

UC Santa Barbara

Spatial Data Science Symposium 2023 Short Paper Proceedings

Title

ChatGPT as a mapping assistant: A novel method to enrich maps with generative AI and content derived from street-level photographs

Permalink

<https://escholarship.org/uc/item/64h832hd>

Authors

Juhász, Levente
Mooney, Peter
Hochmair, Hartwig H.
[et al.](#)

Publication Date

2023-09-05

DOI

10.25436/E2ZW27

Peer reviewed

ChatGPT as a mapping assistant: A novel method to enrich maps with generative AI and content derived from street-level photographs

Levente Juhász¹[0000-0003-3393-4021], Peter Mooney²[0000-0002-2389-3783],
Hartwig H. Hochmair³[0000-0002-7064-8238], and Boyuan Guan¹

¹ GIS Center, Florida International University, Miami, FL 33199, USA
{ljuhasz,bguan}@fiu.edu

² Department of Computer Science, Maynooth University, Co. Kildare, Ireland
peter.mooney@mu.ie

³ Geomatics Sciences, University of Florida, Ft. Lauderdale, FL 33144, USA
hhhochmair@ufl.edu

Abstract. This paper explores the concept of leveraging generative AI as a mapping assistant for enhancing the efficiency of collaborative mapping. We present the results of an experiment that combines multiple sources of volunteered geographic information (VGI) and large language models (LLMs). Three analysts described the content of crowdsourced Mapillary street-level photographs taken along roads in a small test area in Miami, Florida. GPT-3.5-turbo was instructed to suggest the most appropriate tagging for each road in OpenStreetMap (OSM). The study also explores the utilization of BLIP-2, a state-of-the-art multimodal pre-training method as an artificial analyst of street-level photographs in addition to human analysts. Results demonstrate two ways to effectively increase the accuracy of mapping suggestions without modifying the underlying AI models: by (1) providing a more detailed description of source photographs, and (2) combining prompt engineering with additional context (e.g. location and objects detected along a road). The first approach increases the accuracy of the suggestion by up to 29%, and the second one by up to 20%.

Keywords: ChatGPT · OpenStreetMap · Mapillary · LLM · volunteered geographic information · mapping

DOI: 10.25436/E2ZW27

1 Introduction

Generative artificial intelligence (AI) is a type of AI that can produce various types of content, including text, images, audio, code, and simulations. It has gained enormous attention since the public release of ChatGPT in late 2022. ChatGPT is an example of a Large Language Model (LLM), which is a form of generative AI that produces human-like language. Since the launch of ChatGPT,

researchers, including the geographic information science (GIScience) community, have been trying to understand the potential role of AI for research, teaching, and applications. ChatGPT can be used extensively for Natural Language Processing (NLP) tasks such as text generation, language translation, writing software code, and generating answers to a plethora of questions, engendering both positive and adverse impacts [2]. The emergence of generative AI has introduced transformative opportunities for spatial data science. In this paper, we explore the potential of generative AI to assist human cartographers and GIS professionals in increasing the quality of maps, using OSM as a test case (Figure 1).




Fig. 1. OpenStreetMap roads and Mapillary images in the study area near Downtown Miami

GeoAI has been part of the GIScience discourse in recent years. For example, Janowicz et al. [7] elaborated on whether it was possible to develop an artificial GIS analyst that passes a domain specific Turing test. Although these questions are still largely unanswered, utilizing LLMs and foundational models in geospatial contexts contributes to this direction. Despite the challenges due to the different nature of LLM training methodologies and human learning of spatial concepts [15], these tools and methods are being explored, for example to generate maps [18]. Our study fits into this direction by using an LLM (ChatGPT) and a multimodal pre-training method (BLIP-2) to connect visual and

language information in the context of mapping. We explore the larger question of whether *generative AI is a useful tool in the context of creating and enriching map databases* and more specifically investigate the following research questions:

1. Is generative AI capable of turning natural language text descriptions into the correct attribute tagging of road features in digital maps?
2. For this problem, can the accuracy of suggestions be improved through prompt engineering [17]?
3. To what extent can the work of human analysts be substituted with generative AI approaches within these types of mapping processes?

Furthermore, our approach focuses on the fusion of freely available volunteered geographic information (VGI) [5] data sources (OSM, Mapillary) and off-the-shelf AI tools to present a potentially low-cost and uniformly available solution. OSM is a collaborative project that aims to create a freely accessible worldwide map database (<https://openstreetmap.org>), while Mapillary crowdsources street-level photographs (<https://mapillary.com>) that power mapping and other applications, such as object detection, semantic segmentation and other computer vision algorithms to extract semantic information from imagery [3]. While the use of VGI has not yet been explored in the context of generative AI, previous studies demonstrated the practicability of combining multiple sources of VGI to improve the mapping process [13]. More specifically, Mapillary street-level images are routinely used to enhance OSM [8,10].

2 Study setup

2.1 Data sources and preparation

In OSM, geographic features are annotated with key-value pairs to assign the correct feature category to them, a process called tagging [14]. For example, roads are assigned a `"highway"=<value>` tag where `<value>` indicates a specific road category, such as `"residential"` for a residential street.


OSM data is not homogeneous, and individual users may perceive roads differently and therefore assign different `"highway"` values to the same type of road. A list of `"highway"` tag values was established to better describe the meaning of each road category in OSM. Furthermore, the difference between some road categories, for example, primary and secondary, is more of an administrative nature rather than visual appearance. For example, a 2-lane road in rural areas could be considered primary, whereas a more heavily trafficked road in an urban environment might be categorized secondary. To consider semantic road categories rather than individual `"highway"` values as one of the evaluation methods, `"highway"` tag values representing similar roads in our dataset were grouped into four categories (Table 1).

Figure 2 shows the methodology to obtain OSM roads of interest with the corresponding Mapillary street-level images. First, all OSM roads with a `"highway"=*`

Table 1. Grouping distinct "highway" tag values into semantically similar categories.

Category name	OSM "highway"	# of roads
Major, access controlled road	motorway trunk	0
Main road	primary secondary tertiary	81
Regular road	residential unclassified service	4
Not for motorized traffic	pedestrian footway cycleway	9

tags were extracted within the study area. Then, short sections (<50m), inaccessible roads, sidewalks along roadways and roads without street-level photo coverage were excluded. Retained OSM roads were matched with corresponding Mapillary photographs, so that each road segment would have at least one representative Mapillary image. Lastly, a list of objects detected in the corresponding image was also extracted from the Mapillary API. These inputs were further used as described in Section 2.3.

**Fig. 2.** Workflow for preparing input from Mapillary

2.2 Resources

AI tools and models We utilize GPT-3.5-turbo, which is an advanced language model developed by OpenAI. It is an upgraded version of GPT-3, designed to offer improved performance and capabilities, and retains the large-scale architecture of its predecessor, enabling it to generate coherent and contextually relevant text [16]. GPT-3.5-turbo serves as a powerful tool for natural language processing, content generation, and other language-related applications. In our study it is used to suggest OSM tagging based on pre-constructed prompts using the content of street-level images. The model was accessed through the OpenAI API.

BLIP-2 [11] is a state-of-the-art, scalable multimodal pre-training method, designed to equip LLMs with the capability to understand images while keeping their parameters entirely frozen. This approach is built upon leveraging frozen pre-trained unimodal models and a proposed Querying Transformer (Q-Former), sequentially pre-trained for vision-language representation learning and vision-to-language generative learning. Despite operating with fewer trainable parameters, BLIP-2 has achieved exceptional performance in a variety of vision-language tasks and has shown potential for zero-shot image-to-text generation [11]. The methodology proposed in BLIP-2 contributes towards the development of an advanced multimodal conversational AI agent. We leverage BLIP-2’s capability to generate image captions as well as to perform visual question-answering (Q&A). A freely available sample implementation of BLIP-2 was used to conduct this experiment (<https://replicate.com/andreasjansson/blip-2>).

Analysts Three analysts were tasked to describe the visual content of street-level images (captioning), and to answer a few questions regarding the image content (visual Q&A). Two (human) analysts were undergraduate students at Florida International University with previous GIS coursework. BLIP-2 was used to perform the same task as human analysts, and its responses were recorded as the third (artificial) analyst. Analysts were deliberately not given any guidelines as to how to describe images so that their answers would not be biased by prior knowledge about OSM and mapping. Table 2 lists questions and tasks performed by analysts.

Table 2. Questions and tasks performed by analysts.

Variable	Question/task	Example response
"caption"	Describe what you see in the photo in your own words.	<i>A city road in an urban area along an elevated railway. There is a wide sidewalk on both sides and trees on the left.</i>
"users"	Who are the primary users of the road that is located in the middle of the photograph? Cars, pedestrians or bicyclists?	<i>Cars</i>
"lanes"	How many traffic lanes are there on the road that is in the middle of the photograph?	<i>3</i>
"surface"	What is the material of the surface of the road that is in the center of the photograph?	<i>Asphalt</i>
"oneway"	Is the road that is in the center of the photograph one-way?	<i>No</i>
"lit"	Are there any street lights in the photograph?	<i>Yes</i>

The answers of analysts differ in level of detail. For example, BLIP-2’s and Analyst #2’s captions were significantly shorter on average (9 and 11 words, respectively) than Analyst #1’s (37 words). BLIP-2’s responses were also found to be more generic (e.g. "a city street with tall buildings in the background") than human analysts’. This allows us to explore the effect of providing increasing detail on tag suggestion accuracy.

2.3 Methodology for suggesting OSM tags

Figure 3 shows the methodology for suggesting tags for an OSM road. For each retained road in the area, the corresponding Mapillary images were shown to analysts described in Section 2.2. Analysts created an image caption and answered simple questions as described in Table 2. These responses in combination with additional context were used to build prompts for an LLM to suggest OSM tags. To explore what influences the accuracy of suggested tags, a series of prompts were developed that differ in the level of detail that is presented to the LLM.

All prompts start with the following message that provides context and instructs the model about the expected output format.

Based on the following context that was derived from a street-level photograph showing the street, recommend the most suitable tagging for an OpenStreetMap highway feature. Omit the 'oneway' and 'lit' tags if the answer to the corresponding questions is no or N/A. Format your suggested key-value pairs as a JSON. Your response should only contain this JSON.

The remainder of individual prompts is organized into four scenarios constructed from the responses of analysts and additional context. Example responses from analysts and additional context are highlighted in bold.

The **Baseline scenario** uses only responses from analysts, and contains the following text in addition to the common message above:

*The content of the photograph was described as follows: **A city road in an urban area along an elevated railway. There is a wide sidewalk on both sides and trees on the left. The road is mainly used by: cars. The surface of the road is: asphalt.***

*When asked how many traffic lanes there are on the road, one would answer: **3.***

*When asked if this street is a one-way road, one would answer: **No.***

*When asked if there are any street lights in the photograph, one would answer: **Yes.***

The **Locational context (LC) enhanced scenario** provides ChatGPT with additional locational context that describes where the roads in questions are located. In addition to the baseline message, it contains the following:

*The photograph was taken near **Downtown Miami, Florida.***

The Object detection (OD) enhanced scenario uses a list of detected objects in addition to the baseline:

*When guessing the correct category, consider that the following list of objects (separated by semicolon) are present in the photograph: **Temporary barrier; Traffic light - horizontal; Traffic light - pedestrian; Signage***

Finally, Object detection and locational context (OD + LC) are combined into a new scenario that supplies both additional contexts for the language model.

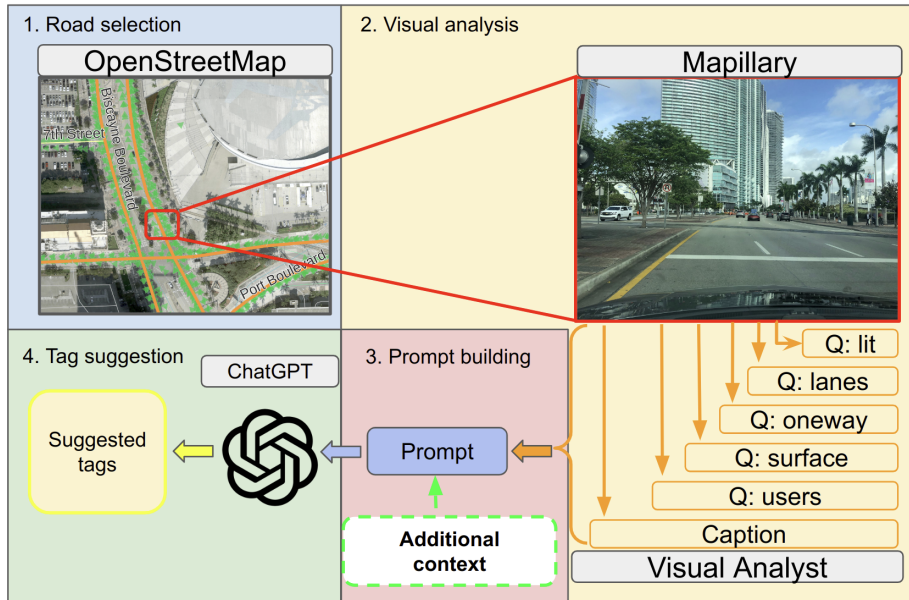


Fig. 3. Workflow of using ChatGPT to suggest OSM "highway" tags

The last step in the process is to supply the prompts described above to GPT-3.5-turbo (ChatGPT for simplicity). The model responds with a JSON document containing the suggested OSM tagging for the roadway, e.g. ("highway"="primary", "lanes"= 3), which can be compared to the original OSM tags of the same roadway.

The final dataset contains 94 OSM highway features and their original tags. For four scenarios and three analysts described above, ChatGPT recommendations based on the corresponding prompts were also recorded for the same roadway, resulting in a total of 12 tagging suggestions. These suggestions are then compared to the original OSM tags to assess the accuracy of a particular scenario and analyst.

3 Results

3.1 Accuracy of suggesting road categories

Table 3 lists the correctness of ChatGPT suggested road categories based on two different methods. First, we consider historical "highway" values of an OSM road. A suggestion was considered correct if the current or any previous versions of the corresponding OSM highway value matched the suggested tag of ChatGPT. This step takes into account differences in how individual mappers may perceive road features (e.g. primary vs. secondary). The second method is based on semantic road categories listed in Table 1. Considering groups of roads as opposed to individual "highway" values mitigates the fact that OSM tagging often follows administrative roles that are difficult to infer from photographs. Table 3 reports the accuracy of individual analysts across the four scenarios as well as the average correctness for analysts (values on bottom) and scenarios (values in different rows).

Table 3. Accuracy score of OSM tags suggested by ChatGPT. (LC = Locational context, OD = Object detection, OD + LC = Object Detection + Locational context)

Based on historical "highway" values						
Scenario	BLIP-2	Analyst #1	Analyst #2	Avg. correct [%]	% change	
Baseline	23.4	37.2	31.9	30.8	-	
LC	24.5	46.8	34.0	35.1	+4.3	
OD	27.7	47.9	39.4	38.3	+7.5	
OD + LC	30.9	47.9	50.0	42.9	+12.1	
Avg. correct [%]	26.6	45.0	38.8			
Based on semantic road categories						
Scenario	BLIP-2	Analyst #1	Analyst #2	Avg. correct [%]	% change	
Baseline	25.5	54.3	41.5	40.4	-	
LC	35.1	64.9	45.7	48.6	+8.2	
OD	29.8	63.8	60.6	51.4	+11.0	
OD + LC	43.6	66.0	70.2	59.9	+19.5	
Avg. correct [%]	33.5	62.3	54.5			

BLIP-2 achieved the lowest accuracy among the three analysts, followed by Analysts #2 and #1 respectively. This resembles the level of detail analysts described photographs with, which suggests that in general, providing more detailed image captions may lead to more accurate tag suggestions by ChatGPT. On average, this method increased the accuracy by up to 28.8% between BLIP-2 (lowest detail) and Analyst #2 (highest detail).

This is further supported by the average accuracy achieved in different scenarios. The baseline scenario, which used prompts purely based on the visual description of street-level photographs, achieved a suggestion accuracy of 30-40% on average from the three analysts. Providing additional context in different scenarios increased this accuracy. Additional location context, i.e. specifying that

the roads are located near Downtown Miami (**LC scenario**) increased suggestion accuracy by 4.3-8.2% on average, depending on the evaluation method. This can potentially be explained by regional differences in OSM tagging practices which are usually determined by local communities. In this scenario, it is possible that the AI model considered these regional differences when suggesting "highway" tags. Providing a list of objects detected in the source photographs (**OD scenario**) increased the average suggestion accuracy by 7.5 - 11.0% compared to the baseline scenario. A potential explanation for this is that objects found on and near roads provide important details that help refine the category of roads. Finally, combining additional locational and object detection contexts (**OD + LC scenario**) with the description of photographs by analysts increases suggestion accuracy by 12.1 - 19.5% on average. It is important to mention that these improvements are observed across all analysts.

3.2 Additional tag suggestions

In addition to the main "highway" category, information about additional characteristics of roads can also be recorded in OSM. To assess such a scenario, we analyze the "lit" tag, which indicates the presence of lighting on a particular road segment. The "lit" tag is set to "yes" if there are lights installed along the roadway. One question explicitly asked analysts whether street lights are visible on street-level photographs. In addition, street lights are a potential object category in Mapillary detections. For the following analysis, we consider the **Object Detection enhanced scenario**. The original dataset contains 24 roadways with "lit"="yes".

Table 4 shows that ChatGPT correctly suggested the presence of "lit" tag between 63% (BLIP-2) and 92% (Analyst #2) of the existing cases. ChatGPT suggested the use of the "lit" tag for an additional 44 - 61 features that are potentially missing from OSM. Among these features, 59 have been suggested based on prompts from at least two analysts, and 36 have been suggested based on all three analysts.

Table 4. ChatGPT suggestions of the "lit" tag.

	BLIP-2	Analyst #1	Analyst #2
Correctly tagged	15 (63%)	20 (83%)	22 (92%)
Additional	58	61	44

3.3 Limitations of the results

One limitation of this study is that we conducted the experiment on a small, geographically limited sample size. The implications of these are well studied in the GIScience literature, such as the uneven coverage of street-level photographs

[9] and the heterogeneity of OSM tagging [4], which could limit the adaptability of our method to different areas. The other group of limitations is largely related to open problems in Computer Science, such as the so-called “hallucinations” of generative AI, which results in content that is false or factually incorrect [1], as well as the non-deterministic nature of ChatGPT’s answers [19].

4 Summary and discussion

The study employs ChatGPT as a mapping assistant to propose OSM road tags based on textual descriptions of street-level photos. It leverages freely available geospatial data (OSM, Mapillary) and off-the-shelf models (GPT-3.5-turbo, BLIP-2) to enhance collaborative mapping. ChatGPT accurately suggests OSM “highway” values from text in 39-45% cases, rising to 55-62% for semantic road categories. Substituting human analysts with BLIP-2 yielded less accurate results (27-34%), potentially due to short captions. The study explores prompt engineering and context addition for improved accuracy. Providing road location raised accuracy by 4-8%, object lists boosted it by 8-11%, and combining both enhanced it by 12-20%. The study also found that increasing the level of detail with which a road scene is described in a prompt increases the accuracy of the suggested highway tags. However, it is expected that there are limitations of this method in terms of the accuracy that can be theoretically achieved. Future research will expand OSM sample size and incorporate refined traffic-related data (e.g. speed limits, turn restrictions).

Although experiments like this are useful initial steps, we urge the GIScience community to go beyond simply applying AI in geographic contexts and focus on synergistic research that advances both the spatial sciences and AI research (see, e.g. [12,6]). There are multiple potential extension of this work along this idea that go beyond the case study presented in this paper. For example, the exploration of a multimodal conversational AI agent for spatial data science is a promising research direction. In theory, incorporating a spatial understanding component in a multimodal AI system allows it to comprehend and analyze geospatial data. This could result in a method for the AI to interpret and interact with geospatial data, similar to how BLIP-2 enables language models to understand images. Future research should focus on exploring the potential of this integration and on deepening our understanding of the theoretical and practical aspects of this fusion. This, in turn will advance the field of research and lay the groundwork for future innovations in comprehensive multimodal AI systems for geospatial science.

Data availability The dataset supporting findings presented in this paper is freely available at: <https://doi.org/10.17605/OSF.IO/M9RSG>.

Acknowledgements Authors would like to thank Mei Hamaguchi and Flora Beleznay for providing image captions and visual Q&A.

References

1. Bang, Y., Cahyawijaya, S., Lee, N., Dai, W., Su, D., Wilie, B., Lovenia, H., Ji, Z., Yu, T., Chung, W., Do, Q.V., Xu, Y., Fung, P.: A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity (Feb 2023). <https://doi.org/10.48550/arXiv.2302.04023>
2. Dwivedi, Y.K., Kshetri, N., Hughes, L., Slade, E.L., Jeyaraj, A., Kar, A.K., Baabdullah, A.M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Al-bashrawi, M.A., Al-Busaidi, A.S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., Carter, L., Chowdhury, S., Crick, T., Cunningham, S.W., Davies, G.H., Davison, R.M., Dé, R., Dennehy, D., Duan, Y., Dubey, R., Dwivedi, R., Edwards, J.S., Flavián, C., Gauld, R., Grover, V., Hu, M.C., Janssen, M., Jones, P., Junglas, I., Khorana, S., Kraus, S., Larsen, K.R., Latreille, P., Laumer, S., Malik, F.T., Mardani, A., Mariani, M., Mithas, S., Mogaji, E., Nord, J.H., O'Connor, S., Okumus, F., Pagani, M., Pandey, N., Papagiannidis, S., Pappas, I.O., Pathak, N., Pries-Heje, J., Raman, R., Rana, N.P., Rehm, S.V., Ribeiro-Navarrete, S., Richter, A., Rowe, F., Sarker, S., Stahl, B.C., Tiwari, M.K., van der Aalst, W., Venkatesh, V., Viglia, G., Wade, M., Walton, P., Wirtz, J., Wright, R.: “so what if chatgpt wrote it?” multidisciplinary perspectives on opportunities, challenges and implications of generative conversational ai for research, practice and policy. *International Journal of Information Management* **71**, 102642 (2023). <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
3. Ertler, C., Mislej, J., Ollmann, T., Porzi, L., Neuhold, G., Kuang, Y.: The Mapillary Traffic Sign Dataset for Detection and Classification on a Global Scale. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) *Computer Vision – ECCV 2020*. pp. 68–84. *Lecture Notes in Computer Science*, Springer International Publishing, Cham, Switzerland (2020). https://doi.org/10.1007/978-3-030-58592-1_5
4. Girres, J.F., Touya, G.: Quality assessment of the French OpenStreetMap dataset. *Transactions in GIS* **14**(4), 435–459 (2010). <https://doi.org/10.1111/j.1467-9671.2010.01203.x>
5. Goodchild, M.F.: Citizens as sensors: the world of volunteered geography. *GeoJournal* **69**(4), 211–221 (2007). <https://doi.org/10.1007/s10708-007-9111-y>
6. Janowicz, K.: Philosophical foundations of geoai: Exploring sustainability, diversity, and bias in geoai and spatial data science (2023). <https://doi.org/10.48550/arXiv.2304.06508>
7. Janowicz, K., Gao, S., McKenzie, G., Hu, Y., Bhaduri, B.: Geoai: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *International Journal of Geographical Information Science* **34**(4), 625–636 (2020). <https://doi.org/10.1080/13658816.2019.1684500>
8. Juhász, L., Hochmair, H.H.: Cross-linkage between Mapillary Street Level Photos and OSM Edits. In: Sarjakoski, T., Santos, M.Y., Sarjakoski, T. (eds.) *Geospatial Data in a Changing World: Selected papers of the 19th AGILE Conference on Geographic Information Science*, vol. *Lecture Notes in Geoinformation and Cartography*, pp. 141–156. Springer, Berlin (2016). https://doi.org/10.1007/978-3-319-33783-8_9
9. Juhász, L., Hochmair, H.H.: User Contribution Patterns and Completeness Evaluation of Mapillary, a Crowdsourced Street Level Photo Service. *Transactions in GIS* **20**(6), 925–947 (2016). <https://doi.org/10.1111/tgis.12190>
10. Juhász, L., Hochmair, H.H.: How do volunteer mappers use crowdsourced Mapillary street level images to enrich OpenStreetMap? In: Bregt, A., Sarjakoski, T.,

- Lammeren, R.v., Rip, F. (eds.) *Societal Geo-Innovation : short papers, posters and poster abstracts of the 20th AGILE Conference on Geographic Information Science*. Wageningen, The Netherlands (2017)
11. Li, J., Li, D., Savarese, S., Hoi, S.: BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models (May 2023). <https://doi.org/10.48550/arXiv.2301.12597>
 12. Li, J., Li, D., Xiong, C., Hoi, S.: BLIP: Bootstrapping Language-Image Pre-training for Unified Vision-Language Understanding and Generation (Feb 2022). <https://doi.org/10.48550/arXiv.2201.12086>
 13. Liu, L., Olteanu-Raimond, A.M., Jolivet, L., Bris, A.I., See, L.: A data fusion-based framework to integrate multi-source vgi in an authoritative land use database. *International Journal of Digital Earth* **14**(4), 480–509 (2021). <https://doi.org/10.1080/17538947.2020.1842524>
 14. Mooney, P., Corcoran, P.: The Annotation Process in OpenStreetMap. *Transactions in GIS* **16**(4), 561–579 (2012). <https://doi.org/10.1111/j.1467-9671.2012.01306.x>
 15. Mooney, P., Cui, W., Guan, B., Juhász, L.: Towards Understanding the Spatial Literacy of ChatGPT – Taking a Geographic Information Systems (GIS) Exam (2023). <https://doi.org/10.31223/X5P38P>, earthArXiv
 16. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askill, A., Welinder, P., Christiano, P., Leike, J., Lowe, R.: Training language models to follow instructions with human feedback (Mar 2022). <https://doi.org/10.48550/arXiv.2203.02155>
 17. Reynolds, L., McDonell, K.: Prompt Programming for Large Language Models: Beyond the Few-Shot Paradigm. In: *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. pp. 1–7. CHI EA '21, Association for Computing Machinery, New York, NY, USA (May 2021). <https://doi.org/10.1145/3411763.3451760>
 18. Tao, R., Xu, J.: Mapping with chatgpt. *ISPRS International Journal of Geo-Information* **12**(7), 284 (2023). <https://doi.org/10.3390/ijgi12070284>
 19. Wang, S., Scells, H., Koopman, B., Zuccon, G.: Can ChatGPT Write a Good Boolean Query for Systematic Review Literature Search? (Feb 2023). <https://doi.org/10.48550/arXiv.2302.03495>