



RELATIONSHIP MODELLING FOR ROAD GEOMETRIC DIGITAL TWINS USING SPATIAL ANALYSIS AND LLMs

Diana Davletshina¹, Varun Kumar Reja¹, and Ioannis Brilakis¹

¹University of Cambridge, Cambridge, UK

Abstract

Geometric Digital Twins offer a robust framework for managing road environments by integrating high-fidelity object geometries and their interrelationships. However, existing relationship modelling approaches rely on manual processes or are not tailored for road objects, creating bottlenecks in large-scale digitisation. This paper proposes a hybrid method combining spatial analysis with large language models to automate topological and functional relationship modelling among road assets. Our approach captures and enriches the spatial context with semantic insights derived from relevant textual data. The resulting holistic GDT leads to more efficient information and asset management, facilitating advanced use cases, including safety assessments and simulations.

Introduction

This paper is about modelling relationships for road geometric digital twins using spatial analysis and large language models (LLMs). A geometric digital twin (GDT) is a digital replica of a physical asset with a focus on the geometric properties of objects and their relationships (Davletshina et al., 2024). GDT forms a twin system along with the physical asset, i.e., the physical twin., with a two-way connection between them (Botín-Sanabria et al., 2022; Jiang et al., 2021). Relationship modelling includes identifying topological and functional relationships between road assets and their explicit digital representation in 2D or 3D (Davletshina and Brilakis, 2024).

Relationships between objects are critical to success in built environment projects because they reveal how individual elements interact and influence one another (Sacks et al., 2020). Along with semantically labelled objects and their geometries, these relationships form the backbone of robust project data, enabling comprehensive representation and supporting advanced analytics — such as artificial intelligence and machine learning — without narrow scope limitations. They also lay the groundwork for constructing digital twins (DTs), specifically GDTs.

GDTs have emerged as a powerful framework for various use cases in road environments, driving improvements in both information and asset management. As digital replicas, they support safety assessments, computer simulations, predictive maintenance, and resource optimization

(Marie d'Avigneau et al., 2025). More broadly, they enable higher levels of digitalization in transportation and construction, ultimately enhancing productivity and efficiency across these industries (McKinsey Global Institute, 2015).

Currently, state-of-practice software like Autodesk InfraWorks (Autodesk, 2023) often models relationships semi-automatically, requiring manual object fitting (Davletshina and Brilakis, 2024). This manual process is time-consuming and creates a critical bottleneck, particularly when digitising thousands of kilometres of roads.

In this paper, we address a key challenge in constructing GDTs by automating the generation of object relationships. Building on our previous work on scene understanding and 3D reconstruction (Davletshina et al., 2024), we complement the automatic construction of road GDTs through relationship modelling. This automation significantly reduces – and in some cases eliminates – the manual labour required, thereby ensuring that GDTs deliver greater overall benefits and a higher return on investment. Our contributions in this paper are as follows:

1. We propose a hybrid method for relationship discovery and modelling for road GDT objects, which combines spatial analysis with natural language processing, resulting in topological and functional relationships;
2. We deliver a holistic road GDT with connected objects, such as shown in Figure 1.

This paper starts with an introduction in this section, followed by a literature review in the Background section. Then, the Proposed solution section discusses our method for modelling relationships. The Research Methodology section covers assumptions about the approach and details the implementation steps. The Results and Discussion section provides the results of the experiments and discusses the outcomes, as well as lists limitations. The paper concludes with the Conclusion section.

Background

Scene Understanding and Relationship Modelling

Scene understanding is a critical aspect of computer vision that encompasses not just object detection and classi-

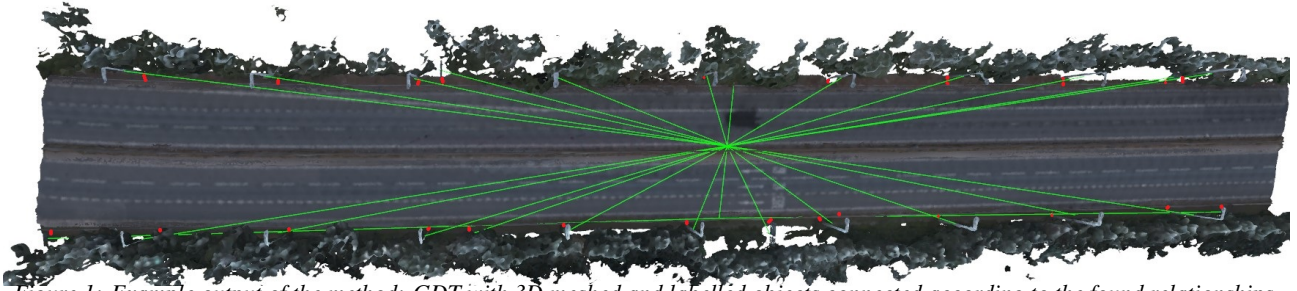


Figure 1: Example output of the method: GDT with 3D meshed and labelled objects connected according to the found relationships. Red lines – “closest to” relationships; Green lines – “nearby” or association relationships.

fication but also the identification of relationships between entities in a scene. One effective method for representing these relationships is through the construction of a scene graph, a structured representation where nodes correspond to objects and edges denote their relationships (Malawade et al., 2022; Drobný et al., 2024).

For example, Malawade et al. (2022) introduced an innovative tool designed to extract scene graphs from video clips capturing road environments. Their approach involved a two-step process: first, they extracted road entities and their relations from the video data; then, they used this information to train a graph neural network (GNN) model. The scope of their work was focused specifically on vehicles and their spatial relationships, particularly their positions within specific road lanes. This constrained domain enabled precise modelling and analysis tailored to the dynamics of traffic flow.

Similarly, Tian et al. (2020) explored the generation of road scene graphs using video imagery, employing a convolutional graph neural network. Their work expanded the domain of scene graphs to include additional road-related entities such as vehicles, pedestrians, traffic signs, and obstacles. This scope provided a more comprehensive understanding of the road environment than the previous work, allowing for the detection of complex interactions among diverse elements. Although they maintained a focus on key traffic elements critical for applications like autonomous driving and traffic management, their scope did not explicitly include other frequent road assets like lamps and guardrails, which are essential when considering wider application ranges.

Modelling Road Relationships Using IFC

Modelling road environments using the Industry Foundation Classes (IFC) format offers a robust framework for representing relationships. The IFC format inherently supports the embedding of relationships within its structure, enabling a hierarchical organization of objects, where they can be decomposed into their constituent parts or explicitly defined using the *IfcRelationship* abstraction. This abstraction allows for the modelling of a wide range of relationships, including assigning, connecting, declaring, associating, decomposing, and defining relationships (BuildingSMART, 2025).

For example, Justo et al. (2021) demonstrated the potential

of this approach by modelling a road segment along with its associated traffic sign and guardrail assets in the IFC format. Their methodology involved leveraging reference point clouds to extract properties necessary for parametric object modelling. While they successfully automated the extraction of properties from point clouds, their overall process was semi-automatic, relying on manual intervention for certain steps. This reliance on manual input poses scalability challenges when attempting to model large and complex road networks, as the effort required increases significantly with scale.

Large Language Models for Spatial Relationship Modelling

Recently, large language models were applied to geospatial tasks. Starace and Di Martino (2024) studied LLMs for various GIS report generation capabilities with case studies on earthquake distribution, traffic accidents and spatio-temporal road network coverage data. In the latter case, they explored the capability of LLMs to process complex spatial data, such as polygon shapes, and deduce spatial relationships between these intricate objects. Although their results indicated that general-purpose LLMs excel at detecting spatio-temporal patterns, identifying anomalies and generating concise, readable reports, challenges were still there. This included variability across executions and difficulty with complex spatial data like polygons.

Using LLMs, Mansourian and Oucheikh (2024) created a tool ChatGeoAI, which enables prompting with natural language queries and getting various GIS tasks involving mapping and spatial relationships done with Llama 2 LLM (Touvron et al., 2023) and QGIS software (QGIS, 2025) integration. Example tasks included queries for listing all pharmacies within a certain range to a place of interest, finding the shortest paths between A and B and filtering places of interest that satisfy certain criteria. Their system demonstrated proficiency in handling these tasks but struggled with some attributes, names and aliases.

Research Objective and Questions

In summary, the state of research offers promising solutions with the scene graph construction and spatial analysis, as well as LLMs supporting geospatial tasks. As found earlier, LLMs could help to parse and interpret textual descriptions of geographic locations or datasets, enabling

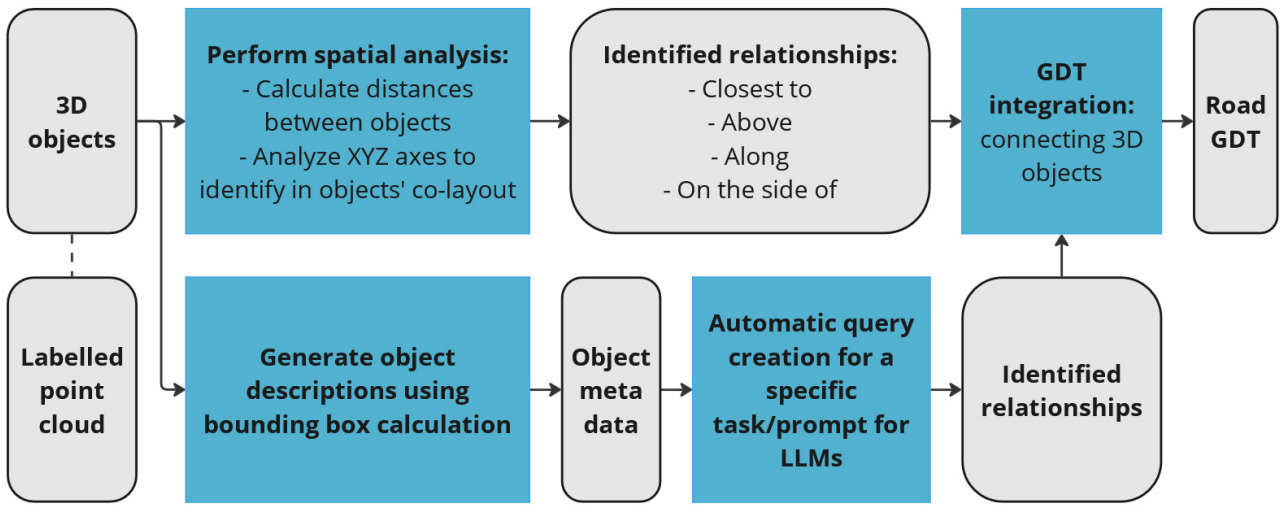


Figure 2: Relationship modelling with a hybrid approach: spatial analysis and LLMs.

automated data enrichment and categorisation of geospatial attributes. Additionally, they helped craft or translate queries for GIS software, reducing barriers for users unfamiliar with complex spatial syntax or workflows. When integrated into spatial analysis pipelines, LLMs offered quick, context-aware assistance in summarising findings, producing accessible explanations, or generating documentation for geospatial projects. However, a few limitations remained, which include the accuracy and repeatability of the tested LLM versions, as well as their limited application to road objects. Thus, the research objective in this paper is to develop a method for relationship modelling for road assets. The following research questions were identified to support the objective:

1. How to model relationships between road assets?
2. To what extent the most recent LLMs can support relationship modelling for road assets?
3. Can LLM-based relationships complement spatial analysis?

Proposed Solution

One effective method for discovering relationships between objects is to conduct spatial analysis. This approach involves examining the position, distance and arrangement of objects within a defined space to identify patterns and correlations that might otherwise remain hidden. Understanding how objects cluster or spread out or how proximity may affect their properties can bring deeper insights into the underlying dynamics of a system. Spatial analysis is particularly useful because it translates abstract locations into tangible, measurable forms — distances and patterns — which can be then analysed with statistical or machine learning tools. It allows us to visualise complex interactions more clearly, understand scene structure and build object hierarchies, revealing meaningful relationships that can be easily overlooked when relying solely

on non-spatial methods. Although spatial analysis is an accurate method for establishing spatial relationships, it remains a rigid approach. This means that any changes in the desirable output require changes in the solution backend. Therefore, LLMs that showed promise in handling geospatial tasks might bring the required flexibility.

LLMs like ChatGPT (OpenAI, 2025) can significantly support geospatial tasks by providing natural language understanding and generation capabilities that streamline data extraction, interpretation and communication. LLMs are able to process an unrestricted range of queries/prompts. Thus, this makes methods based on LLM queries flexible and much less rigid.

In this section, we propose a hybrid solution for modelling relationships, combining both spatial analysis and LLMs-based relationship modelling to support GDT development. The proposed solution will benefit from precise spatial analysis and flexibility of general-purpose LLMs. Figure 2 demonstrates the steps of the solution. The input is 3D objects, either obtained from a road GDT, e.g., constructed by Davletshina et al. (2024), or segmented point cloud, which could be also obtained using the segmentation method of Davletshina et al. (2024). With this input, two workflows are available: 1) spatial analysis for basic but precise topological relationships and 2) automatic prompt generation and LLM API (Application Programming Interface) prompting. The spatial analysis step will result in such relationships as “closest to” based on the distances between object boundaries, “above”, “on the side of” and “along” for prolonged objects like road surfaces and guardrails.

In the meantime, the LLM-based workflow will create automatic prompts based on geometric descriptions of objects and their coordinates. The automatic prompts will target the discovery of spatial relationships as well as functional relationships, as LLMs have context knowledge in the geospatial domain, as seen in the previous section. The prompt will include:

- A list of detected 3D road object instances such as road surfaces, barriers, lamps, traffic signs, etc.
- A list of geometric descriptions for each object instance, such as eight 3D bounding box point coordinates.
- A specific request for a target task, such as
 - Please, analyse a given list of road object instances and their bounding box coordinates and find the main spatial relationships. Then, draw a graph diagram with these relationships.
 - Please, analyse a given list of road object instances and their bounding box coordinates and find the main functional relationships. Then, draw a graph diagram with these relationships.

The specific task prompts could be modified, but the overall solution pipeline remains the same.

The discovered relationships are used further to link objects in GDT in the GDT integration step (Figure 2).

Research Methodology

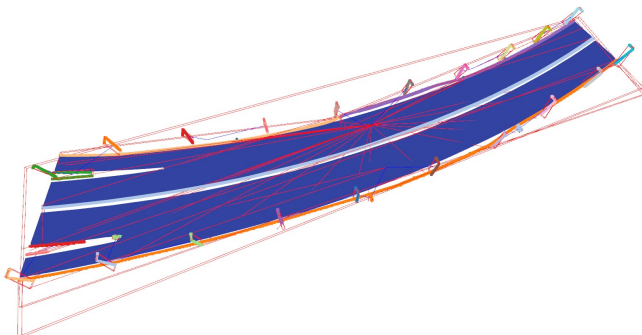


Figure 3: Minimal oriented bounding boxes around road instances.

Assumptions of this research include the existence of road GDTs in the form of 3D meshed objects or semantically labelled point clouds, both of which could be obtained using a solution by Davletshina et al. (2024). Furthermore, we assume the existence of a trained LLM such as ChatGPT (OpenAI, 2025).

We employ the CAMHighways dataset (Marie d’Avigneau et al., 2025) for our experiments on the A12 road segment, as well as the road GDT constructed by Davletshina et al. (2024). In our analyses, we utilise point cloud clusters corresponding to each 3D object when determining various geometric properties. If these point clouds are not stored with the 3D objects, 3D meshes can be used instead, or a point cloud can be sub-sampled from the 3D models to achieve the same objectives.

In order to identify spatial relationships, distances between object instances are calculated in two ways: 1) between object centres as defined by their minimal bounding box centres and 2) between the closest points in the corresponding

point clouds. A minimal oriented bounding box is calculated around each object instance, as exemplified in Figure 3. Open3D python library (Open3D, 2025) is used for this operation. For calculating the shortest distance between objects, KDTree (SciPy, 2025) is built on the object’s points to find the closest point for the more optimal search. Thus, the closest objects based on centres and based on the shortest distances are found.

Among LLMs, we chose one of the state-of-the-art networks, specifically ChatGPT o1 model (OpenAI, 2025). The following prompts were used:

1. *Hi! I have a list of road object instances and their bounding box coordinates. Could you please analyse this data and find the main **functional** relationships? Then, draw a graph diagram with these relationships. Use the listed objects as nodes and relationships - as arrow links. Please add text on nodes with the corresponding objects and instances and text on arrow links- with the corresponding relationship.*

Object barrier Instance 0 Bounding box coordinates ... Object barrier Instance 1 Bounding box coordinates ... etc.

2. *Hi! I have a list of road object instances and their bounding box coordinates. Could you please analyse this data and find the main **spatial** relationships? Then, draw a graph diagram with these relationships. Use the listed objects as nodes and relationships - as arrow links. Please add text on nodes with the corresponding objects and instances and text on arrow links- with the corresponding relationship.*

Object barrier Instance 0 Bounding box coordinates ... Object barrier Instance 1 Bounding box coordinates ... etc.

3. *Thanks! Please create a colourful graph diagram for all the listed objects.*

In addition, it is worth to note that the priorly tested weaker models of ChatGPT such as GPT-4o performed significantly worse. While the o1 model did not ignore any objects and any request in the prompt, the weaker model missed some objects and produced much worse quality graphs (with missed objects and incorrect relationships). Therefore, it was excluded from further experiment.

Results and Discussion

This section discusses the results of relationship modelling using spatial analysis and natural language processing with LLMs.

Spatial Analysis

Figures 4 and 5 demonstrate detected topological relationships, particularly “closest to” relationships with red lines and “nearby”, or association, relationships with green lines. We can say that this approach is reliable as it is based on the straightforward calculation of distances and analysis of objects’ locations and thus provides precise spatial

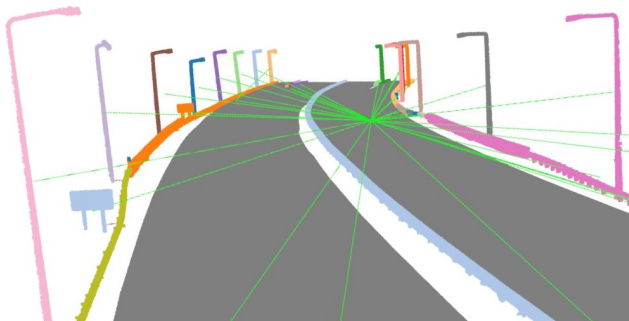
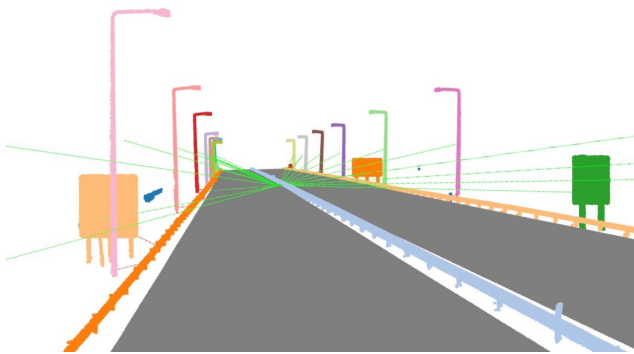
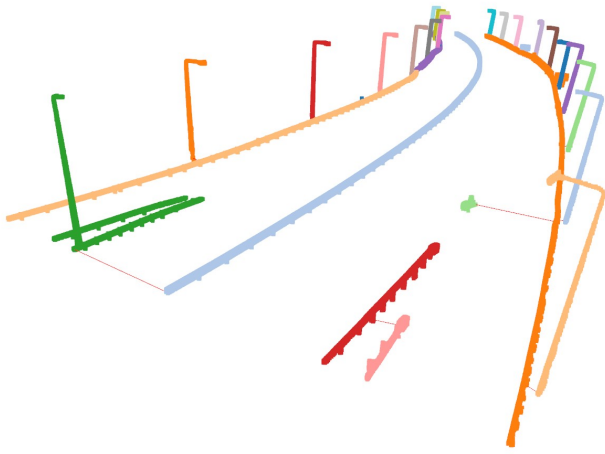


Figure 4: Examples of relationships modelled with spatial analysis using point cloud clusters of the GDT objects. Red lines – “closest to” relationships; Green lines – “nearby” or association relationships.

relationships that can be used further in GDT structuring and scene graphs creation.

Furthermore, these relationships can further be extended to more detailed spatial relationships, such as “above”, “on the side of”, “between” and others. However, this requires explicit calculation of spatial co-location.

LLM Analysis

Spatial relationships

Figure 6 presents a graph diagram generated by ChatGPT’s o1 model, illustrating the spatial analysis of objects and their relationships. Notably, the model accurately identifies each object’s spatial interrelationships and interprets bounding box descriptions using 3D coordinates

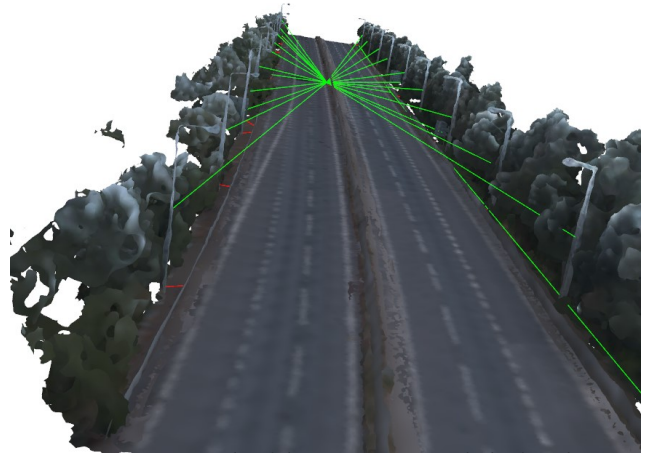


Figure 5: Example of the output GDT with the found topological relationships.

with “overlaps” and “on” relationships. All given objects were taken into account in this example (1 road segment, 20 lamps, 2 traffic signs, 8 barriers).

Furthermore, a concise and accurate summary was created by the LLM in addition to the graph output. We further provide this unedited summary text for the sample:

- “All barriers (instances 0–7) are on road 0.
- All traffic signs (instances 0–1) are on road 0.
- All lamps (instances 0–19) are on road 0.
- Multiple objects’ bounding boxes also overlap each other (e.g., barriers with barriers, barriers with signs, etc.), indicating they occupy closely adjacent or intersecting volumes in 3D.”

Unlike conventional spatial analysis methods constrained by predefined relationships, ChatGPT allows for extended or alternative analysis through additional prompts or re-runs of the same prompts. For example, the relations of being “very close” (when objects were touching each other), “close”, “near”, “far” with respect to the main road segment were also identified by the LLM in another run, with thresholds for closeness being also decided by the LLM. Furthermore, these prompts can be triggered automatically to produce deeper or more diverse insights.

Functional relationships

Next, we instructed ChatGPT o1 to explore functional relationships (Figure 7), thereby assessing its capacity to interpret the functional context of road environments. The model accurately linked each object instance to its corresponding road segment. For instance, it recognised that lamps *illuminate* the road, barriers *provide protection* or delineation, and traffic signs *convey information*. It also correctly identified every object along with its primary function relative to the road segment. When supplemented with additional context – such as whether the objects are operational or damaged – this graph can deliver significant value to asset managers.

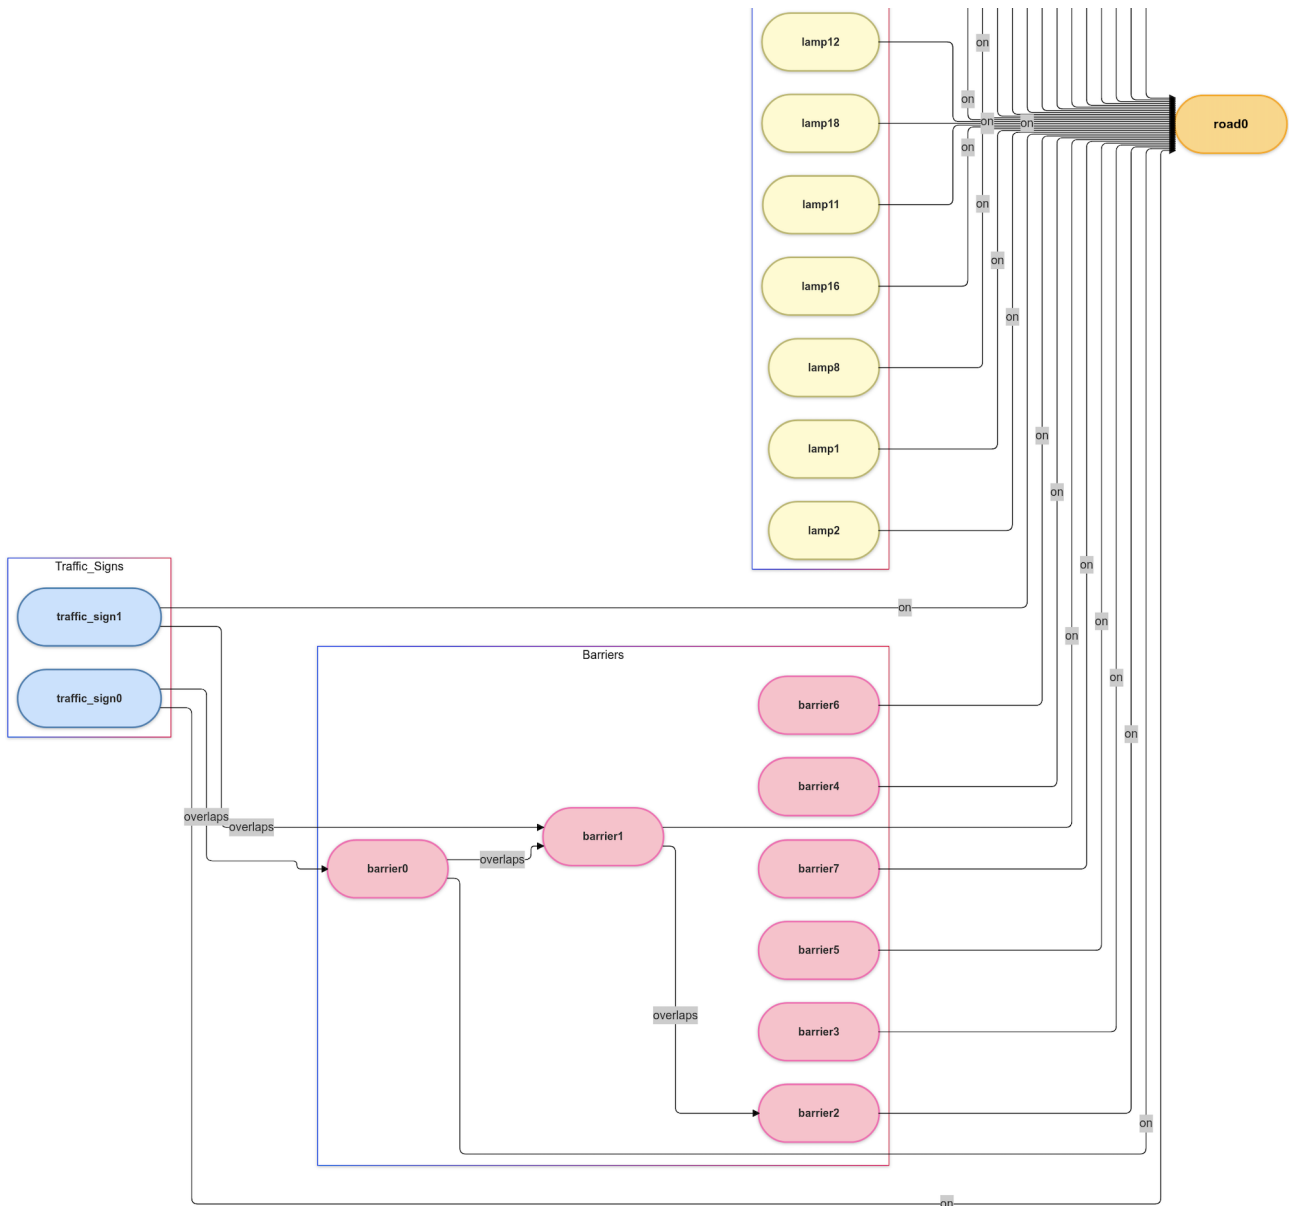


Figure 6: ChatGPT o1 prompt output (unredacted) for spatial relationships. Only a part of the output is shown as the image was cropped to fit the page. Visualised with Mermaid Chart.

Besides producing the graph output, this model was additionally able to provide an executive summary of the given data and its interpretation as follows (unredacted output for the sample):

“From the coordinate ranges given, it appears:

- All barrier instances (0–7) cluster around the same large region spanned by the single road instance (0). This suggests that each barrier is functionally “along” or “protecting”/“separating” the road.
- All lamp instances (0–19) also share overlapping X/Y extents with that same road region, at higher Z coordinates (suggesting elevation). In a typical scenario, lamps are “installed” above or next to the road for illumination.

- Traffic sign instances (0–1) likewise lie within the same broad coordinate space, consistent with being “placed on the roadside.”

Hence, at a high level, each of these objects (barriers, lamps, and traffic signs) functionally relates to the road.”

Summary

Overall, the outputs generated by the LLM demonstrate strong potential for geospatial applications, particularly in modelling relationships between road infrastructure elements. Remarkably, even without task-specific fine-tuning, this general-purpose language model exhibits accurate reasoning about spatial co-locations and functional relationships. Its ability to interpret and infer spatial connections suggests that LLMs can both complement and,

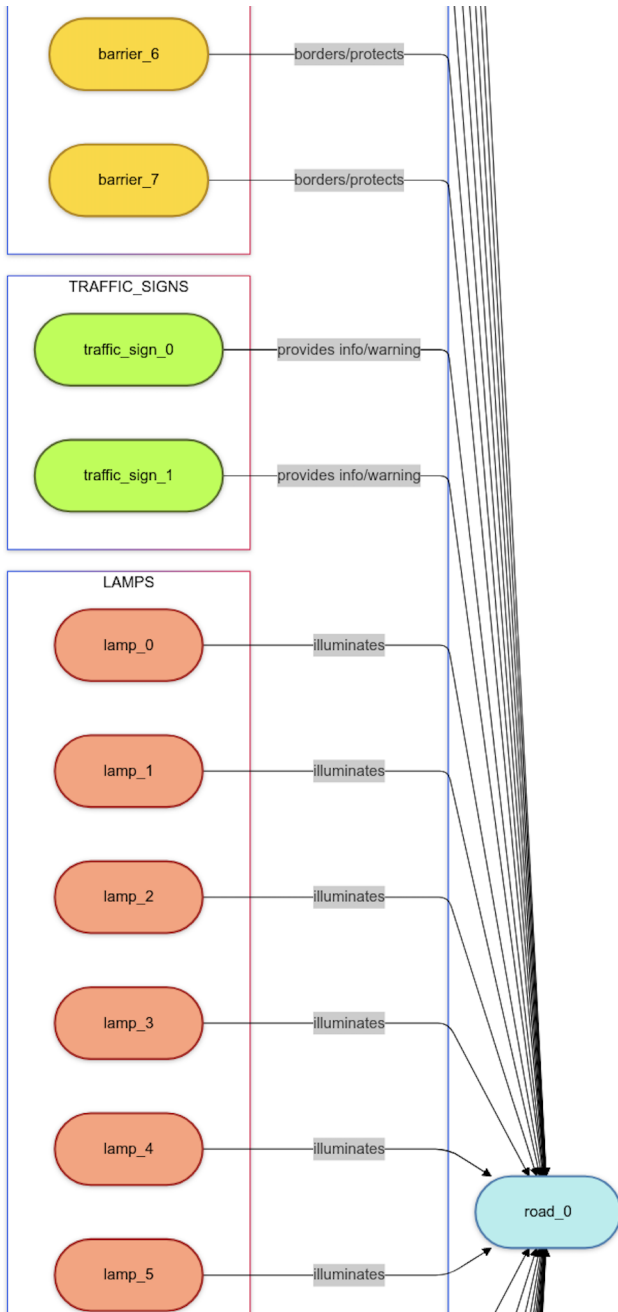


Figure 7: ChatGPT o1 prompt output (unredacted) for functional relationships. Only a part of the output is shown as the image was cropped to fit the page. Visualised with Mermaid Chart.

in some cases, serve as an alternative to traditional spatial analysis techniques. This capability holds significant promise for use in digital twinning and automated scene graph generation, where understanding the spatial and functional interplay between objects is essential.

Limitations

Although the proposed method successfully merges the precision and flexibility of two distinct approaches, some limitations remain. First, the lack of ground-truth labels for relationships restricts our evaluation to the following:

1. We evaluated the outputs generated by the LLM using spatial analysis, which serves as a valid and robust approach for assessing the quality and contextual accuracy of geospatial predictions. This method is grounded in fundamental spatial principles, such as distance measurements and the analysis of relative positions, allowing for objective and interpretable evaluation of the model’s spatial reasoning.
2. LLM-based functional relationships were evaluated manually by the authors of this paper, i.e., instances of object types were searched in the output and matched with their predicted function.

Future research should address this gap by annotating relationship data and including more relationship types, thereby enabling more rigorous evaluation and the supervised training of other deep learning models. Second, the current context is limited to geometric information without incorporating condition data. Expanding the contextual scope could further enhance both the accuracy and applicability of the method.

Conclusion

In this paper, we presented a novel approach to relationship modelling for road GDTs, integrating spatial analysis and LLMs to automate the identification of both topological and functional relationships among road assets. This hybrid method addresses a critical bottleneck in current practice – namely, the extensive manual labour required to model and maintain relationships – and thereby accelerates large-scale digitisation of road networks. By building upon our prior work on scene understanding and 3D reconstruction, we have demonstrated that it is possible to achieve a holistic, interconnected digital twin that offers not only accurate geometric representation but also meaningful semantic context.

Moreover, our approach is not limited by geographic regions and can potentially accommodate varying levels of asset complexity. It can be scaled to different road networks, offering a flexible, adaptable solution for diverse roadway environments.

The outcomes of our approach confirm the value of combining precise spatial techniques with the flexibility and interpretive power of LLMs. As a result, GDTs can now more effectively support advanced use cases, including offering more structured data, safety assessments and predictive maintenance. Our solution thus lays the groundwork for broader adoption and further innovation in digital twin technology for transportation and construction industries. Future research will benefit from expanding the contextual scope, e.g., incorporating asset condition data, and establishing annotated relationship datasets to enable supervised deep learning for various tasks such as enhancing anomaly detection, object detection and others. Ultimately, automating the modelling of road relationships paves the way for more efficient asset management, improved decision-making and a higher return on investment

in digital twin systems.

Acknowledgments



Funded by
the European Union

This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) Industrial CASE in partnership with National Highways and Costain [Grant No. EP/V056441/1]. This project also

has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101034337.

References

- Autodesk (2023). Autodesk. www.autodesk.co.uk/collections/architecture-engineering-construction.
- Botín-Sanabria, D. M., Mihaita, A.-S., Peimbert-García, R. E., Ramírez-Moreno, M. A., Ramírez-Mendoza, R. A., and Lozoya-Santos, J. d. J. (2022). Digital twin technology challenges and applications: A comprehensive review. *Remote Sensing*, 14(6).
- BuildingSMART (2025). Ifcrelationship. <https://ifc43-docs.standards.buildingsmart.org/IFC/RELEASE/IFC4x3/HTML/lexical/IfcRelationship.htm>.
- Davletshina, D. and Brilakis, I. (2024). A review on constructing and maintaining geometric digital twins of highways. *Transportation Research Board 103rd Annual Meeting/Transportation Research Board*, (TRBAM-24-00258).
- Davletshina, D., Reja, V. K., and Brilakis, I. (2024). Automating construction of road digital twin geometry using context and location aware segmentation. *Automation in Construction*, 168:105795.
- Drobný, V., Li, S., and Brilakis, I. (2024). Connectivity detection for automatic construction of building geometric digital twins. *Automation in Construction*, 159:105281.
- Jiang, F., Ma, L., Broyd, T., and Chen, K. (2021). Digital twin and its implementations in the civil engineering sector. *Automation in Construction*, 130:103838.
- Justo, A., Soilán, M., Sánchez-Rodríguez, A., and Riveiro, B. (2021). Scan-to-bim for the infrastructure domain: Generation of ifc-compliant models of road infrastructure assets and semantics using 3d point cloud data. *Automation in Construction*, 127:103703.
- Malawade, A. V., Yu, S.-Y., Hsu, B., Kaeley, H., Karra, A., and Al Faruque, M. A. (2022). *roadscene2vec*: A tool for extracting and embedding road scene-graphs. *Knowledge-Based Systems*, 242:108245.
- Mansourian, A. and Oucheikh, R. (2024). Chatgeoai: Enabling geospatial analysis for public through natural language, with large language models. *ISPRS International Journal of Geo-Information*, 13(10).
- Marie d'Avigneau, A., Potseluyko, L., Anvo, N. R., Taha, H. M., Reja, V. K., Davletshina, D., Lam, P., de Silva, L., Al-Tabbaa, A., and Brilakis, I. (2025). *Camhighways*: The cambridge highways dataset. *Advanced Engineering Informatics*, 64:103036.
- McKinsey Global Institute (2015). *Digital America: A tale of the haves and have-mores*.
- Open3D (2025). *Open3D: A Modern Library for 3D Data Processing*. <https://www.open3d.org/docs/release/introduction.html>.
- OpenAI (2025). *Learning to reason with LLMs*. <https://openai.com/index/learning-to-reason-with-llms/>.
- QGIS (2025). *QGIS*. <https://qgis.org/>.
- Sacks, R., Girolami, M., and Brilakis, I. (2020). Building information modelling, artificial intelligence and construction tech. *Developments in the Built Environment*, 4:100011.
- SciPy (2025). *kd-tree for quick nearest-neighbor lookup*. <https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.cKDTree.html>.
- Starace, L. L. L. and Di Martino, S. (2024). Can large language models automatically generate gis reports? In Lotfian, M. and Starace, L. L. L., editors, *Web and Wireless Geographical Information Systems*, pages 147–161, Cham. Springer Nature Switzerland.
- Tian, Y., Carballo, A., Li, R., and Takeda, K. (2020). Road scene graph: A semantic graph-based scene representation dataset for intelligent vehicles. *arXiv preprint arXiv:2011.13588*.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. (2023). *Llama 2: Open foundation and fine-tuned chat models*. *arXiv preprint arXiv:2307.09288*.