Geospatial Active Search for Preventing Evictions

Extended Abstract

Anindya Sarkar*
Washington University in St. Louis
St. Louis, United States
anindyasarkar.ece@gmail.com

Alex DiChristofano*
Washington University in St. Louis
St. Louis, United States
a.dichristofano@wustl.edu

Sanmay Das
George Mason University
Fairfax, United States
sanmay@gmu.edu

Patrick J. Fowler
Washington University in St. Louis
St. Louis, United States
pjfowler@wustl.edu

Nathan Jacobs
Washington University in St. Louis
St. Louis, United States
jacobsn@wustl.edu

Yevgeniy Vorobeychik
Washington University in St. Louis
St. Louis, United States
yvorobeychik@wustl.edu

ABSTRACT

Evictions are a threat to housing stability and a major concern for many cities. An open question is whether data-driven methods can enhance door-to-door outreach programs to target at-risk tenants. We model this problem using a new framework we term geospatial active search. Geospatial Active Search integrates visual information such as satellite imagery along with tabular data such as property and neighborhood-level information to create an online exploration plan. We develop an approach for the implementation of Geospatial Active Search in St. Louis to find properties containing tenants who will have an eviction filed against them.

CCS CONCEPTS

• **Computing methodologies → Active learning settings**; **Sequential decision making**; *Computer vision*; *Neural networks*.

KEYWORDS

active search; reinforcement learning; eviction; housing; geospatial

ACM Reference Format:
Anindya Sarkar*, Alex DiChristofano*, Sanmay Das, Patrick J. Fowler, Nathan Jacobs, and Yevgeniy Vorobeychik. 2024. Geospatial Active Search for Preventing Evictions: Extended Abstract. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

1 INTRODUCTION

This work establishes the problem of Geospatial Active Search (GAS) and solves it in the context of tenant outreach. GAS projects points in the active search problem to a spatial domain, and an agent queries these points by traveling between them. The effort required for the agent to move between two points is modeled as the cost of querying. A trained GAS model is a plan which, given the current state of exploration, returns the next point to query, and, when a point is queried, updates the model, estimating which points are members of the target class. Our contributions are: (1) We propose a real-world problem of geospatial exploration; (2) We define this problem using a framework we term Geospatial Active Search; (3) We construct an end-to-end pipeline to solve GAS that leverages a pretraining phase while allowing for test-time adaptation; (4) We show the potential of GAS by applying it to the task of uncovering properties with tenants at risk of eviction in St. Louis.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

2 RELATED WORK

Geospatial Active Search builds on Active Search, first proposed by Garnett et al.. Previous work on Active Search has focused on developing nonmyopic algorithms [14, 16, 17] and minimizing the cost of discovery of a set number of targets [13]. Recent approaches learn policies that adapt at test time using only visual data [18, 19].

Our work is aligned with Mashiat et al., who use the same tabular data for offline eviction prediction. It also builds on a larger body of literature focusing on geospatial applications of optimization and artificial intelligence in humanitarian domains. This includes solutions for collaborative recycling [12], the redistribution of food donations [3], the routing of disaster relief [8], predicting micronutrient deficiency [7], and anti-poaching measures [4–6, 9, 10, 21].

3 PROBLEM AND PROPOSED APPROACH

A geospatial search task consists, in part, of properties embedded as points in a geographic region. The classes these points belong to are not known apriori but can be discovered via sequential queries. The goal is to uncover as many properties belonging to a prespecified class as possible. Formally, our search task involves K parcels, defined as $x = [x_1, \ldots, x_K]$. We formalize this by associating each parcel j with a binary label $y^{(j)} \in \{0, 1\}$, where $y^{(j)} = 1$ iff parcel index j contains a parcel with future eviction filing. When we inquire about parcel i, we both acquire the corresponding label $y^{(i)}$ and gain utility if at least one eviction is filed at parcel i within the next three months. Denoting a query performed in step t as q_t and $c(i, j)$ as the cost associated with querying parcel j from parcel i, our optimization objective is:

$$\max_{\{q_t\}} U(x; \{q_t\}) \equiv \sum_t y^{(q_t)} \text{s.t.} : \sum_{t \geq 0} c(q_{t-1}, q_t) \leq C \qquad (1)$$

With this objective in mind, we suggest an approach for training a search policy that acquires the skill of efficient exploration through learning from similar pre-labeled search tasks. Following previous work [19], we model multi-modal active search as a Budget

Table 1: Average Number of Targets (ANT) Found by Search Task Parameters and Solution Method

	Average Positive Rate of 2.5%			Average Positive Rate of 5%			Average Positive Rate of 10%		
Search Budget	15	20	25	15	20	25	15	20	25
Random	0.328	0.472	0.652	0.656	0.856	1.276	1.428	1.940	2.428
Greedy	0.540	0.632	0.804	1.660	2.040	2.308	3.796	4.296	4.780
Greedy Adaptive	0.556	0.716	0.828	1.788	2.124	2.400	3.872	4.452	5.020
GAS-TAB (no images)	**1.008**	**1.184**	**1.340**	**1.972**	**2.336**	**2.640**	**4.216**	**5.004**	**5.501**
GAS	**1.052**	**1.288**	**1.352**	**2.240**	**2.596**	**2.828**	**4.372**	**5.204**	**5.632**

Constraint MDP. In this MDP, the **input state** at time t includes: 1) the multi-modal input feature corresponds to K parcel, denoted as x^t, 2) the outcomes of past search queries o^t, and 3) the remaining budget $B^t \leq C$. Since we possess visual data for each individual parcel as well as tabular features containing past eviction records for the corresponding parcels, we utilize a widely adopted multi-modal transformer architecture as described in [20]. Each element of search query history o corresponds to a parcel index j, so that $o^j = 2y^{(j)} - 1$. In this MDP, the **actions** are simply choices over which parcels to query next. We assign an **immediate reward** for query a parcel j as $R(j) = y^{(j)}$. The **state transition** process involves both updating the remaining budget and incorporating the result of the most recent search query into the state. Armed with this MDP problem representation, we next describe our proposed deep RL approach for learning a search policy that makes use of a dataset D of past search tasks. Specifically, we use the REINFORCE to directly learn a search policy $\psi_\theta(x, o, B)$. In order to utilize the valuable information we acquire during search, we develop a search policy comprised of two key components: 1) the prediction module represented by $f_\phi(x, o)$ and 2) the search module denoted as $g_\zeta(p, o, B)$, where ϕ and ζ represent trainable parameters, where $p = f_\phi(x, o)$ is the vector of predicted eviction probabilities with $p^{(j)}$ the predicted probability of at least one eviction in parcel index j. In a conceptual sense, f_ϕ generates predictions by exclusively considering the task features x and previous search outcomes o, whereas g_ζ depends solely on information pertinent to the search process itself, including the predicted eviction probabilities p, o and B. The resulting search policy is a combination of these modules, expressed as $\psi(x, o, B) = g_\zeta(f_\phi(x, o), o, B)$. Our ultimate objective is to train both g_ζ and f_ϕ in a manner where f_ϕ supports adaptive search by continuously updating itself during a search task. To address this, at the start of each task, ϕ is set to ϕ_{start} and is *subsequently updated as labels become available after each query* during both training and inference using BCE loss. In the case of the search module, we calculate the cumulative sum of rewards $R_r = \sum_j y^{(j)}$ for the parcels j explored during the episode and employ the RL loss $\mathcal{L}_{RL}$ based on the REINFORCE algorithm. The proposed approach balances the RL and supervised loss through the loss function: $\mathcal{L}_{GAS} = (\mathcal{L}_{RL} + \lambda\mathcal{L}_{BCE})$

4 EXPERIMENTS AND RESULTS

Utilizing real eviction filing data from 2021 and 2022, we implement GAS to find properties at risk of eviction, defined as having a filing in the next three months. We assess our models using the average number of targets found (ANT), which is defined as the targets found across all tasks divided by the number of search tasks. We consider uniform query costs: $c(i, j) = 1 \;\forall\; i, j$, where C is the number of queries. While this analysis focuses on uniform costs, our framework supports non-uniform query costs without any further modifications.

Baseline Search Methods: We evaluate the performance of GAS using the following baselines: ***random***, where we randomly choose parcels from those unexplored; ***greedy***, in which we train a classifier ψ_{greedy} to predict eviction risk and search highest-risk properties first; and ***greedy adaptive***, similar to *greedy* but with updates to ψ_{greedy} at each time step based on query outcomes.

Data: The tabular features used are drawn from a companion study and are originally derived from municipal sources across St. Louis [15]. They cover eviction court filings, owner information, and property-level attributes. Neighborhood features are obtained from the American Community Survey [1]. Satellite imagery data comes from the National Agriculture Imagery Program (NAIP) [2].

Search Task Design: We construct a series of search tasks, each containing 100 properties. We bootstrap properties from the original 26,700 and control for the positive rate. The mean positive rates are 2.5% (σ = 0.005%) (near the true base rate), 5% (σ = 0.01%), and 10% (σ = 0.02%). Experiments are run for budgets of 15, 20, and 25 queries, and each parameterization is trained and tested on 250 search tasks.

Results: We evaluate GAS against baselines and present the results in Table 1. We observe a substantial improvement using GAS compared to all baselines, ranging from approximately 16% - 90% when compared to the most competitive approach, *greedy adaptive*. We also observe that *greedy adaptive* outperforms the *greedy* method across all experiments. This highlights the importance of inference time adaptive strategies for optimizing search efficiency. The use of satellite imagery is supported by an ablation study, where GAS with tabular and visual data (**GAS**) outperforms GAS using only tabular data (**GAS-TAB**). Finally, the performance gap between our proposed method and the baseline approaches is more significant when evictions are scarce and resources are limited, which reflects the real world.

ACKNOWLEDGMENTS

This project was partially supported by the NSF (IIS-1905558, IIS-1903207, IIS-2214141, IIS-1939677, IIS-2127752), ARO (W911NF1810208), Amazon, NVIDIA, and JP Morgan Chase. We thank the St. Louis Regional Response Team and members of the collaborative network for ongoing assistance in the research design.

ETHICS STATEMENT

While the work presented here is not yet ready to be deployed in the field, any future outreach performed using these methods must be carefully considered. When answering the question of who gets outreach, one inevitably decides who doesn't. St. Louis has a long history of housing discrimination, and without careful analysis of mechanisms for bias, work in this domain risks repeating errors of the past. Yet, despite this, approaches such as the one presented here offer promise in being able to more adeptly anticipate those in need and those who will be able to benefit the most from intervention.

REFERENCES

[1] 2021. American Community Survey 5-Year Estimates, 2021. Data retrieved from the U.S. Census Bureau website. https://data.census.gov/cedsci/ Accessed on October 6, 2023.

[2] 2022. National Agriculture Imagery Program - NAIP Hub Site. https://naip-usdaonline.hub.arcgis.com/.

[3] Burcu Balcik, Seyed Iravani, and Karen Smilowitz. 2014. Multi-Vehicle Sequential Resource Allocation for a Nonprofit Distribution System. *IIE Transactions* 46, 12 (Dec. 2014), 1279–1297. https://doi.org/10.1080/0740817X.2013.876240

[4] Elizabeth Bondi, Debadeepta Dey, Ashish Kapoor, Jim Piavis, Shital Shah, Fei Fang, Bistra Dilkina, Robert Hannaford, Arvind Iyer, Lucas Joppa, and Milind Tambe. 2018. AirSim-W: A Simulation Environment for Wildlife Conservation with UAVs. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3209811.3209880

[5] Elizabeth Bondi, Fei Fang, Mark Hamilton, Debarun Kar, Donnabell Dmello, Jongmoo Choi, Robert Hannaford, Arvind Iyer, Lucas Joppa, Milind Tambe, and Ram Nevatia. 2018. SPOT Poachers in Action: Augmenting Conservation Drones With Automatic Detection in Near Real Time. *Proceedings of the AAAI Conference on Artificial Intelligence* 32, 1 (April 2018). https://doi.org/10.1609/aaai.v32i1.11414

[6] Elizabeth Bondi, Raghav Jain, Palash Aggrawal, Saket Anand, Robert Hannaford, Ashish Kapoor, Jim Piavis, Shital Shah, Lucas Joppa, Bistra Dilkina, and Milind Tambe. 2020. BIRDSAI: A Dataset for Detection and Tracking in Aerial Thermal Infrared Videos. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 1736–1745. https://doi.org/10.1109/WACV45572.2020.9093284

[7] Elizabeth Bondi-Kelly, Haipeng Chen, Christopher D. Golden, Nikhil Behari, and Milind Tambe. 2023. Predicting Micronutrient Deficiency with Publicly Available Satellite Data. *AI Magazine* 44, 1 (2023), 30–40. https://doi.org/10.1002/aaai.12080

[8] Luis E. de la Torre, Irina S. Dolinskaya, and Karen R. Smilowitz. 2012. Disaster Relief Routing: Integrating Research and Practice. *Socio-Economic Planning Sciences* 46, 1 (March 2012), 88–97. https://doi.org/10.1016/j.seps.2011.06.001

[9] Fei Fang, Thanh Nguyen, Rob Pickles, Wai Lam, Gopalasamy Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. 2016. Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security. *Proceedings of the AAAI Conference on Artificial Intelligence* 30, 2 (Feb. 2016), 3966–3973. https://doi.org/10.1609/aaai.v30i2.19070

[10] Fei Fang, Peter Stone, and Milind Tambe. 2015. When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI'15)*. AAAI Press, Buenos Aires, Argentina, 2589–2595.

[11] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. 2012. Bayesian Optimal Active Search and Surveying. In *Proceedings of the 29th International Coference on International Conference on Machine Learning (ICML'12)*. Omnipress, Madison, WI, USA, 843–850.

[12] Vera Hemmelmayr, Karen Smilowitz, and Luis de la Torre. 2017. A Periodic Location Routing Problem for Collaborative Recycling. *IISE Transactions* 49, 4 (April 2017), 414–428. https://doi.org/10.1080/24725854.2016.1267882

[13] Shali Jiang, Roman Garnett, and Benjamin Moseley. 2019. Cost Effective Active Search. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc.

[14] Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, and Roman Garnett. 2017. Efficient Nonmyopic Active Search. In *Proceedings of the 34th International Conference on Machine Learning*. PMLR, 1714–1723.

[15] Tasfia Mashiat, Alex DiChristofano, Patrick J. Fowler, and Sanmay Das. 2024. Beyond Eviction Prediction: Leveraging Local Spatiotemporal Public Records to Inform Action. https://doi.org/10.48550/arXiv.2401.16440 arXiv:2401.16440 [cs]

[16] Quan Nguyen and Roman Garnett. 2023. Nonmyopic Multiclass Active Search with Diminishing Returns for Diverse Discovery. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*. PMLR, 5231–5249.

[17] Quan Nguyen, Arghavan Modiri, and Roman Garnett. 2021. Nonmyopic Multifidelity Acitve Search. In *Proceedings of the 38th International Conference on Machine Learning*. PMLR, 8109–8118.

[18] Anindya Sarkar, Nathan Jacobs, and Yevgeniy Vorobeychik. 2023. A Partially Supervised Reinforcement Learning Framework for Visual Active Search. *arXiv preprint arXiv:2310.09689* (2023).

[19] Anindya Sarkar, Michael Lanier, Scott Alfeld, Roman Garnett, Nathan Jacobs, and Yevgeniy Vorobeychik. 2022. A Visual Active Search Framework for Geospatial Exploration. *arXiv preprint arXiv:2211.15788* (2022).

[20] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, Vol. 2019. NIH Public Access, 6558.

[21] Lily Xu, Shahrzad Gholami, Sara McCarthy, Bistra Dilkina, Andrew Plumptre, Milind Tambe, Rohit Singh, Mustapha Nsubuga, Joshua Mabonga, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Tom Okello, and Eric Enyel. 2020. Stay Ahead of Poachers: Illegal Wildlife Poaching Prediction and Patrol Planning Under Uncertainty with Field Test Evaluations (Short Version). In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. 1898–1901. https://doi.org/10.1109/ICDE48307.2020.00198