

HeatPrompt: Zero-Shot Vision-Language Modeling of Urban Heat Demand from Satellite Images

Kundan Thota, Xuanhao Mu, Thorsten Schlachter, Veit Hagenmeyer

Institute for Automation and Applied Informatics (IAI), Karlsruhe Institute of Technology (KIT), Germany

{kundan.thota, xuanhao.mu, thorsten.schlachter, veit.hagenmeyer}@kit.edu

Abstract—Accurate heat-demand maps play a crucial role in decarbonizing space heating, yet most municipalities lack detailed building-level data needed to calculate them. We introduce HeatPrompt, a zero-shot vision-language energy modeling framework that estimates annual heat demand using semantic features extracted from satellite images, basic Geographic Information System (GIS), and building-level features. We feed pretrained Large Vision Language Models (VLMs) with a domain-specific prompt to act as an energy planner and extract the visual attributes such as roof age, building density, etc, from the RGB satellite image that correspond to the thermal load. A Multi-Layer Perceptron (MLP) regressor trained on these captions shows an R^2 uplift of 93.7% and shrinks the mean absolute error (MAE) by 30% compared to the baseline model. Qualitative analysis shows that high-impact tokens align with high-demand zones, offering lightweight support for heat planning in data-scarce regions.

Index Terms—Heat demand estimation, VLMs, Remote sensing, Energy modeling, Semantic features, Zero-shot learning.

I. INTRODUCTION

Minimizing dependency on fossil fuels is a crucial step towards fighting climate change. Space heating accounts for a larger share of energy use in both the commercial and residential sectors, making it a key focus for decarbonization efforts. In Europe, the policy goal is to optimize energy supply and use to reduce greenhouse gas emissions by 40% by 2030 and 80% by 2050 [1], [2]. To achieve this, reducing heat demand through efficient upgrades to district heating systems and building retrofits is recognized as an effective strategy, given the carbon intensity of the current energy mix [3]. Assessing heat demand across the city is a crucial factor in selecting a decarbonization strategy. Certain high-precision heat demand evaluations depend on comprehensive geographic and energy datasets.

Nevertheless, key building attributes for energy consumption, such as construction year, geometry, insulation level, and heating system, are frequently incomplete, outdated, or restricted by data-protection regulations. To overcome these data gaps, prior work has relied on either bottom-up or top-down heat-mapping approaches [3]. Their accuracy, however, is constrained by the availability of comprehensive building metadata.

Therefore, researchers have pivoted to publicly available remote sensing satellite imagery, which has become a foundation




Fig. 1: Energy-relevant features such as solar-ready surfaces, green roof areas, reflective roofing, and surrounding vegetation that support heat demand estimation.

for urban analytics, from land-use classification and building-footprint extraction to population estimation. However, converting pixel data into quantitative energy metrics continues to present significant challenges [4]–[6]. As a proxy, other features, such as roof material, building density, and vegetation presence, derived from images as shown in Fig. 1 can be used to estimate heat demand.

To bridge this gap, we propose HeatPrompt, a zero-shot pipeline that extracts energy-relevant semantic features from satellite images and uses them to estimate continuous annual heat demand. Our main contributions include the following:

- **Zero-Shot Semantic Heat Mapping:** We introduce HeatPrompt, the first framework to model annual heat demand for municipal zones using RGB satellite images, binary isoline masks, and zero-shot semantic captions from a VLM, eliminating the need for manual labeling of energy-level features.
- **Interpretable Pipeline:** HeatPrompt not only outperforms Convolutional Neural Networks (CNN)-based regressors but also highlights human-readable semantic features that reveal the key drivers of heat demand, providing planners with transparent, actionable insights rather than opaque black-box predictions.
- **Reproducible Benchmark:** We provide open-source code ¹ for data preparation, stratified cross-validation, and paired t -tests on absolute errors, establishing the

*corresponding author kundan.thota@kit.edu

¹<https://github.com/kundanthota/HeatPrompt>

first vision-language benchmark for urban heat-demand estimation.

The paper is organized as follows: Section II reviews related work on energy demand models and remote sensing data. Section III details the data acquisition and the proposed method. Section IV shows the experimental results and an ablation study. In contrast, Section V concludes the paper and outlines the direction of future research.

II. RELATED WORK

A. Heat Demand Mapping and Simulation Approaches

Heat demand mapping serves as a basis for municipal heat transition planning, and traditional approaches fall into two categories: bottom-up and top-down models. Bottom-up models estimate thermal energy needs by aggregating detailed building-level features such as floor area, usage type, construction year, and insulation levels. For instance, Meha et al. [7] applied a bottom-up model for computing heating demand at the building level using geometric and usage parameters, then aggregating the results to a 100×100 m resolution map. These methods offer high spatial accuracy but are constrained by their reliance on comprehensive local data, which is often missing or outdated.

Top-down models address data scarcity by disaggregating aggregate heat consumption to finer scales using features such as land-use, population density, or floor space distribution. Novosel et al. [3] demonstrated such a top-down mapping methodology for data-poor areas by calibrating national consumption data with minimal local information. Recent strategic studies and national data reports highlight that accurate spatial heat demand information is key to cost-efficient, optimised district heating networks. [1], [8], [9]

B. Machine Learning for Energy Demand Prediction

Machine learning techniques are increasingly used to estimate energy consumption and heating demand, particularly when data is missing in simulation models. Early works explored regression techniques and time-series models such as SARIMA to predict heat loads [10]. More recent studies utilize deep neural networks (DNNs), Long Short-Term Memory (LSTMs), and CNNs for short-term load forecasting [11], leveraging meteorological and historical sensor data.

While these methods perform well, they often depend on historical data, limiting their use for estimating energy consumption in unmeasured buildings. To address this, Gao et al. [12] and Morteza et al. [13] applied transfer learning to estimate building energy consumption in data-scarce settings. However, these techniques require at least partial building-level features such as insulation levels or age.

C. Remote Sensing and VLMs in Urban Analytics

Remote sensing has been used for urban classification, land cover analysis, and socio-economic inference [5], [12]. High-resolution aerial and satellite images are now routinely used to identify building footprints, roof types, and vegetation

coverage. Wang et al. [6] leverage UNetFormer for efficient semantic segmentation, enabling pixel-level information on rooftops, infrastructure, and urban context, which is used for various energy-relevant applications.

More recently, foundation models such as OpenAI’s Contrastive Language-Image Pre-training (CLIP) [14] have enabled powerful vision-language alignment via contrastive pre-training. CLIP and its geospatial adaptations can perform zero-shot classification of satellite images by matching them against textual descriptions, eliminating the need for task-specific labels. RemoteCLIP, trained on a large corpus of remote sensing image-text pairs, significantly outperforms generic VLMs in scene understanding tasks and enables the extraction of rich semantic cues from RGB imagery [15]–[17]. Instruction-tuned VLMs such as Large Language and Vision Assistant (LLaVA) [18] and multimodal LLMs like GPT-4o [19], Qwen2.5-VL [20] push this further by enabling expert-like reasoning over images through tailored prompts.

Despite their proven effectiveness in remote sensing and general vision tasks, these models have not been widely applied in energy modeling. Our research addresses this gap by explaining that features extracted from VLMs, such as roof condition, building age, and urban density, meaningfully correlate with various energy-related parameters. These insights serve as a semantic proxy, enriching estimation processes by providing valuable information that is typically difficult to quantify without extensive field surveys. By leveraging these inferred features, we enhance traditional models, particularly in urban areas characterized by diverse structures and visual characteristics.

III. PROPOSED METHOD

A. Dataset Acquisition

Satellite images are captured using the Esri World Imagery basemap [21], and municipal heat planning isolines are overlaid onto these images as shown in Fig. 2. The objective of this study is to predict the annual heat demand based on the resulting satellite imagery.

Accordingly, each sample consists of an RGB image and a corresponding binary mask representing the region of interest, defined by a multipolygon isoline. As an example data source, we use the Rheinland-Palatinate (RLP) Energy Atlas [22], which provides municipal heat consumption isolines as geospatial multipolygons along with associated metadata, including the annual heat demand y_i (MWh/a), planimetric area A_i (m²), and perimeter length ℓ_i (m).

Each polygon P_i is reprojected to the Web Mercator coordinate reference system (EPSG:3857), following standard geospatial preprocessing practices [23]. We compute its centroid and define a square sampling window of side

$$s_i = 1.05 \max\{\text{width}(P_i), \text{height}(P_i)\}$$




Fig. 2: Illustration of HeatPrompt’s semantic feature extraction. An isoline mask (red boundary) is first overlaid on the 512×512 px satellite image, defining the region of interest. Then, a VLM is prompted to extract visual attributes and associated confidence scores from this RGBA composite. These semantic features are then used for heat-demand regression.

to introduce a 5% contextual margin and capture a 512×512 px RGB satellite Image I_i , ensuring sub-meter ground sampling distance.

To isolate the region of interest, we rasterize P_i at the exact resolution, producing a binary mask $M_i \in \{0, 1\}^{512 \times 512}$ where $M_i(x, y) = 1$ only if the pixel lies inside the isoline. The resulting dataset consists of 1,677 samples with aligned triplets I_i, M_i, y_i , along with per-polygon GIS attributes (A_i, ℓ_i) , and forms the foundation for feature extraction and regression. In addition to polygon-level metadata, we enrich the region of interest with detailed building-level footprints extracted from OpenStreetMap (OSM) such as building types (e.g: Single Family House (SFH), Multi Family House (MFH), Terraced Row House (TR), Non-Residential (NR)) [24], Census 2011 [25], supplemented by LOD2 building height wireframes [26], as shown in the Algorithm 1. By mapping these features with each isoline, we construct a high-resolution dataset that serves as both input and ground truth for model validation.

B. HeatPrompt: Semantic Feature Extraction

For each 512×512 satellite image I_i and its aligned binary mask M_i , we construct an RGBA composite by inserting M_i into the alpha channel, thereby highlighting the region of interest. This composite image is fed to a VLM under a domain-specific prompt that instructs the model to act as a municipal heat planner, as shown in Fig. 2, and list the five most salient visual factors affecting heating demand. The resulting output is transformed into a fixed-length vector $z_i^{\text{sem}} \in \mathbb{R}^{512}$ using Nomic text-embeddings [27]. These semantic embeddings capture human-interpretable features, such as building density, roof characteristics, and vegetation cover, along with the prediction confidence level, that can be fed into a regressor without any additional fine-tuning or pixel-level supervision.

C. Regression on Semantic Embeddings

Let $\mathbf{x}_i^{(0)} = [A_i, \ell_i] \in \mathbb{R}^2$ denote the GIS attributes from the RLP Energy Atlas, and let $z_i^{\text{sem}} \in \mathbb{R}^{512}$ be the semantic embedding derived from the VLM captions. Additionally, let $\mathbf{x}_i^{\text{comp}} \in \mathbb{R}^d$ represent the building composition features derived from LOD2 geometry, OpenStreetMap tags, and census metadata. These features include aggregated descriptors such

Algorithm 1 Construction of Aligned Heat Demand Dataset with Auxiliary Attributes

Input: Polygons $\{P_i\}$ from Energicatlas with metadata y_i, A_i, ℓ_i

Output: Triplets $\{I_i, M_i, y_i\}$ with auxiliary building-level attributes

for all P_i **do**

Reproject $P_i \rightarrow \text{EPSG:3857}$

Define sampling window of side

$s_i = 1.05 \cdot \max\{\text{width}(P_i), \text{height}(P_i)\}$

Fetch RGB satellite image $I_i \in \mathbb{R}^{512 \times 512 \times 3}$

Rasterize binary mask $M_i \in \{0, 1\}^{512 \times 512}$ from P_i

end for

for all buildings $B_j \subseteq P_i$ **do**

Extract OSM tags $\rightarrow$ estimate use-type

$t_j \in \{\text{SFH, TH, MFH, NR}\}$

Link LOD2 roof wireframe $\rightarrow$ compute height h_j

Nearest-neighbor match to census data $\rightarrow$ age class e_j

Compute floor area $f_j = \text{area}(B_j) \cdot \lfloor h_j / h_{\text{floor}} \rfloor$

end for

for all P_i **do**

Append derived features (A_i, ℓ_i) to $\{I_i, M_i, y_i\}$

end for

return Full dataset $\{I_i, M_i, y_i, A_i, \ell_i, \text{building features}\}$

as height, floor count, floor area, construction decade, and building type ratios per isoline region.

We define the full feature vector as:

$$\mathbf{x}_i^{\text{full}} = \begin{bmatrix} \mathbf{x}_i^{(0)} \\ z_i^{\text{sem}} \\ \mathbf{x}_i^{\text{comp}} \end{bmatrix} \in \mathbb{R}^{514+d},$$

where d is the dimensionality of the building composition vector (e.g., $d \approx 11$ after feature aggregation).

We train and evaluate a Multi-Layer Perceptron (MLP) with two hidden layers of 64 units each, a dropout rate of 0.2, and optimized using a learning rate of 1×10^{-4} with weight decay for up to 250 epochs with early stopping. The model is trained on the dataset

$$X_{\text{full}} = \{\mathbf{x}_i^{\text{full}}\}_{i=1}^N,$$

using stratified five-fold cross-validation (CV) based on heat-demand quintiles to ensure that each fold preserves the overall distribution of target values.

For each fold, we report the coefficient of determination (R^2) and Mean Absolute Error (MAE).

IV. RESULTS AND DISCUSSION

A. Structured Feature Regression Performance

Table I compares the performance of various regression models trained on structured building attributes using stratified 5-fold cross-validation. The bottom-up simulation baseline, based on physical energy modeling, achieved an average R^2 of 0.32 and MAE of 287.1, reflecting its limited predictive capacity in the absence of fine-grained or updated building information. Among traditional ML models, Ridge regression performed best ($R^2=0.47$), with a 46.9% improvement over the baseline. Tree-based methods such as Random Forests and Gradient Boosting yielded marginal improvements. A simple MLP model achieved the highest performance, with an average R^2 of 0.51 and an MAE of 239.0, resulting in a 59.4% increase over the simulation baseline, making it the base regressor for further ablation and feature expansion.




Fig. 3: Semantic trends across heat demand terciles. Visual signals, such as tree cover and roof condition, correlate with modeled demand.

Model	$R^2 (\pm \text{std})$	MAE ($\pm \text{std}$)	% R^2
Baseline simulation	0.32 ± 0.30	287.1 ± 56.4	—
Linear Regression [28]	0.43 ± 0.15	244.7 ± 39.3	+34.4%
Ridge Regression [29]	0.47 ± 0.15	240.9 ± 38.8	+46.9%
Random Forest [30]	0.35 ± 0.12	238.1 ± 42.6	+8.6%
HistGradBoost [31]	0.34 ± 0.41	237.8 ± 49.0	+6.3%
MLP [32]	0.51 ± 0.17	239.0 ± 37.7	+59.4%

TABLE I: **Baseline Regression Performance Using Structured Building Features.** Comparison of linear, tree-based, and neural regression models evaluated with 5-fold cross-validation.

Features	R^2	MAE	% R^2
Baseline simulation	0.32 ± 0.30	287.0 ± 56.4	—
+ ResNet50 [33]	0.59 ± 0.18	205.3 ± 27.3	+84.4%
+ EfficientNet-B0 [34]	0.58 ± 0.17	208.6 ± 32.6	+81.3%
+ DenseNet121 [35]	0.57 ± 0.18	213.3 ± 38.4	+78.1%
+ CLIP [14]	0.60 ± 0.18	205.4 ± 22.3	+87.5%
+ Qwen2.5-VL [20]	0.58 ± 0.22	210.4 ± 27.1	+81.3%
+ GPT-4o [19]	0.62 ± 0.17	200.7 ± 33.2	+93.7%

TABLE II: **Impact of Semantic Vision Features on Energy Demand Prediction Performance.** Ablation study comparing baseline simulation with vision-augmented models using 5-fold cross-validation.

B. Impact of Semantic Features from Vision Models

To evaluate the contribution of visual semantics, we enrich the structured features with embeddings from various pre-trained vision-language and CNN backbones. Table II reports the resulting performance. Each model uses the same MLP architecture and is evaluated via stratified 5-fold cross-validation. Across all configurations, augmenting the simulation baseline with latent visual features leads to substantial gains. GPT-4o-derived text embeddings show the most improvement, with an R^2 of 0.62 and MAE of 200.7, marking a 93.7% relative gain over the bottom-up simulation. CLIP also achieves competitive performance, strengthening the idea that semantically rich features can serve as proxies for metadata that can't be manually extracted, which are energy-oriented.

C. Semantic Trends and Visual Interpretability

In addition to numerical gains, the semantic features exhibit interpretable relationships with known energy drivers. As shown in Fig. 3, areas with high predicted heat demand exhibit lower vegetation presence, fewer solar panels, and higher structural density, all of which are visually inferred via GPT-4o.

D. Qualitative Comparison: Role of Visual Cues

Fig. 4 provides a qualitative comparison of two geometrically similar buildings. Despite having similar size and shape, they differ significantly in visual characteristics: one has a modern roof with high greenery, while the other has an old tiled roof and bare surroundings. Traditional simulation models would treat both identically. In contrast, our approach, through zero-shot VLM prompting, identifies such cues and yields



New Roof
High vegetation cover
Estimated Demand:
100 kWh/m²

Old Roof (tile)
Low vegetation cover
Estimated Demand:
150 kWh/m²

Fig. 4: Visual comparison of similar buildings showing semantic cues influencing heat demand.

distinctly different heat demand predictions, reflecting more realistic outcomes.

These results strengthen the value of using vision-based semantic features in such pipelines, making it easier and more automated for energy planners to consider as many factors as possible that directly correlate with energy demand. Meanwhile, embeddings such as CLIP achieve notable results, making them a practical substitute for many physical variables, such as roof condition and land cover, and even for masking features for data protection.

V. CONCLUSION

We introduce HeatPrompt, a zero-shot vision-language framework that uses GPT-4o-generated captions along with basic GIS and building-level features to estimate annual heat demand from satellite images. This addition of semantic features yields a consistent increase in the explained variance R^2 by 93.7% from (0.32 ± 0.30) to (0.62 ± 0.17) and cuts down the MAE by 30%. Unlike CNN-based regressors, HeatPrompt also produces human-readable tokens, such as building density, roof type, and vegetation cover, thereby offering a user-friendly energy modeling pipeline.

Looking forward, we will develop energy-aware visual-language embeddings, inspired by CLIP and tailored to urban-energy semantics. This will enable planners to directly retrieve attributes such as roof age, insulation properties, PV coverage, and vegetation shielding from satellite images, as well as indirect proxies for less observable factors like insulation quality. These embeddings will capture latent correlations among building morphology, construction era, and neighborhood context that influence thermal performance. These vectors can be integrated into energy demand models without task-specific retraining. In parallel, we will explore lightweight fine-tuning of promptable VLMs to focus the embedding space on energy-relevant factors. Together, these efforts aim to provide a reusable and interpretable embedding layer for data-light energy modeling that planners can query, visualize, and use.

ACKNOWLEDGMENTS

This work was carried out within the project "New Data for the Energy Transition" (NEED) [36] and was supported by BMW under Grant No. 03EN3077J. Portions of the manuscript text (e.g., wording suggestions and Grammar) were drafted with ChatGPT's Assistance. The authors reviewed, edited, and take full responsibility for all content.

REFERENCES

- [1] Xenia Malcher and Miguel Gonzalez-Salazar. Strategies for decarbonizing european district heating: Evaluation of their effectiveness in sweden, france, germany, and poland. *Energy*, 306:132457, 2024.
- [2] Annika Boldrini, JP Jiménez Navarro, WHJ Crijns-Graus, and Ma Van den Broek. The role of district heating systems to provide balancing services in the european union. *Renewable and Sustainable Energy Reviews*, 154:111853, 2022.
- [3] Tomislav Novosel, Tomislav Pukšec, Neven Duić, and Julije Domac. Heat demand mapping and district heating assessment in data-poor areas. *Renewable and Sustainable Energy Reviews*, 131:109987, 2020.
- [4] Yonghao Xu, Bo Du, Liangpei Zhang, Daniele Cerra, Miguel Pato, Emiliano Carmona, Saurabh Prasad, Naoto Yokoya, Ronny Hänsch, and Bertrand Le Saux. Advanced multi-sensor optical remote sensing for urban land use and land cover classification: Outcome of the 2018 ieee grss data fusion contest. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(6):1709–1724, 2019.
- [5] Jorge E Patino and Juan C Duque. A review of regional science applications of satellite remote sensing in urban settings. *Computers, Environment and Urban Systems*, 37:1–17, 2013.
- [6] Libo Wang, Rui Li, Ce Zhang, Shenghui Fang, Chenxi Duan, Xiaoliang Meng, and Peter M Atkinson. Unetformer: A unet-like transformer for efficient semantic segmentation of remote sensing urban scene imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 190:196–214, 2022.
- [7] Drilon Meha, Tomislav Novosel, and Neven Duić. Bottom-up and top-down heat demand mapping methods for small municipalities, case gllogoc. *Energy*, 199:117429, 2020.
- [8] Alicia Lerbinger, Ivalin Petkov, Georgios Mavromatidis, and Christof Knoeri. Optimal decarbonization strategies for existing districts considering energy systems and retrofits. *Applied Energy*, 352:121863, 2023.
- [9] Umwelt Bundesamt. Energieverbrauch für fossile und erneuerbare wärme, 2023.
- [10] Tingting Fang and Risto Lahdelma. Evaluation of a multiple linear regression model and sarima model in forecasting heat demand for district heating system. *Applied energy*, 179:544–552, 2016.
- [11] Lilli Frison, Simon Götzhäuser, Moritz Bitterling, and Wolfgang Kramer. Evaluating different artificial neural network forecasting approaches for optimizing district heating network operation. *Energy*, 307:132745, 2024.
- [12] Yuan Gao, Yingjun Ruan, Chengkuan Fang, and Shuai Yin. Deep learning and transfer learning models of energy consumption forecasting for a building with poor information data. *Energy and Buildings*, 223:110156, 2020.
- [13] Azita Morteza, Amir Abbas Yahyaiean, Marzieh Mirzaeibonehkhater, Sina Sadeghi, Ali Mohaimeni, and Saman Taheri. Deep learning hyperparameter optimization: Application to electricity and heat demand prediction for buildings. *Energy and Buildings*, 289:113036, 2023.
- [14] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmlR, 2021.
- [15] Xiang Li, Congcong Wen, Yuan Hu, and Nan Zhou. Rs-clip: Zero shot remote sensing scene classification via contrastive vision-language supervision. *International Journal of Applied Earth Observation and Geoinformation*, 124:103497, 2023.
- [16] Fan Liu, Delong Chen, Zhangqingyun Guan, Xiaocong Zhou, Jiale Zhu, Qiaolin Ye, Liyong Fu, and Jun Zhou. Remoteclip: A vision language foundation model for remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–16, 2024.

- [17] Hui Lin, Danfeng Hong, Shuhang Ge, Chuyao Luo, Kai Jiang, Hao Jin, and Congcong Wen. Rs-moe: A vision-language model with mixture of experts for remote sensing image captioning and visual question answering. *IEEE Transactions on Geoscience and Remote Sensing*, 2025.
- [18] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.
- [19] Raisa Islam and Owana Marzia Moushi. Gpt-4o: The cutting-edge advancement in multimodal llm. *Authorea Preprints*, 2024.
- [20] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- [21] Esri. World imagery. ArcGIS Online basemap, 2024. Sources: Esri, Maxar, Earthstar Geographics, CNES/Airbus DS, USDA, USGS, AeroGRID, IGN, and the GIS User Community.
- [22] Ministerium für Wirtschaft, Verkehr, Landwirtschaft und Weinbau Rheinland-Pfalz. Energieatlas rheinland-pfalz, 2025. Accessed: 2025-08-07.
- [23] Thomas F Sturm. Some principles of web mercator maps and their computation. 2020.
- [24] Peter Mooney, Marco Minghini, et al. A review of openstreetmap data. *Mapping and the citizen sensor*, pages 37–59, 2017.
- [25] Statistisches Bundesamt (Destatis). Zensus 2011 – bevölkerungsstand deutschland. <https://www.destatis.de/DE/Themen/Gesellschaft-Umwelt/Bevoelkerung/Bevoelkerungsstand/Glossar/zensus2011.html>, 2011. Zugriff: 09. Dezember 2025.
- [26] Bundesamt für Kartographie und Geodäsie (BKG). 3d-gebäudemodelle lod2 deutschland (lod2-de). <https://gdz.bkg.bund.de/index.php/default/3d-gebaedemodelle-lod2-deutschland-lod2-de.html>, 2025. accessed: 09 December 2025.
- [27] Zach Nussbaum, John X. Morris, Brandon Duderstadt, and Andriy Mulyar. Nomic embed: Training a reproducible long context text embedder, 2024.
- [28] Xiaogang Su, Xin Yan, and Chih-Ling Tsai. Linear regression. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4(3):275–294, 2012.
- [29] Gary C McDonald. Ridge regression. *Wiley Interdisciplinary Reviews: Computational Statistics*, 1(1):93–100, 2009.
- [30] Steven J Rigatti. Random forest. *Journal of insurance medicine*, 47(1):31–39, 2017.
- [31] Aleksei Guryanov. Histogram-based algorithm for building gradient boosting ensembles of piecewise linear decision trees. In *International Conference on Analysis of Images, Social Networks and Texts*, pages 39–50. Springer, 2019.
- [32] Marius-Constantin Popescu, Valentina E Balas, Liliana Perescu-Popescu, and Nikos Mastorakis. Multilayer perceptron and neural networks. *WSEAS Transactions on Circuits and Systems*, 8(7):579–588, 2009.
- [33] Sheldon Mascarenhas and Mukul Agarwal. A comparison between vgg16, vgg19 and resnet50 architecture frameworks for image classification. In *2021 International conference on disruptive technologies for multi-disciplinary research and applications (CENTCON)*, volume 1, pages 96–99. IEEE, 2021.
- [34] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
- [35] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [36] Markus Duchon et al. A platform ecosystem providing new data for the energy transition. In *DACH+ Energy Informatics Conference*, Lugano, Switzerland, October 2024.