Bursty Spatial and Temporal Activity Resulting from Social Search

Mahi Luthra and Peter M. Todd

Cognitive Science Program, Indiana University Bloomington (IUB)
Department of Psychological and Brain Sciences, IUB
mkluthra@iu.edu

Abstract

Social systems often display striking patterns of bursty spatial and temporal activity—for example, search activity for certain social resources (e.g., information) increases rapidly in quick bursts followed by lulls. Here we explore the role of social search (i.e., using social information to locate resources) in producing bursty spatial and temporal behavior using agent-based models where consumer agents search for resources. By varying the degree to which agents rely on other agents to locate resources and measuring the resulting burstiness, we find that as social search increases, so does burstiness (as indicated by previous literature); however, we also see that at higher levels of social search, burstiness appears to decline. In further simulations, we enable agents to evolve optimal social search values in response to the environment. Here we find that the social search values that evolve coincide with those that produce the highest burstiness. This suggests that burstiness is a characteristic of adaptive search strategies.

Introduction

Social search—or using social information to locate resources—is an effective search strategy evolved among several species. Humans, particularly, develop extensive social networks to find resources (e.g., information, food, and even other social connections). Prior research has suggested that such social networking produces striking patterns of bursty search activity in social systems—short sudden increases of search for particular resources followed by lulls (Karsai, Jo, & Kaski, 2018). For instance Ratkiewicz et al. (2010) found that Wikipedia pages experience bursts of popularity (i.e., number of visits over time) as people seek socially popular information. Detailed analyses find that such activity follows a heavy-tailed distribution (e.g., power law), indicating frequent occurrence of small amounts of activity and lower, but only slowly decreasing, likelihoods of larger activity. Similar burstiness patterns are found over time in popularity of baby names (Hahn & Bentley, 2003), dog breeds (Herzog, Bentley, & Hahn, 2004), terms of speech (Lappas et al., 2012) and other informational domains where social search is used. In addition to such burstiness occurring across time, studies have also found evidence of spatial burstiness (Fang et al., 2018), where small patches of informational landscapes experience large spikes of search activity while other regions are largely underutilized (e.g., at a particular timepoint, only a few clothing styles are trendy while most others are "out of fashion").

Although extensive previous research has focused on modelling and simulating existing dynamics of bursty social systems, few studies have systematically explored the role of social search in producing such activity. In the current work, we fill this gap using agent-based models where consumer agents search for attractive resources. By varying the degree to which agents rely on social information (i.e., locations of other consumers) to detect resources, we analyze the influence of social search on spatial and temporal burstiness.

We also run simulations where social search evolves across time to see how burstiness emerges. Here, social search parameter values are not fixed—rather, they can evolve in response to the evolving social search of other agents and characteristics of the environment (and burstiness of the environment in turn changes based on the search strategies, creating a mutual interaction of search and environment; Luthra, Izquierdo, & Todd, 2020). These simulations are representative of real-world search where humans and other species can vary the degree to which they rely on social information (e.g., by learning or evolving better social search strategies) while interacting with a resource domain.

Methods

Agents search for resources in 100×100 2D landscapes. On every timestep, they may gain energy by consuming resources located near them. To achieve this, agents can perceive resources and other searching agents and by attaching weights to each they can make decisions to move toward or away from them. In our simulations, resource weights are always fixed at +1 (therefore they always have a predisposition to move toward resources) while weights attached to other consumer agents can vary between +1 and -1. Social search here refers to the weights attached to these other agents. Therefore, in positive social search ($0<\text{social search}<+1$), agents move toward other agents, who might serve as indicators of attractive landscape regions. Under negligible social search (~0), agents detect only resources and move independent of other agents. In negative social search ($-1<\text{social search}<0$), agents move away from other agents (potentially to avoid competition).

Depending on the energy agents accrue (by consuming resources), they may reproduce (if energy crosses a high threshold) or die (if energy falls below 0). In fixed social search simulations, social search parameters of all agents are fixed at various levels across experimental conditions to analyze its influence on burstiness. However, in evolving social search,

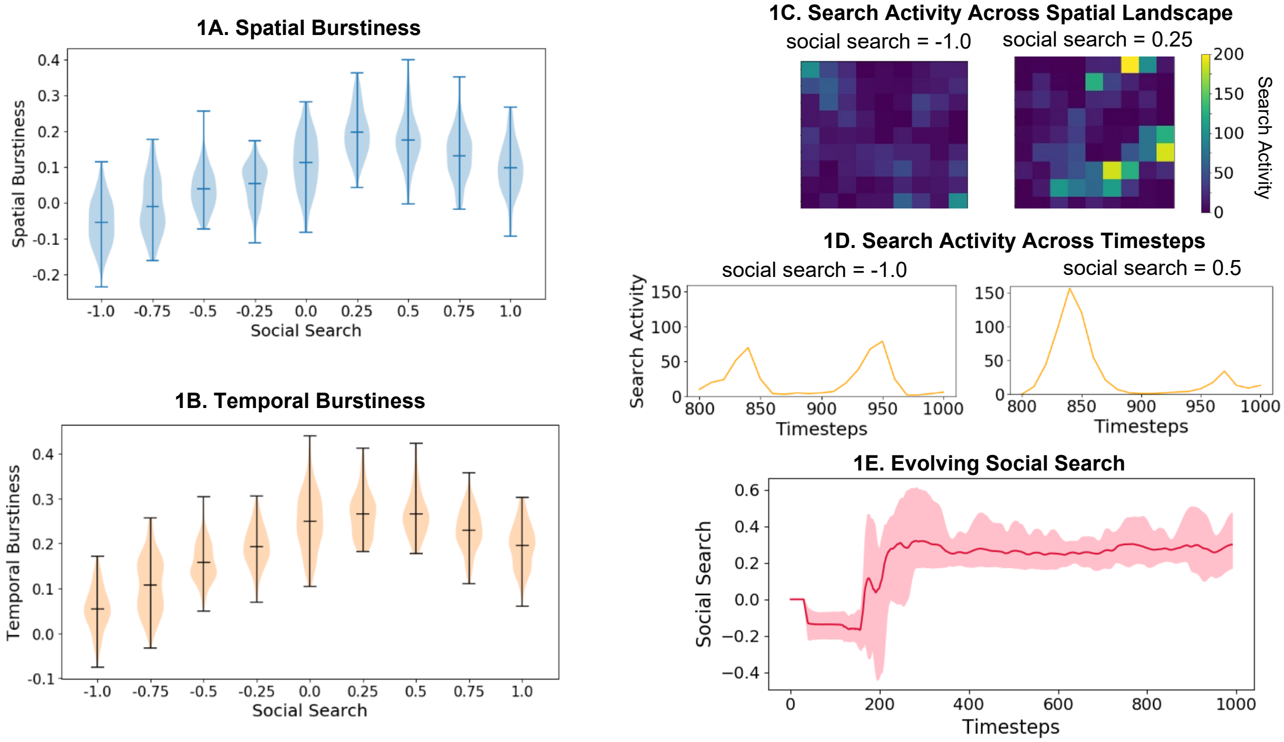

Figure 1: 1A. Distributions of spatial burstiness across social search conditions. Burstiness measure was applied to search activity across 10×10 sections of the landscape at each timestep. Violin plots represent 2000 datapoints of burstiness values (last 200 timesteps $\times$ 10 simulation runs) 1B. Distributions of temporal burstiness across social search conditions. Burstiness measure was applied to search activity across 200 timesteps for each 10×10 landscape section. Violin plots represent 1000 datapoints of burstiness values (100 sections $\times$ 10 simulation runs). 1C. Search activity in an entire landscape at timestep 800 for social search conditions -1.0 and 0.25 each in a single simulation run. 1D. Search activity across timesteps of a single landscape section for social search conditions -1.0 and 0.5 each in a single simulation run. 1E. Aggregate social search parameters evolved across 10 simulation runs.

agents pass on their social search parameters to their offspring with some mutation. This enables social search to evolve across time to optimal values in response to the environment. Our recent work (Luthra and Todd, in press) provides detailed descriptions of the model for further reference.

To measure burstiness across simulations, we used Goh and Barabási's (2008) measure:

$$B = \frac{\sigma - \mu}{\sigma + \mu}$$

where σ and μ denote the standard deviation and mean of search activity. For spatial burstiness, we estimated search activity, defined as number of agents searching within an area, within 10×10 unit sections of the resource landscape and at each timestep applied the burstiness measure across these spatial search activity values. For temporal burstiness, we applied the burstiness measure to search activity collated across time for each 10×10 section. Simulations were run for 1000 timesteps and results of the last 200 timesteps for 10 simulations are reported.

Results

Fixed Social Search

Social search was fixed at 0.25 intervals between -1 and +1 and resulting burstiness was measured. Figures 1A and 1B display spatial and temporal burstiness respectively across social search weights. We find that as social search increases, spatial and temporal burstiness also generally increase. However, spatial burstiness appears to peak at social search=0.25 after which it declines; while temporal burstiness peaks at social search=0.5 and then declines. Figure 1C displays the search activity of the landscape at timestep 800 of one simulation run for each of the social search conditions producing the lowest and highest burstiness (-1 and 0.25). For social search=0.25 a few landscape sections have extreme search activity while others have low activity; for social search=-1, most sections have a similar level of search activity. Figure 1D displays search activity across time in one simulation with social search=0.5 showing bursty peaks, while another run with social search=-1 has smaller peaks. We also fit all search activity values with both power-law and less heavy-tailed exponential distributions (across time and space), finding that the former provided better fits of data for social search ≥ 0 while the latter fit simulations with social search ≤ -0.25 better.

Evolving Social Search

Figure 1E displays social search parameters evolving through time (aggregated across 10 simulations). Social search tends to stabilize around 0.27. This is close to the values that produced the highest spatial and temporal burstiness in the fixed social search simulations (Figures 1A and 1B), which also correspond to the aggregate spatial and temporal burstiness levels in these simulations of 0.21 and 0.26, respectively.

Discussion

The current study systematically analyzed the influence of social search on burstiness, finding an inverse U-shaped relationship between the two. Initially, as social search increased, so did burstiness. This can be attributed to the rich-get-richer effect of social search—as agents accumulate in attractive resource regions, more agents follow, producing sudden bursts (Ratkiewicz et al, 2010). However, at higher values of social search, there are declines in burstiness. This could potentially be due to an equalizing effect of high social search—as agents chase after one another, they consume resources inefficiently across the landscape which might result in more even distributions of resources and hence, more even distributions of clustered agents (since they are simultaneously attracted to resources). Future work should closely analyze the dynamics of social aggregation at higher social search values.

When social search evolved across time, we found that agents tended to evolve parameter values around 0.27. Interestingly, these values are close to those producing the highest spatial and temporal burstiness in the fixed social-search simulations. Potentially, social systems with high burstiness are adaptive—they allow agents to optimally exploit popular resource locations during bursts while exploring the landscape more evenly during lulls to find regions of further fertility, effectively balancing the exploration/exploitation trade-off (Hills et al., 2015).

Supplemental Materials

Supplemental materials and code for this project are available online at https://github.com/mahiluthra/social-search-clustering

Acknowledgments

This research was supported in part by the John Templeton Foundation grant, “What drives human cognitive evolution”.

References

Fang, A., Ounis, I., Macdonald, C., Habel, P., Xiong, X., & Yu, H.-T. (2018). An effective approach for modelling time features for classifying bursty topics on twitter. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (pp. 1547–1550). ACM.

Goh, K. I., & Barabási, A. (2008). Burstiness and memory in complex systems. *A Letters Journal Exploring the Frontiers of Physics*, *81*, 48002.

Hahn, M. W., & Bentley, R. A. (2003). Drift as a mechanism for cultural change: an example from baby names. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *270*, S120–S123.

Herzog, H. A., Bentley, R. A., & Hahn, M. W. (2004). Random drift and large shifts in popularity of dog breeds. *Proceedings of the Royal Society: Biology Letters*, *271*, S353–S356.

Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., Bateson, M., … Wolfe, J. W. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences, 19(1)*, 46-54.

Karsai, M., Jo, H.-H., & Kaski, K. (2018). *Bursty human dynamics*. Springer International Publication.

Lappas, T., Vieira, M. R., Gunopulos, D., & Tsotras, V. J. (2012). On the spatiotemporal burstiness of terms. *Proceedings of the VLDB Endowment*, *5*(9), 836–847.

Luthra, M., Izquierdo, E. J., & Todd, P. M. (2020). Cognition evolves with the emergence of environmental patchiness. In J. Bongard, J. Lovato, L. Herbert-Dufresne, R. Dasari, & L. Soros (Eds.), *Artificial Life Conference Proceedings* (pp. 450–458). MIT Press.

Luthra, M. & Todd, P. M. (in press). Social search evolves with the emergence of clustered environments. In *Artificial Life Conference Proceedings.* MIT Press.

Ratkiewicz, J., Fortunato, S., Flammini, A., Menczer, F., & Vespignani, A. (2010). Characterizing and modeling the dynamics of online popularity. *Physical Review Letters*, *105*, 158701.

Ratkiewicz, J., Menczer, F., Fortunato, S., Flammini, A., & Vespignani, A. (2010). Traffic in social media II: Modeling bursty popularity. In *IEEE Second International Conference on Social Computing* (pp. 393–400).