, John T. Jost, Joshua A. Tucker, and Jay J. Van Bavel. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28):7313–7318, 2017.
- [9] Simon Carrignon, R. Alexander Bentley, and Damian Ruck. Modelling rapid online cultural transmission: evaluating neutral models on twitter data with approximate bayesian computation. *Palgrave Communications*, 5:83, 07 2019.
- [10] Robert F. Lachlan, Oliver Ratmann, and Stephen Nowicki. Cultural conformity generates extremely stable traditions in bird song. *Nature Communications*, 9, 2018.
- [11] Mason Youngblood and David Lahti. Content bias in the cultural evolution of house finch song. *Animal Behaviour*, 185:37–48, 2022.
- [12] Aaron M. McCright and Riley E. Dunlap. Cool dudes: The denial of climate change among conservative white males in the united states. *Global Environmental Change*, 21(4):1163–1172, 2011.
- [13] C. Hutto and Eric Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1):216–225, May 2014.
- [14] Stefan T. Radev, Ulf K. Mertens, Andreas Voss, Lynton Ardizzone, and Ullrich Köthe. BayesFlow: Learning complex stochastic models with invertible neural networks. *IEEE transactions on neural networks and learning systems*, 33(4):1452–1466, 2020.
- [15] Stefan T. Radev, Marvin Schmitt, Lukas Schumacher, Lasse Elsemüller, Valentin Pratz, Yannik Schälte, Ullrich Köthe, and Paul-Christian Bürkner. BayesFlow: Amortized Bayesian workflows with neural networks. *Journal of Open Source Software*, 8(89):5702, 2023.
- [16] Mason Youngblood and Sam Passmore. Simulation-based inference with deep learning suggests speed climbers combine innovation and copying to improve performance, Jun 2025.
- [17] Jay Pitocchelli, Adam Albina, R Alexander Bentley, David Guerra, and Mason Youngblood. Temporal stability in songs across the breeding range of *geothlypis philadelphia* (mourning warbler) may be due to learning fidelity and transmission biases. *Ornithology*, 142(1):ukae046, 09 2024.