The Cartographer as Auditor: From Invisible Uncertainty to Reliable Spatial Knowledge

Peng Luo [1,2]* and Xiayin Lou [3]

[1] Senseable City Lab, Massachusetts Institute of Technology, Cambridge, USA
[2] School of Earth, Environment, and Sustainability, University of Iowa, Iowa City, USA
[3] Chair of Cartography and Visual Analytics, Technical University of Munich, Munich, Germany
* pengluo@mit.edu

Keywords: uncertainty visualization, geographic bias, human-in-the-loop, reliable spatial analytics, visual analytics

Introduction: Invisible Uncertainty in Spatial Analysis

Spatial analysis follows a pipeline of selecting data, building models, and interpreting results to generate spatial knowledge (Goodchild, 2007). While GeoAI has revolutionized this pipeline with remarkable predictive performance (Li & Hsu, 2022; Janowicz et al., 2020), the black-box nature of these models, combined with an incomplete understanding of spatial data, results in endogenous biases at each stage that remain invisible to analysts (Chen et al., 2026). For example, at the data level, spatial heterogeneity can produce Simpson's Paradox (Simpson, 1951); at the modeling level, methods like GWR (Fotheringham et al., 2002) assume local linearity that may not hold; and at the interpretation level, the Modifiable Areal Unit Problem (MAUP; Openshaw, 1984) demonstrates how grouping choices can alter conclusions derived from identical data.

These biases, rooted in standard analytical practices, introduce inherent uncertainty throughout the spatial analysis pipeline, potentially rendering results unreliable and posing ethical risks when such analysis informs public policy. To ensure that uncertainty and bias are systematically addressed and mitigated, we propose a framework (Figure 1) for reliable spatial knowledge discovery that keeps the human analyst meaningfully in the loop. This framework integrates two enabling technologies—uncertainty quantification (UQ) and visual analytics (VA)—to empower the analyst as a critical auditor who leverages uncertainty-aware visualizations to rigorously select data, modify models, and refine interpretations.

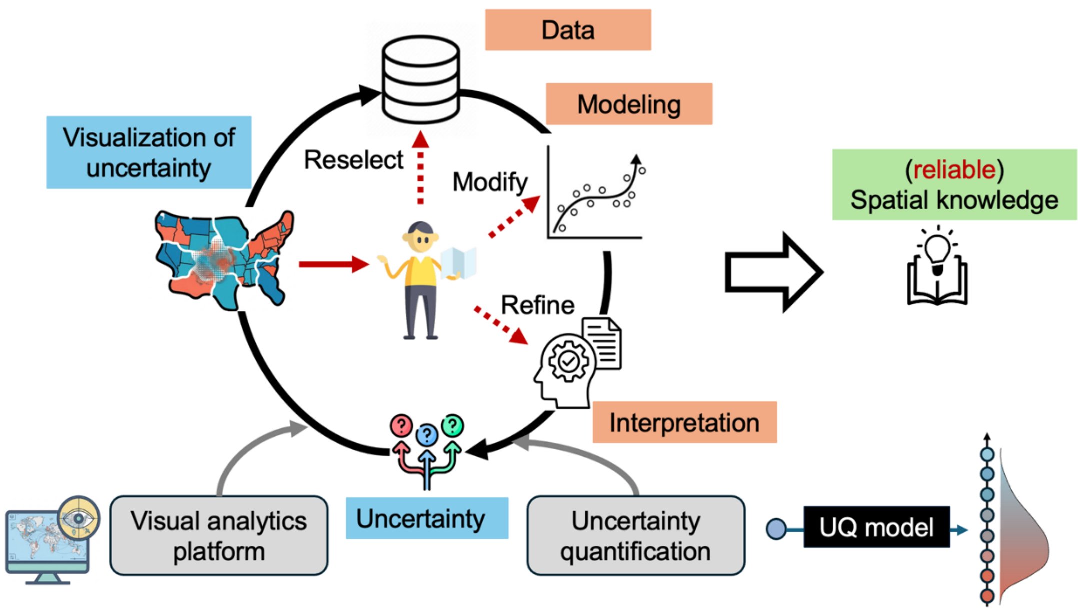

Figure 1. A human-in-the-loop framework for reliable spatial knowledge discovery. The analyst occupies the center of the spatial analysis loop (Data, Modeling, Interpretation), supported by uncertainty quantification models and a visual analytics platform.

Uncertainty Quantification: Making Confidence Measurable

The first enabling technology quantifies the reliability of GeoAI outputs. Various UQ methods are available for spatial analysis; for instance, conformal prediction (Vovk et al., 2005; Angelopoulos & Bates, 2023) provides distribution-free prediction intervals with finite-sample guarantees and has recently been integrated into geospatial studies. GeoConformal Prediction (Lou et al., 2025a) adapts this to spatial data through geographic weighting grounded in Tobler's First Law (Tobler, 1970), while GeoXCP (Lou et al., 2025b) extends uncertainty quantification to Explainable AI (XAI) explanations themselves, producing calibrated confidence intervals for SHAP-based feature attributions at each location. Without such measures, analysts cannot distinguish reliable explanations from unreliable ones (Rudin, 2019). Our vision (Lou & Luo, 2025c) is that every spatial prediction and explanation should carry its uncertainty, transforming point estimates into interval estimates.

Visual Analytics: Making Uncertainty Visible and Actionable

Uncertainty quantification is a necessary but incomplete step. It must be coupled with visual analytics to translate abstract uncertainty into actionable insights, enabling analysts to identify biases and optimize the end-to-end spatial analysis pipeline. Such a human-centric approach is critical for deriving robust and reliable spatial knowledge. Uncertainty visualization is a recognized cross-domain challenge. Pang et al. (1997) established foundational taxonomies, Brodlie et al. (2012) distinguished visualization of uncertainty from uncertainty of visualization, and Padilla et al. (2018) showed that visual encoding of uncertainty significantly affects decision quality. In the geospatial domain, traditional VA focuses on spatial distributions of variables (Andrienko et al., 2010; MacEachren & Kraak, 2001), but responsible spatial analysis requires visualizing the analytical process itself—model parameters, XAI outputs, and uncertainty measures

(Keim et al., 2008). The GeoVisX framework (Luo et al., 2025e) demonstrates this process-oriented approach by spatially mapping SHAP values and prediction errors, while Chen et al. (2026) introduce tools for detecting data heterogeneity and auditing modeling assumptions. This shift from result-oriented to process-oriented visualization transforms the analyst from a passive consumer into an active auditor (Sacha et al., 2014).

The Human as Auditor: Closing the Loop

By leveraging UQ for quantification and VA for visibility, the human analyst can effectively audit the spatial analysis pipeline, identifying hidden biases and refining the workflow (Lou et al., 2025d; Endert et al., 2017). As Figure 1 illustrates, the analyst occupies the center of the loop, interacting through three modes: re-selecting data to address heterogeneity, modifying models when visualized parameters reveal assumption violations, and refining interpretations when uncertainty overlays expose unreliable conclusions. This enables tasks no automated system can perform: distinguishing genuine patterns from modeling artifacts, evaluating whether high-importance features in uncertain regions should inform policy, and recognizing when aggregates mask vulnerable subgroups (Crawford & Finn, 2015).

Case Study: The Spatial Lens Platform

Our open-source Spatial Lens platform (Figure 2; https://vezarachan.github.io/GeoXCPWeb/) operationalizes this framework through an interactive Radial Focus Explorer. For each dataset, an XGBoost regression model is trained on a set of predictor variables to estimate a target variable; SHAP values are then computed to quantify each feature's local importance, and GeoXCP produces calibrated confidence intervals for those importance scores. Users click any map location to inspect results within an adjustable radius: left arcs encode feature importance magnitude, right arcs encode the associated geo-uncertainty. Figure 2 shows three scenarios: (a) census-tract median income in NYC (2 km), (b) housing prices in Seattle (1 km), and (c) county-level Democratic vote share across the US (80 km). Dataset details are available on the platform website. These cross-domain, cross-scale examples demonstrate how uncertainty-aware visual analytics enable real-time auditing of spatial analysis.

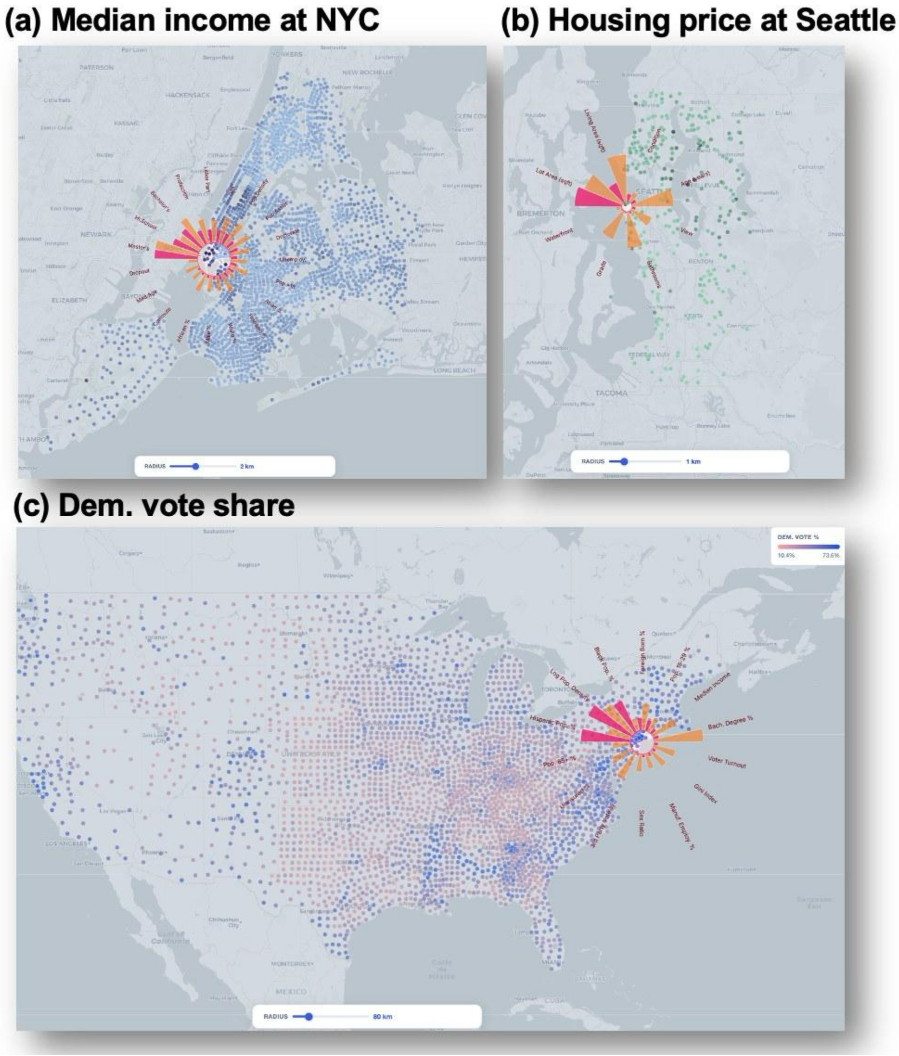

Figure 2. The Spatial Lens platform's Radial Focus Explorer. (a) Median income at NYC, (b) Housing price at Seattle, (c) Dem. vote share. Left arc = mean |SHAP|; right arc = GeoXCP prediction interval.

Conclusion

In the GeoAI era, the cartographer evolves into a critical auditor who, by synthesizing uncertainty quantification with process-oriented visual analytics, can diagnose where analytical processes introduce bias—positioning the human expert as the indispensable guarantor of reliable knowledge discovery rather than a bottleneck. Realizing this vision requires a research agenda that scales uncertainty visualization for high-dimensional data (Mai et al., 2023), formalizes visual auditing through graphical inference (Wickham et al., 2010), and embeds ethical auditing into GIScience pedagogy (Bearman et al., 2016). Ultimately, interdisciplinary co-design with domain experts is essential to ensure these tools effectively support the ethical complexities of spatial analysis in public policy (Singleton et al., 2021).

References

Andrienko, G., Andrienko, N., Demsar, U., Dransch, D., Dykes, J., Fabrikant, S. I., ... & Tominski, C. (2010). Space, time and visual analytics. International Journal of Geographical Information Science, 24(10), 1577–1600.

Angelopoulos, A. N., & Bates, S. (2023). Conformal prediction: A gentle introduction. Foundations and Trends in Machine Learning, 16(4), 494–591.

Bearman, N., Jones, N., André, I., Cachinho, H. A., & DeMers, M. (2016). The future role of GIS education in creating critical spatial thinkers. Journal of Geography in Higher Education, 40(3), 394–408.

Brodlie, K., Allendes Osorio, R., & Lopes, A. (2012). A review of uncertainty in data visualization. In J. Dill et al. (Eds.), Expanding the Frontiers of Visual Analytics and Visualization (pp. 81–109). Springer.

Chen, C., Luo, P., Zhao, B., Wang, S., Feng, Y., & Meng, L. (2026). Ethical spatial analysis: Addressing endogenous bias through visual analytics. Journal of Geovisualization and Spatial Analysis, 10(4). doi: 10.1007/s41651-025-00243-2

Crawford, K., & Finn, M. (2015). The limits of crisis data: Analytical and ethical challenges of using social and mobile data to understand disasters. GeoJournal, 80(4), 491–502.

Endert, A., Ribarsky, W., Turkay, C., Wong, B. W., Nabney, I., Blanco, I. D., & Rossi, F. (2017). The state of the art in integrating machine learning into visual analytics. Computer Graphics Forum, 36(8), 458–486.

Fotheringham, A. S., Brunsdon, C., & Charlton, M. (2002). Geographically Weighted Regression: The Analysis of Spatially Varying Relationships. Wiley.

Goodchild, M. F. (2007). Citizens as sensors: The world of volunteered geography. GeoJournal, 69(4), 211–221.

Janowicz, K., Gao, S., McKenzie, G., Hu, Y., & Bhaduri, B. (2020). GeoAI: Spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. International Journal of Geographical Information Science, 34(4), 625–636.

Keim, D. A., Andrienko, G., Fekete, J. D., Görg, C., Kohlhammer, J., & Melançon, G. (2008). Visual analytics: Definition, process, and challenges. In A. Kerren et al. (Eds.), Information Visualization (pp. 154–175). Springer.

Li, W., & Hsu, C. Y. (2022). GeoAI for large-scale image analysis and machine vision: Recent progress of artificial intelligence in geography. ISPRS International Journal of Geo-Information, 11(7), 385.

Lou, X., Luo, P., & Meng, L. (2025a). GeoConformal Prediction: A model-agnostic framework for measuring the uncertainty of spatial prediction. Annals of the American Association of Geographers. doi: 10.1080/24694452.2025.2457900

Lou, X., Luo, P., Li, Z., Gao, S., & Meng, L. (2025b). GeoXCP: Uncertainty quantification of spatial explanations in explainable AI. International Journal of Geographical Information Science, 39. doi: 10.1080/13658816.2025.2574900

Lou, X., & Luo, P. (2025c). Towards the uncertainty-aware geospatial artificial intelligence. In Proceedings of the 8th ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery (GeoAI ’25). doi: 10.1145/3764912.3770821

Lou, X., Luo, P., & Meng, L. (2025d). Uncertainty-aware geospatial visual analytics. Abstracts of the International Cartographic Association, 10(182). doi: 10.5194/ica-abs-10-182-2025

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems, 30 (pp. 4765–4774).

Luo, P., Chen, C., Gao, S., Zhang, X., Majok Chol, D., Yang, Z., & Meng, L. (2025e). Understanding of the predictability and uncertainty in population distributions empowered by visual analytics. International Journal of Geographical Information Science, 39(3), 675–705. doi: 10.1080/13658816.2024.2427870

MacEachren, A. M., & Kraak, M. J. (2001). Research challenges in geovisualization. Cartography and Geographic Information Science, 28(1), 3–12.

Mai, G., Huang, W., Sun, J., Song, S., Mishra, D., Liu, N., ... & Lao, N. (2023). On the opportunities and challenges of foundation models for GeoAI. ACM Transactions on Spatial Algorithms and Systems, 10(2), 1–46.

Openshaw, S. (1984). The Modifiable Areal Unit Problem. Concepts and Techniques in Modern Geography, 38. Geo Books.

Padilla, L. M., Creem-Regehr, S. H., Hegarty, M., & Stefanucci, J. K. (2018). Decision making with visualizations: A cognitive framework across disciplines. Cognitive Research: Principles and Implications, 3(1), 29.

Pang, A. T., Wittenbrink, C. M., & Lodha, S. K. (1997). Approaches to uncertainty visualization. The Visual Computer, 13(8), 370–390.

Robinson, A. C., Demsar, U., Moore, A. B., Buckley, A., Jiang, B., Field, K., ... & Sluter, C. R. (2017). Geospatial big data and cartography: Research challenges and opportunities for making maps that matter. International Journal of Cartography, 3(sup1), 32–60.

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1(5), 206–215.

Sacha, D., Stoffel, A., Stoffel, F., Kwon, B. C., Ellis, G., & Keim, D. A. (2014). Knowledge generation model for visual analytics. IEEE Transactions on Visualization and Computer Graphics, 20(12), 1604–1613.

Simpson, E. H. (1951). The interpretation of interaction in contingency tables. Journal of the Royal Statistical Society: Series B, 13(2), 238–241.

Singleton, A. D., Spielman, S. E., & Folch, D. C. (2018). Urban Analytics. SAGE Publications.

Tobler, W. R. (1970). A computer movie simulating urban growth in the Detroit region. Economic Geography, 46(sup1), 234–240.

Vovk, V., Gammerman, A., & Shafer, G. (2005). Algorithmic Learning in a Random World. Springer.

Weiskopf, D. (2022). Uncertainty visualization: Concepts, methods, and applications in biological data visualization. Frontiers in Bioinformatics, 2, 793819. doi: 10.3389/fbinf.2022.793819

Wickham, H., Cook, D., Hofmann, H., & Buja, A. (2010). Graphical inference for infovis. IEEE Transactions on Visualization and Computer Graphics, 16(6), 973–979.