

From Street Views to Urban Science: Discovering Road Safety Factors with Multimodal Large Language Models

Anonymous authors

Paper under double-blind review

Abstract

Urban and transportation research has long sought to uncover statistically meaningful relationships between key variables and societal outcomes such as road safety, to generate actionable insights that guide the planning, development, and renewal of urban and transportation systems. However, traditional workflows face several key challenges: (1) reliance on human experts to propose hypotheses, which is time-consuming and prone to confirmation bias; (2) limited interpretability, particularly in deep learning approaches; and (3) underutilization of unstructured data that can encode critical urban context. Given these limitations, we propose a Multimodal Large Language Model (MLLM)-based approach for interpretable hypothesis inference, enabling the automated generation, evaluation, and refinement of hypotheses concerning urban context and road safety outcomes. Our method leverages MLLMs to craft safety-relevant questions for street view images (SVIs), extract interpretable embeddings from their responses, and apply them in regression-based statistical models. URBANX supports iterative hypothesis testing and refinement, guided by statistical evidence such as coefficient significance, thereby enabling rigorous scientific discovery of previously overlooked correlations between urban design and safety. Experimental evaluations on Manhattan street segments demonstrate that our approach outperforms pretrained deep learning models while offering full interpretability. Beyond road safety, URBANX can serve as a general-purpose framework for urban scientific discovery, extracting structured insights from unstructured urban data across diverse socioeconomic and environmental outcomes. This approach enhances model trustworthiness for policy applications and establishes a scalable, statistically grounded pathway for interpretable knowledge discovery in urban and transportation studies.

1 Introduction

Understanding how cities' physical structures shape societal outcomes is central to urban science. Across transportation, planning, and policy, researchers seek links between urban form and social indicators such as traffic safety (Yu et al., 2024), walkability (Ewing & Handy, 2009), equity (Guzman & Bocarejo, 2017), and environmental health (Majchrowska et al., 2022). A key challenge is discovering generalizable, interpretable factors that explain urban phenomena (Batty, 2024). However, urban form is complex and heterogeneous, with relevant information often stored in unstructured formats such as street-level imagery and visual cues (Biljecki & Ito, 2021), which are difficult to analyze using traditional methods.

Despite progress in urban analytics, identifying new interpretable factors remains challenging. Existing methods often rely on expert-defined features, black-box models, or




Figure 1: Real vs. synthetic mobility patterns.

handcrafted metrics (Xia et al., 2025), each with limitations. First, hypothesis generation is manual and prone to cognitive bias (Gettys & Fisher, 1979). Second, deep models usually lack interpretability. Third, unstructured data like SVIs are underused due to difficulties in extracting meaningful structure (Tang et al., 2025). These limitations hinder scalable and transparent urban analysis. Foundation models such as Large Language Models (LLMs) (Naveed et al., 2023) have transformed data-driven reasoning. Trained on large text corpora, LLMs perform flexible and context-aware inference (Wei et al., 2022). Recent extensions to visual inputs have produced Multimodal LLMs (MLLMs) (Wu et al., 2023), which jointly process images and text for tasks like scene interpretation and visual reasoning. MLLMs align visual content with language, allowing them to generate human-interpretable variables from raw imagery.

Given these, we present URBANX, a framework for hypothesis-driven urban discovery powered by MLLMs. As illustrated in Figure 1, URBANX treats machine learning as a collaborator in scientific inquiry. It iteratively generates hypotheses, derives variables from multimodal data, and evaluates their statistical relevance. Weak hypotheses are discarded, and new ones are proposed, gradually refining a set of interpretable, empirically supported factors. We apply URBANX to urban road safety, where interpretability is crucial. In a Manhattan case study, the framework uncovers novel visual variables from SVIs that correlate with crash rates. Our approach surpasses deep learning baselines such as ResNet and Vision Transformer, while maintaining transparency. Our contributions are:

- We frame scientific discovery in urban contexts as inference over a hypothesis space, enabling machines to generate, test, and refine hypotheses using available data.
- We propose using MLLMs as semantic engines that transform unstructured inputs, such as SVIs, into interpretable variables based on natural-language hypotheses.
- We design an interpretable, nonparametric, iterative framework that approximates the posterior over hypotheses, enabling scalable and statistically grounded discovery of novel urban factors.
- We demonstrate the effectiveness of our framework on road safety in Manhattan, where it discovers visual predictors of crash rates that outperform vision baselines while offering interpretable results. The framework generalizes to other domains in urban science.

2 Related Work

Understanding how urban form influences outcomes like public health, equity, and road safety is central to urban science and transportation research (Hall, 2012). Traditional approaches rely on statistical models that relate expert-defined variables to outcomes (Santamouris, 2013). In road safety, for example, street design, traffic calming, and pedestrian infrastructure have been linked to crash rates (Ewing & Dumbaugh, 2009). These methods, however, face several limitations. Hypothesis generation is often manual and based on intuition, making it slow, biased, and narrow in scope (Xia et al., 2025). This can restrict the discovery of less obvious relationships. While deep learning models can enhance prediction, they usually lack interpretability (Goodfellow et al., 2016), making it difficult to identify causal drivers or support policy decisions. Their opacity can reduce trust, particularly in critical applications (Benara et al., 2024). Another challenge lies in the underused unstructured data. SVIs contain rich visual details about the urban environment, such as infrastructure condition and safety cues (Biljecki & Ito, 2021). Yet, their use in quantitative research is limited by issues of image consistency, variability, and the difficulty of extracting structured variables (Tang et al., 2024). Current methods often use general computer vision models that require extensive tuning and still may miss subtle, context-specific signals. Recent advances in AI-driven discovery offer promising directions. Some work uses LLMs for causal inference in urban settings (Xia et al., 2025), or combines LLMs with knowledge graphs for hypothesis generation in other domains (Lopez et al., 2025). These efforts show AI’s potential to support scientific reasoning, but a transparent framework that generates and tests hypotheses directly from SVIs remains largely unexplored. Our work addresses this gap by using MLLMs to build interpretable hypotheses linking urban visuals to road safety outcomes.

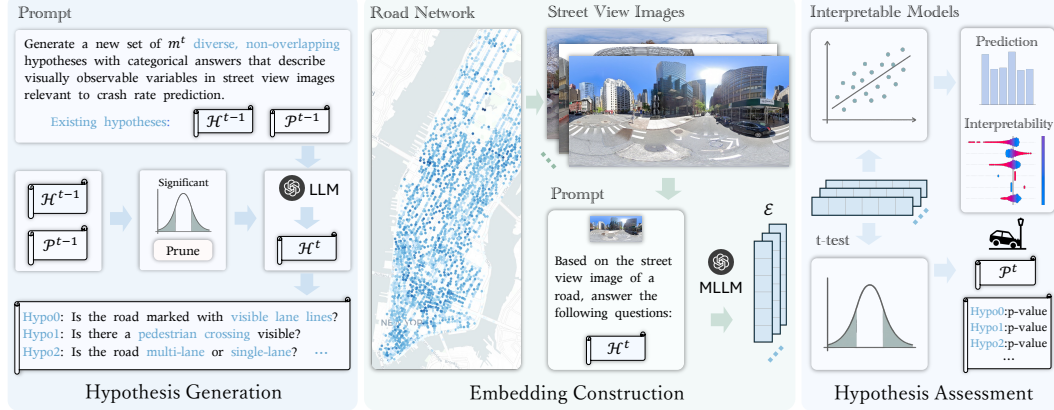


Figure 2: The URBANX framework consists of three iterative modules: (1) Hypothesis Generation using LLMs, (2) Embedding Construction via MLLM-based VQA on SVIs, and (3) Hypothesis Assessment using interpretable regression analysis.

3 Methodology

Overview Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$ denote a dataset of n SVIs x_i and their associated road-level crash rates $y_i \in \mathbb{R}$. We define a hypothesis space \mathcal{H} comprising all natural-language queries that describe visually observable variables potentially related to road safety. Our objective is to uncover an optimal subset of hypotheses $\mathcal{H}^* = \{h_1, h_2, \dots, h_k\} \subset \mathcal{H}$ that captures meaningful visual semantics from each SVI and enables interpretable, accurate prediction of y_i . We formalize this as a posterior mode estimation problem over the hypothesis space: $\mathcal{H}^* = \arg \max_{\mathcal{H}' \subseteq \mathcal{H}} P(\mathcal{H}' | \mathcal{D}) \propto P(\mathcal{D} | \mathcal{H}') \cdot P(\mathcal{H}')$, where \mathcal{H}' is a candidate hypothesis subset. The likelihood $P(\mathcal{D} | \mathcal{H}')$ captures how well the hypothesis-derived variables explain variation in crash rates, typically assessed via a regression model. The prior $P(\mathcal{H}')$ encodes structural preferences over hypothesis subsets and is implicitly governed by the generative behavior of the MLLM. Each hypothesis $h_j \in \mathcal{H}^*$ corresponds to a semantically meaningful question with a categorical answer that could be inferred from an SVI using an MLLM. Applying these k hypotheses to each image x_i yields a k -dimensional interpretable embedding $\phi(x_i, \mathcal{H}^*) \in \mathbb{R}^k$, where each component reflects the MLLM’s answer to the corresponding hypothesis. We denote the complete embedding matrix as $\mathcal{E} \in \mathbb{R}^{n \times k}$, where $e_i = \phi(x_i, \mathcal{H}^*)$ is the embedding vector for the i -th image. Since exact Bayesian inference over all subsets of \mathcal{H} is intractable due to the large search space and unknown likelihoods, we use an approximate approach and cast the problem as nonparametric structure learning. Starting with an initial set \mathcal{H}^0 sampled from an LLM, we iteratively evaluate and refine it using a linear regression model. For each hypothesis h_j , we test the significance of its regression coefficient via a two-sided t -test under the null hypothesis that the coefficient is zero. This gives a p -value vector $\mathcal{P} = \{p_1, p_2, \dots, p_k\}$, where each p_j represents the probability of observing the result under the null. Hypotheses with $p_j > \alpha$ (typically 0.05) are treated as statistically insignificant and removed. New hypotheses are generated to replace them, forming an iterative refinement loop. This procedure approximates Bayesian inference over \mathcal{H} using statistical evidence and LLM priors, in a nonparametric and data-driven way. An overview of the URBANX framework is shown in Figure 2.

Hypothesis Generation At each iteration t , the framework refines the hypothesis set \mathcal{H}^{t-1} using statistical evidence derived from the previous assessment. For each hypothesis $h_j \in \mathcal{H}^{t-1}$, we compute a p -value p_j using a two-sided t -test on the coefficient estimated by a regression model, where the input variable is derived from the MLLM-inferred categorical responses to h_j across all SVIs. The detailed procedure for constructing hypothesis-driven embeddings is described in a later subsection. Hypotheses with $p_j > \alpha$ (typically $\alpha = 0.05$) are considered statistically insignificant. While the prompt for the LLM includes the full set of previous hypotheses \mathcal{H}^{t-1} and their p -values \mathcal{P}^{t-1} , only m^t new hypotheses are generated, where m^t equals the number of pruned hypotheses. This maintains a

fixed hypothesis set size while ensuring that each iteration incorporates empirical feedback into the generative process. Formally, the hypothesis generation step is given by: $\mathcal{H}^t \sim \text{LLM}(\text{Prompt}_{\text{HypoGen}}(\mathcal{H}^{t-1}, \mathcal{P}^{t-1}, m^t))$, where m^t is the number of new hypotheses to generate. The prompt is constructed to elicit m^t diverse, categorical, and visually inferable questions that are relevant to crash prediction. By conditioning on statistically grounded examples, the LLM acts as a posterior-informed generator, implicitly sampling from a distribution biased toward hypotheses that are both semantically coherent and empirically promising. This design allows the system to balance exploration of new concepts with exploitation of previously validated structure, enabling effective refinement of the hypothesis space over time. An illustration of this process is shown in the left panel of Figure 2.

Embedding Construction To leverage the generated hypotheses $\mathcal{H}^t = \{h_1^t, h_2^t, \dots, h_k^t\}$ for downstream modeling, we must transform their semantic content into structured, machine-interpretable representations. Traditional deep models rely on latent high-dimensional features extracted from images, which lack transparency and hinder hypothesis-driven analysis. In contrast, our goal is to construct a hypothesis-guided embedding that is transparently aligned with the semantics of each generated question. For each image x_i , we use an MLLM to answer all questions in \mathcal{H}^t based on the visual content of the image. These categorical answers are then encoded into a k -dimensional vector $e_i^t \in \mathbb{R}^k$, where each element corresponds to the response to hypothesis h_j^t . Formally, we define: $e_i^t \sim \text{MLLM}(x_i, \text{Prompt}_{\text{Embed}}(\mathcal{H}^t))$, where $\text{Prompt}_{\text{Embed}}(\mathcal{H}^t)$ denotes the prompt that queries the MLLM to answer all hypotheses in \mathcal{H}^t based on the visual content of x_i . This embedding ensures full semantic traceability and supports interpretable downstream modeling.

This procedure yields a hypothesis-aligned, semantically interpretable embedding for each image, where each dimension has a well-defined linguistic meaning. It enables transparent variable construction while supporting statistical assessment and iterative refinement in subsequent stages. The embedding process is illustrated in the center panel of Figure 2.

Hypothesis Assessment After constructing semantically aligned embeddings \mathcal{E}^t for all SVIs based on the current hypothesis set \mathcal{H}^t , the next step is to assess which hypotheses meaningfully explain variation in crash rates. Rather than focusing solely on predictive accuracy, we adopt interpretable models that provide transparent, decomposable attribution of outcomes to individual hypotheses. This is particularly important in societal domains such as road safety, where policy decisions and public accountability require not only reliable predictions but also actionable explanations. Each input dimension in the model corresponds to a specific hypothesis, enabling us to quantify its effect and assess statistical significance. In our implementation, we use linear regression as the default method due to its analytical tractability and well-established inference procedures.

For each embedding dimension, we perform a two-sided t -test on its regression coefficient to test the null hypothesis that the coefficient is zero. This yields a p -value for each hypothesis h_j^t , denoted p_j^t , indicating the likelihood that the observed effect is due to chance. The resulting vector $\mathcal{P}^t = \{p_1^t, p_2^t, \dots, p_k^t\}$ serves as the statistical evidence that guides the next iteration: hypotheses with $p_j^t > \alpha$ (typically $\alpha = 0.05$) are deemed statistically insignificant and are replaced by new hypotheses in the next round. This assessment mechanism plays a dual role: it enables interpretability by quantifying the contribution of each hypothesis, and it drives hypothesis refinement by filtering out those that lack explanatory value. The right panel of Figure 2 illustrates this evaluation process.

Iterative Posterior Approximation The overall framework is executed through an iterative loop that approximates posterior inference over the hypothesis space by alternating between generation, embedding, and statistical assessment. This process reflects a structure-learning approach where hypothesis subsets are progressively refined based on empirical evidence. Unlike standard optimization methods such as gradient descent or expectation maximization, where the objective function is guaranteed to monotonically improve, our setting

Algorithm 1 Iterative Posterior Approximation

Require: Dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$; number of total hypotheses k ; number of iterations T ; interpretable model \mathcal{M}

Ensure: Final hypothesis set \mathcal{H}^T and embedding matrix \mathcal{E}^T

- 1: Initialize $\mathcal{H}^0 \sim \text{LLM}(\text{Prompt}_{\text{HypoGen}}(k))$
- 2: **for** $t = 1, 2, \dots, T$ **do**
- 3: $\mathcal{H}^t \sim \text{LLM}(\text{Prompt}_{\text{HypoGen}}(\mathcal{H}^{t-1}, \mathcal{P}^{t-1}, m^t))$ ▷ Hypothesis Generation
- 4: $\mathcal{E}^t = \{e_i^t = \text{MLLM}(x_i, \text{Prompt}_{\text{Embed}}(\mathcal{H}^t))\}_{i=1}^n$ ▷ Embedding Construction
- 5: $\{\hat{y}_i\}_{i=1}^n, \mathcal{P}^t \leftarrow \mathcal{M}(\mathcal{E}^t)$ ▷ Hypothesis Assessment
- 6: **end for**

involves sampling from a nonparametric, LLM-driven space that lacks such guarantees. To mitigate the risk of degeneracy or performance collapse, we adopt a conservative update rule: new hypotheses \mathcal{H}^t are only retained if they yield improved predictive performance on the validation set compared to the previous iteration.

Algorithm 1 outlines the overall iterative procedure. In each iteration, insignificant hypotheses from the previous set \mathcal{H}^{t-1} are filtered based on their p -values \mathcal{P}^{t-1} . The remaining hypotheses serve as context for LLM-based generation of new candidates. The resulting hypothesis set \mathcal{H}^t is then used to construct interpretable embeddings \mathcal{E}^t via MLLM-based reasoning, which are subsequently used to train an interpretable model and evaluate statistical significance. This iterative process continues for a predefined number of iterations.

4 Experiments

Settings We study road segments in Manhattan, New York City, using crash records, traffic volume data, and street-view imagery from 2013 to 2019. For each segment, we compute the crash rate as in prior work (Yu et al., 2024): $CR_i = \frac{\text{No_crash}_i}{AADT_i \times L_i \times \frac{365}{1,000,000}}$, where No_crash_i is the average yearly number of crashes, $AADT_i$ is the average annual daily traffic, and L_i is the segment length in kilometers. Crash data comes from NYC Open Data, and $AADT$ is from the New York State Department of Transportation. SVIs were sampled every 15 meters along road centerlines using ArcGIS, and images were collected via the Google Street View API. After filtering, we used 16,000 panoramic SVIs for training, 2,000 for validation, and 2,000 for testing. Unless stated otherwise, we use GPT-4o as the LLM for hypothesis generation and InternVL2.5-78B (Chen et al., 2024) as the MLLM for hypothesis answering during embedding generation. We run MLLMs using LMDeploy, an optimized inference engine for efficient deployment. For evaluation and comparison, we also collect 58 conventional built environment variables from public sources. These fall into five categories: (1) road attributes, (2) land use, (3) point-of-interest (POI) features, (4) traffic-related facilities, and (5) visual indices from panoptic segmentation of SVIs. These features are from sources including PLUTO, CommonPlace, Geofabrik.

Predictive Performance We first evaluate the predictive performance of our interpretable embedding framework by comparing it with conventional vision-based baselines. Figure 3 reports results across three standard metrics: root mean square error (RMSE), mean absolute error (MAE), and the coefficient of determination (R^2). For baselines, we use two representative pretrained image encoders: ResNet50, a widely adopted convolutional architecture, and ViT-Base (ViT-B/16), specifically the vit_base_patch16_224 variant that seg-




Figure 3: Performance comparison between ResNet, ViT, and our interpretable embedding-based models using linear regression (LR) and LightGBM (LGBM).

ments each image into 16×16 patches and processes them with transformer blocks. These models are fine-tuned to predict crash rates directly from raw SVIs. We compare these against two variants of our framework that rely on interpretable embeddings constructed from MLLM responses: one using linear regression (LR) and another using LightGBM (LGBM) as the downstream predictor. Across all metrics, our method consistently outperforms the deep learning baselines while maintaining transparency and semantic interpretability. The LightGBM variant achieves the strongest overall results. These results demonstrate that the embeddings retain sufficient information to make accurate predictions while also enabling interpretability.

Discovered Factors A central goal of URBANX is not just to predict societal outcomes, but to enable interpretable, data-driven discovery of unstructured or previously overlooked urban factors. To assess whether our framework successfully identifies meaningful visual variables, we examine the learned regression model using SHAP (SHapley Additive exPlanations) analysis, which quantifies each variable's marginal contribution to the prediction.

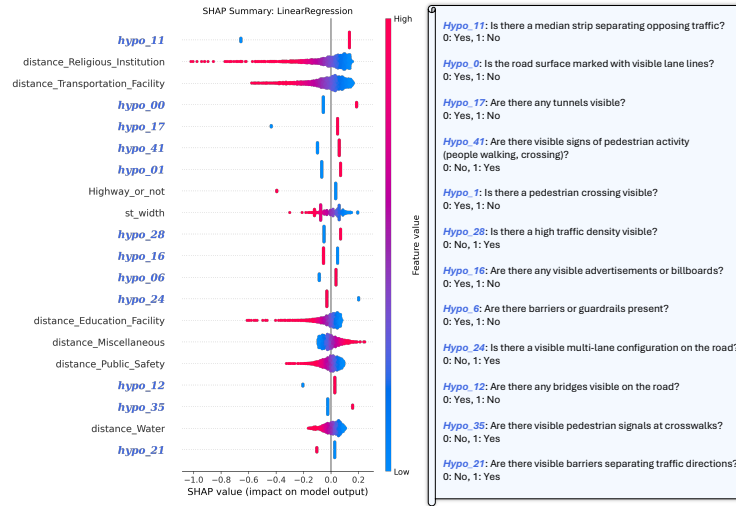


Figure 4: SHAP plot of the regression model with both traditional built environment variables and discovered hypotheses.

Figure 4 presents a top-20 ranked summary of both traditional (existing) built environment variables and the automatically discovered hypotheses. Remarkably, a majority of the top-ranked variables by explanatory power are generated by our LLM-based hypothesis pipeline. This highlights URBANX’s capacity to uncover impactful, interpretable factors that are not present in standard urban datasets, supporting its role as a scientific discovery tool rather than a black-box predictor. Many of the discovered hypotheses align with well-established urban safety principles, validating the model’s ability to recover known but unstated domain knowledge. For example, Hypo_11 (“Is there a median strip separating opposing traffic?”) and Hypo_0 (“Is the road surface marked with visible lane lines?”) are both highly ranked and show negative SHAP contributions when absent, suggesting their presence is associated with lower crash risk. These align with conventional traffic engineering wisdom on lane separation and visual guidance.

At the same time, URBANX also surfaces more nuanced or less commonly considered factors. Several high-ranking hypotheses relate to pedestrian visibility and activity, such as Hypo_1 (pedestrian crossing), Hypo_41 (pedestrian presence), and Hypo_35 (pedestrian signals). These factors may have complex and context-sensitive relationships with safety outcomes, underscoring the value of semantically grounded, hypothesis-level variables. In addition, URBANX identifies less conventional features that might escape manual enumeration. For instance, Hypo_16 (“Are there any advertisements or billboards?”) and Hypo_6 (“Are there barriers or guardrails present?”) point to visual distractions and physical protection measures that may subtly influence crash risk. These hypotheses extend the scope of interpretable modeling into environmental and perceptual dimensions that are often hard to encode using conventional GIS-based variables. Compared to traditional indicators such as street width or proximity to facilities, our hypotheses are more granular, semantically aligned, and directly grounded in what is observable in urban space. This illustrates the unique advantage of URBANX in supporting structured discovery over unstructured inputs. Taken together, these results show that URBANX is capable not only of achieving competitive predictive accuracy but also of surfacing novel, interpretable factors that enrich the understanding of urban safety.

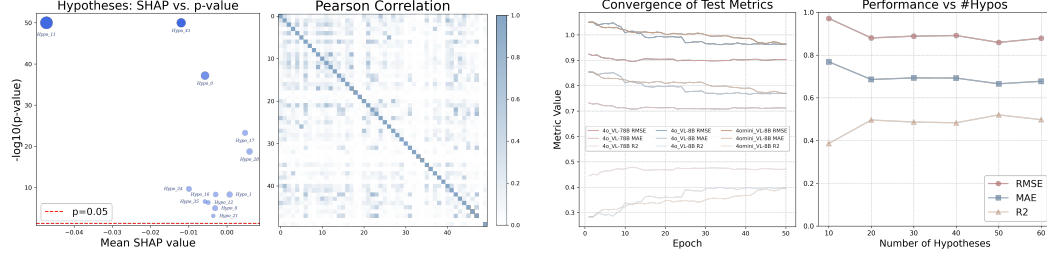


Figure 5: (Left) Variable quality analysis: SHAP vs. statistical significance and pairwise correlation. (Right) Robustness analysis: model performance across different LLM/MLLM settings and hypothesis counts.

Variable Quality, Robustness, and Practical Implications We assess the quality and robustness of URBANX’s hypotheses based on significance, independence, and sensitivity to model and hypothesis settings. As shown in the left panel of Figure 5, we visualize each hypothesis by its average SHAP value and the negative log of its p -value from linear regression. This allows us to assess both predictive contribution and statistical significance. Hypotheses such as Hypo_11, Hypo_41, and Hypo_0 score highly on both axes. These correspond to interpretable road safety features such as the presence of a median strip, pedestrian visibility, and clear lane markings. These results confirm the alignment between model-generated variables and established domain knowledge. The same panel also includes the pairwise Pearson correlation matrix, where low off-diagonal values indicate that the discovered variables are largely uncorrelated. This structural independence enhances interpretability and reduces the risk of multicollinearity.

The right panel of Figure 5 presents a robustness analysis. We vary the capacities of the LLM used for hypothesis generation and the MLLM used for answering those hypotheses. Larger MLLMs (such as InternVL2.5-78B) consistently produce better predictive accuracy and faster convergence than smaller ones (such as 8B), emphasizing the importance of visual reasoning capability. The size of the LLM has a smaller effect. GPT-4o leads to faster training convergence compared to GPT-4o-mini, but both eventually reach similar levels of performance. We also examine the impact of hypothesis count. Accuracy improves steadily up to around 50 hypotheses, after which performance plateaus or slightly declines. This reflects a tradeoff between representational richness and statistical noise.

5 Conclusion

We introduced URBANX, a framework that uses MLLMs and interpretable modeling to automate hypothesis-driven discovery from urban data. Applied to road safety in Manhattan, URBANX generates natural-language hypotheses, extracts visual embeddings, and evaluates them using transparent regression, outperforming deep learning baselines while revealing meaningful, interpretable variables. URBANX represents a new approach to urban science that unifies perception, language, and reasoning. Its generality supports applications beyond road safety, including walkability, equity, and environmental quality. Future directions include extending to dynamic data, integrating causal inference, and leveraging advances in foundation models for scalable, interpretable discovery.

References

- Michael Batty. *The computable city: histories, technologies, stories, predictions*. MIT Press, 2024.
- Vinamra Benara, Chandan Singh, John X Morris, Richard J Antonello, Ion Stoica, Alexander G Huth, and Jianfeng Gao. Crafting interpretable embeddings for language neuroscience by asking llms questions. *Advances in neural information processing systems*, 37: 124137, 2024.
- Filip Biljecki and Koichi Ito. Street view imagery in urban analytics and gis: A review. *Landscape and Urban Planning*, 215:104217, 2021.

- 322 Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu,
323 Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of
324 open-source multimodal models with model, data, and test-time scaling. *arXiv preprint*
325 *arXiv:2412.05271*, 2024.
- 326 Reid Ewing and Eric Dumbaugh. The built environment and traffic safety: a review of
327 empirical evidence. *Journal of Planning Literature*, 23(4):347–367, 2009.
- 328 Reid Ewing and Susan Handy. Measuring the unmeasurable: Urban design qualities related
329 to walkability. *Journal of Urban design*, 14(1):65–84, 2009.
- 330 Charles F Gettys and Stanley D Fisher. Hypothesis plausibility and hypothesis generation.
331 *Organizational behavior and human performance*, 24(1):93–110, 1979.
- 332 Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, vol-
333 ume 1. MIT press Cambridge, 2016.
- 334 Luis A Guzman and Juan P Bocarejo. Urban form and spatial urban equity in bogota,
335 colombia. *Transportation research procedia*, 25:4491–4506, 2017.
- 336 Randolph Hall. *Handbook of transportation science*, volume 23. Springer Science & Business
337 Media, 2012.
- 338 Vanessa Lopez, Lam Hoang, Marcos Martinez-Galindo, Raúl Fernández-Díaz, Marco Luca
339 Sbodio, Rodrigo Ordonez-Hurtado, Mykhaylo Zayats, Natasha Mulligan, and Joao
340 Bettencourt-Silva. Enhancing foundation models for scientific discovery via multimodal
341 knowledge graph representations. *Journal of Web Semantics*, 84:100845, 2025.
- 342 Sylwia Majchrowska, Agnieszka Mikołajczyk, Maria Ferlin, Zuzanna Klawikowska, Marta A
343 Plantykwow, Arkadiusz Kwasigroch, and Karol Majek. Deep learning-based waste detec-
344 tion in natural and urban environments. *Waste Management*, 138:274–284, 2022.
- 345 Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad
346 Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. A comprehensive overview of
347 large language models. *arXiv preprint arXiv:2307.06435*, 2023.
- 348 Matheos Santamouris. *Energy and climate in the urban built environment*. Routledge, 2013.
- 349 Yihong Tang, Zhaokai Wang, Ao Qu, Yihao Yan, Zhaofeng Wu, Dingyi Zhuang, Jushi Kai,
350 Kebin Hou, Xiaotong Guo, Jinhua Zhao, et al. Itinera: Integrating spatial optimization
351 with large language models for open-domain urban itinerary planning. In *Proceedings of*
352 *the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pp.
353 1413–1432, 2024.
- 354 Yihong Tang, Menglin Kong, and Lijun Sun. Large language models for data synthesis.
355 *arXiv preprint arXiv:2505.14752*, 2025.
- 356 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani
357 Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of
358 large language models. *arXiv preprint arXiv:2206.07682*, 2022.
- 359 Jiayang Wu, Wensheng Gan, Zefeng Chen, Shicheng Wan, and Philip S Yu. Multimodal large
360 language models: A survey. In *2023 IEEE International Conference on Big Data (BigData)*,
361 pp. 2247–2256. IEEE, 2023.
- 362 Yutong Xia, Ao Qu, Yunhan Zheng, Yihong Tang, Dingyi Zhuang, Yuxuan Liang, Shen-
363 hao Wang, Cathy Wu, Lijun Sun, Roger Zimmermann, and Jinhua Zhao. Reimagining
364 urban science: Scaling causal inference with large language models. *arXiv preprint*
365 *arXiv:2504.12345*, 2025.
- 366 Xujing Yu, Jun Ma, Yihong Tang, Tianren Yang, and Feifeng Jiang. Can we trust our eyes?
367 interpreting the misperception of road safety from street view images and deep learning.
368 *Accident Analysis & Prevention*, 197:107455, 2024.