

# MARL: Multi-scale Archetype Representation Learning for Urban Building Energy Modeling

Xinwei Zhuang\*, Zixun Huang\*, Wentao Zeng, Luisa Caldas  
University of California, Berkeley

{xinwei\_zhuang, zixun, wentao.zeng, lcaldas}@berkeley.edu

## Abstract

*Building archetypes, representative models of building stock, are crucial for precise energy simulations in Urban Building Energy Modeling. The current widely adopted building archetypes are developed on a nationwide scale, potentially neglecting the impact of local buildings' geometric specificities. We present Multi-scale Archetype Representation Learning (MARL), an approach that leverages representation learning to extract geometric features from a specific building stock. Built upon VQ-AE, MARL encodes building footprints and purifies geometric information into latent vectors constrained by multiple architectural downstream tasks. These tailored representations are proven valuable for further clustering and building energy modeling. The advantages of our algorithm are its adaptability with respect to the different building footprint sizes, the ability for automatic generation across multi-scale regions, and the preservation of geometric features across neighborhoods and local ecologies. In our study spanning five regions in LA County, we show MARL surpasses both conventional and VQ-AE extracted archetypes in performance. Results demonstrate that geometric feature embeddings significantly improve the accuracy and reliability of energy consumption estimates. Code, dataset and trained models are available on the project page: <https://github.com/ZixunHuang1997/MARL-BuildingEnergyEstimation>.*

## 1. Introduction

The built environment plays a significant role in global climate change, accounting for approximately 34% of worldwide final energy consumption in 2020 [21]. In light of the urgent worldwide issues concerning sustainability and energy efficiency, there has been a concentrated focus on research in Urban Building Energy Modeling (UBEM) [3]. UBEM has evolved into a crucial computational tool

that enables the analysis and forecast of energy consumption at an urban scale. Given the impracticability or infeasibility of modeling each individual building in minute detail, developing building archetypes for the study area has become a critical determinant for precision for urban scale energy estimation. These archetypes represent specific groups of buildings sharing similar characteristics and energy performance within the area of interest. Their strategic use in UBEM allows for efficient and effective energy modeling at the urban scale.

Recent studies have established global building archetypes [10, 18, 22]. However, many of these are designed at a country scale, depend on expert input [10], and don't integrate actual building geometry, leading to potential inaccuracies and limiting their scalability, particularly in areas where computational resources and data accessibility present challenges. Additionally, urban-scale energy estimation demands significant computational resources, often inaccessible to disadvantaged communities and neighborhoods. Such disparities can introduce data biases, curtail global applicability, and limit the UBEM's diversity mainly to major cities like Boston [4, 6] and San Francisco [6]. These limitations hinder the effectiveness of building archetypes in energy-efficient design and implementation.

In response, we present Multi-scale Archetype Representation Learning (MARL, Figure 1), a method designed to automate local building archetype construction through representation learning. Our proposed method addresses the aforementioned challenges by refining the essential elements of building archetypes for UBEM. Our main contributions are summarized as follows:

- We introduce building geometry into the building archetype construction process.
- We provided a cost-effective solution for Urban Building Energy Modeling (UBEM).
- We propose an image reconstruction-based framework that streamlines the archetype creation process.

---

\*Equal Contribution

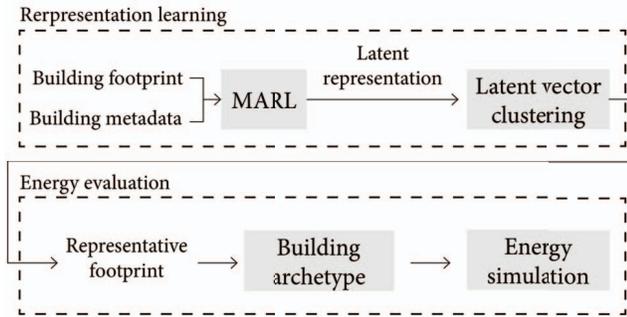


Figure 1. **Building Archetype Finding Pipeline.** Besides data collection, our method consists of 2 stages: archetype representation learning and clustering, and energy simulation.

- We introduce the integration of downstream tasks into the building archetype construction process, which increases the accuracy of our models, thereby enabling more precise and reliable energy consumption estimates at an urban scale.
- We conduct open-set experiments to demonstrate that our trained model possesses transferable and reusable characteristics in similar urban and geographical environments.

## 2. Related Studies

### 2.1. Representation learning

Representation learning offers effective strategies for dealing with high-dimensional data forms such as images and 3D models. Integral to these methods are a variety of neural compression algorithms [27], such as normalizing flows [24], variational autoencoders [16], diffusion probabilistic models [15], and generative adversarial networks [13]. Recent advancements have also seen variations of these models, such as Vector Quantized Variational Autoencoders (VQ-VAEs) [25] and Auto Decoders [20]. These techniques leverage the latent space as a compact representation where the high-dimensional data is encoded into lower-dimensional vectors, preserving the underlying features. A wealth of recent research showcases various application areas ranging from natural language processing to image classification [12, 14, 17].

### 2.2. Building archetype for energy estimation

Building archetypes are a representative subset of buildings to model an entire building stock, and are critical to the process of Urban Building Energy Modeling (UBEM). The prevailing methodologies for constructing the representative building archetype leverage statistical analysis or use generalized assumptions informed by expert knowledge

[5, 11], and remain a time-consuming task [11]. Existing archetypes define representative buildings based on categorical attributes such as construction material, vintage, area, stories, and energy system [10, 22]. Nevertheless, these models usually do not incorporate the actual geometries of buildings into the construction of building archetypes, thus potentially overlooking the nuanced variations and distributions of prevailing geometries among different building stocks.

Recent years have witnessed exploration into data-driven methods for the creation of building archetypes [3, 19, 23], but most of these methodologies continue to exclude building geometry parameters, largely as a result of technical intricacies [5]. Instead, the building geometry is simplified to numerical data, such as areas and shape ratio, with limited studies directly addressing the actual geometry [9]. Despite the significant impact of building geometry on energy consumption ([5, 7], its exclusion raises the potential for inaccuracies and unreliability for energy estimation in UBEM. This underscores the necessity for methods capable of incorporating building geometry to represent building stock comprehensively, thereby augmenting the precision and granularity of locally tailored building archetypes.

## 3. Method

Besides data collection, our method consists of 2 stages: archetype representation learning and clustering, and energy simulation. Our method follows the bottom-up approach for constructing a building archetype model [23] with a representation learning module to characterize the building stock. Multi-scale Architectural Representation Learning (MARL) is developed to encapsulate the implicit features of the input building stock into a compact, reduced-dimensional space, i.e., the latent space. Following this, we apply k-means clustering in the latent space to identify the representative building footprints, which are then converted into building archetypes with the same parameters documented in PBM [10].

We then conduct energy simulations on PBM and our building archetypes and compare the aggregated energy profiles. The aim is to investigate whether utilizing real-world building geometries for a specific district enhances the accuracy of urban-scale energy estimation and to evaluate the extent to which the accuracy of such energy modeling can be improved.

### 3.1. Archetype learning and clustering

Our model for representation learning (Figure. 2) consists of an auto-encoder and an optional downstream task pool (DTP). The auto-encoder aims at compressing and purifying the geometric contours of the building, while the downstream tasks add a restriction and punishment on it to conserve geometric information related to vital building

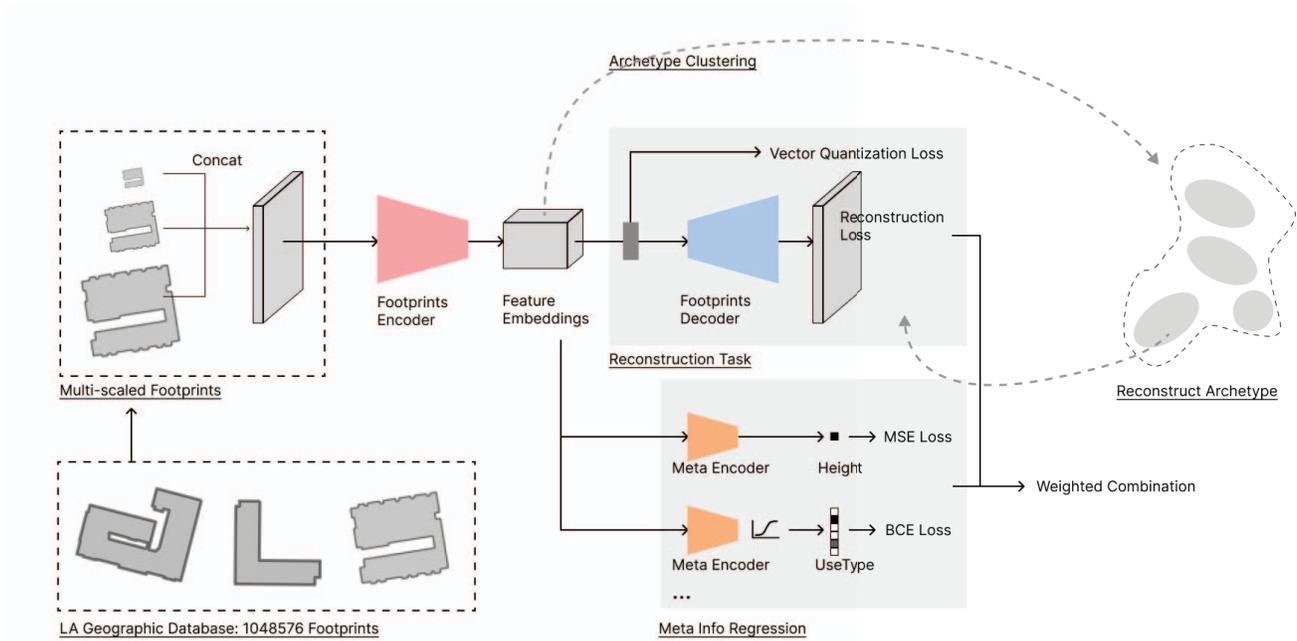


Figure 2. **Model Architecture for multi-scale archetype learning and clustering.** This stage consists of one auto-encoder and multiple downstream tasks.

properties. After the representation learning is finished, we then feed all footprints in a certain region back into our trained encoder and get the clustering centers of the corresponding latent vectors to acquire archetypes in this region.

### 3.1.1 Footprint reconstruction

Since the sizes of different footprints vary widely, to ensure that our model is flexible enough in its choice of data about building contours at different scales, a single footprint is scaled into 3 different sizes. Detailed implementation is described in Section 4.2. Both the encoder and decoder for the image reconstruction task are built with a CNN-based structure with residual connections. The scaled footprints are fused together and encoded into a latent vector with a dimension of  $28 \times 28 \times 32$  by the encoder. A vector quantizer [25] is then used to decouple the implicit representation and map it to discrete embedding vectors (a code book) prior to forwarding the vector to the decoder. During the quantization step, each continuous vector output from the encoder is matched to the nearest vector in the code book, which is then forwarded to the decoder. This helps to better control the quality and diversity of the generated data based on the clustering centers in further steps.

Based on VQ-AE, our image reconstruction process incorporates both reconstruction loss and vector quantization loss in the proposed representation learning. The training

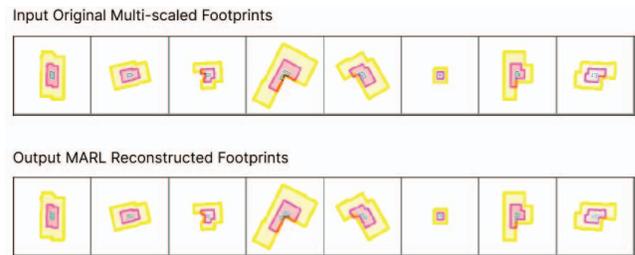


Figure 3. Examples of MARL footprint reconstruction.

objective for the reconstruction part becomes equation 1:

$$L = \log(p(x|z_q(x))) + \|sg(z_e(x)) - e\|_2^2 \quad (1)$$

The first term corresponds to the reconstruction loss, while the second term represents the codebook loss. Here, we use the notations  $x$ ,  $z_q(x)$ ,  $z_e(x)$ , and  $e$  to refer to the inputs, decoder inputs, encoder outputs, and embedding vectors, respectively. And the operator  $sg$  stands for the stop gradient operator.

### 3.1.2 Downstream task pool

Supervised by a downstream task pool (DTP), our MARL model excels in producing a more valuable latent space compared to normal auto-encoders, making it ideal for architectural purpose-driven applications. By leveraging the

insights from the additional tasks, our model is capable of preserving implicit geometry features closely related to crucial meta information of buildings.

The primary task focuses on predicting the building’s program or purpose, while the secondary task involves predicting the ‘vintage’ of a building, defined as the year it was constructed. Some other simple tasks are incorporated as well to gain further constraints on the encoder, such as predicting building heights from image greyscales. All tasks are meticulously tailored and hold a strong connection to traditional building energy modeling. Moreover, they directly or indirectly influence and even shape architectural geometry to a significant extent.

The first task aims at predicting the building program, specifically within the context of residential categories in Los Angeles. These categories are manifold, spanning from Mobile Homes, Units, and Rooming Houses to Apartments. The wide scope of categories caters to the diverse types of residential architecture prevalent within the region. For the second task, we operationalize the concept of ‘vintage’ as a sequence of categorical variables that encapsulate distinct temporal spans: pre-1980, 1980-2004, 2004-2013, and 2013 to the present. These divisions correspond with significant revisions to the California Building Code, serving as crucial epochs for analyzing the energy performance of structures. This segmentation is used for energy performance analysis [26].

By incorporating these tasks, our model is expected to gain a more nuanced understanding of building properties, thereby contributing to more precise and sophisticated predictions. These tasks reinforce the model’s robustness and increase its applicability across a range of building types.

### 3.2. Energy Simulation

After generating the representative footprints, we extrude them to their corresponding heights, which are coded in greyscale. We then apply energy parameters, including the window-wall ratio, U-value of the envelope, heating type, etc., consistent with the PBM [10]. Subsequently, we employ Climate Studio, an EnergyPlus plug-in for Rhino3D, to simulate the annual energy use intensity (EUI) of each representative building under the LA airport climate data sourced from [2]. After this, we aggregate the annual EUIs for each cluster based on their respective total areas and sum these to calculate the cumulative EUI for the entire neighborhood.

## 4. Experiments

### 4.1. Geographic dataset

We selected five distinct neighborhoods within Los Angeles County, each encompassing an area in excess of 100 km<sup>2</sup> and comprising a dense residential building popula-

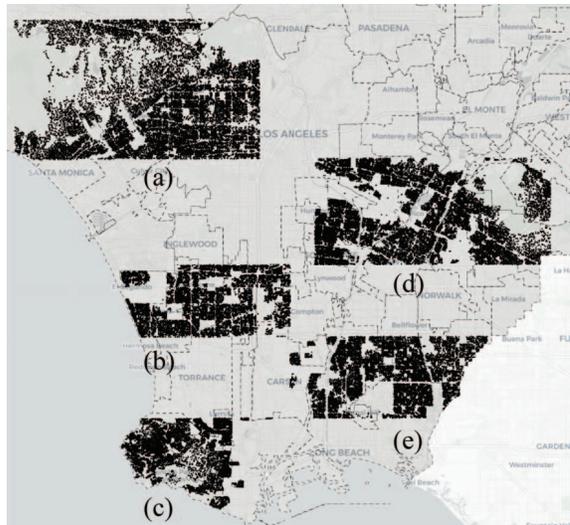


Figure 4. **Dataset Visualization.** Selected Neighborhood in Los Angeles County. (a) Santa Monica, Westwood, Brent Wood, Beverly Hills, west Hollywood (b) Manhattan Beach, Hawthorne, El Segundo, Lawndale Redondo Beach and Gardena (c) Rancho Palos and Rolling Hills (d) Downey, south gate, Commerce, East Los Angeles, Pico Rivers, Whitter (e) Long Beach and Lakewood

tion surpassing 10,000 units. These specific areas have been graphically illustrated in Figure 4. The building footprint, alongside other building metadata including vintage, height, and building type, are derived from the Assessor Parcels Data provided by Los Angeles County [8]. In the interest of maintaining focus on residential aspects, we narrowed down our dataset to exclusively incorporate buildings designated for residential usage (single and multi-family buildings). Building heights were normalized into the range of 0-255 and translated into greyscale values. This color-coded scheme allows for intuitive visualization of height distribution within residential buildings. Representative samples of these building footprints are showcased in Figure 5.

Region (b), (c), and (e) each consist of 6959, 7736, 6064 footprints. Among them, 90% of randomly selected footprints are utilized for the training process of our model, while the remaining portion is allocated for validation. Regions (a) and (d) are held out and encompass 9870, 11767 footprints each, intended for conducting open-set experiments.

### 4.2. Implementation details and Metrics

A single footprint with dimensions of 1410 × 1410 is initially divided into three sizes: 700 × 700, 224 × 224, and 112 × 112. Subsequently, all of these sizes are uniformly resized to 112 × 112 before being combined into a concatenated format.

The image embedding network consists of a 3-layer CNN and 1 residual block serving as the encoder, along

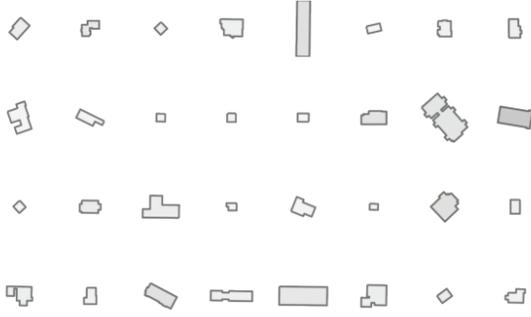


Figure 5. Examples of the building footprints.

with a vector quantizer, 3 up-sampling layers, and 1 residual block as the decoder. All meta-information encoders in the downstream task pool are constructed with only 1 convolution layer and 1 fully-connected layer. Feature embedding fed into both the vector quantizer and the downstream task pool is of dimension  $28 \times 28 \times 32$ .

Regarding the training strategy for our MARL integrated with DTP, we adopted a pre-training phase focusing solely on the reconstruction loss. Following this, we fully fine-tune the entire model’s weights for one epoch through a weighted sum of the reconstruction loss and the downstream task losses.

**Ground truth.** Given the privacy concerns, obtaining real-world data regarding building energy consumption poses a challenge. As a workaround, we used a simulated dataset, the Integrated Multisector Multiscale Modeling (IMMM) data [26], as our ground truth for energy estimation. This dataset presents a thorough simulation of energy consumption for every building across the entire Los Angeles County.

**Baseline.** For comparison, we adopt two prototype models from Prototype Building Models (PBM) set forth by [10]: single-family housing model and multi-family housing model under climate zone 3C for Los Angeles County. The PBM is a well-established building archetype dataset applicable nationwide and adaptable to various climate conditions. We conduct energy simulations for both single and multi-family models. The energy consumption estimation for the baseline is the sum of the product of energy consumption and the respective area for both single-family and multi-family housing models.

**Metrics.** To evaluate the efficacy of the proposed algorithm, we use the accuracy of the aggregated energy consumption for a selected area as the evaluation metrics (3). The absolute error  $AE$  is denoted as the difference between the total estimated energy consumption  $EC_{est}$  and the ground truth  $EC_{gt}$ . With the building stock clustered and the representative footprint selected, create a building archetype model based on these footprints, supplementing

with energy-related parameters that are identical to PBM. In this way, the only variable parameter remains the building geometry. Following this, we conduct energy simulations for each archetype model and aggregate the energy consumption results by the total area the cluster represents. The resulting aggregated energy consumption is compared with the ground truth and the energy baseline in Table 2.

$$AE = |EC_{est} - EC_{gt}| \quad (2)$$

$$Accuracy = 1 - \frac{AE}{EC_{gt}} \quad (3)$$

### 4.3. Ablation experiments

In this experimental section, we would like to answer the following questions: (1) To what extent can our MARL framework improve the accuracy of current UBEM in a fully self-supervised manner with only reconstruction tasks? (2) Whether our MARL framework, when constrained by downstream tasks with a more disciplinary scope, can further improve our accuracy of UBEM?

To answer these two questions, we evaluated the performance of our method with or without DTP in each of the three neighborhoods of Los Angeles County as shown in Figure 4: (b) Manhattan Beach, etc., (c) Rancho Palo, etc., and (e) Long Beach, etc. We input the archetype derived from our method into the energy simulations and compare the results to the one derived from the PBM [10].

As residential structures represent a significant portion of energy consumption in densely populated urban zones, our analysis concentrates on these residential areas, particularly single-family housing (SFH) and multi-family housing (MFH). We feed the data from residential buildings across various regions into our trained encoder, thereby securing the latent representation for each structure.

For each representative footprint within the clustered building stock, we construct a building archetype for energy estimation. We use the building height and footprint to reconstruct the building envelope. Other specifications, such as the window-wall ratio and material are in alignment with the PBM building archetype from zone 3C [10]. Following the archetype construction, we perform energy simulation under Los Angeles International Airport climate data as extracted from [1].

#### 4.3.1 Single archetype

The archetypes provided by PBM only contain one each for SFH and MFH [10]. So we first compute one center of the latent vectors for SFH and one for MFH respectively in Rancho Palos. Taking SFH as an example, when we obtain the center of the latent vectors corresponding to all the buildings in Rancho Palos, we input the nearest samples adjacent to that center as archetypes into the UBEM to obtain

Archetype offered by		EUI ( $kWh/m^2$ )	Building Area ( $m^2$ )	Energy ( $kWh$ )	Accuracy (%)
PBM [10]	MFH	75.14	861123.64	137344567	71.62
	SFH	60.79	1194889.74		
MARL (Ours)	MFH	89.3	861123.64	179539369	93.62
	SFH	85.9	1194889.74		
MARL + DTP (Ours)	MFH	92.5	861123.64	183609344	<b>95.74</b>
	SFH	87	1194889.74		
Energy Consumption GT [26]			2056013.38	191779982	/
Energy Estimation Accuracy Boosted by			<b>Our Reconstruction Task</b>		22.00 ↑
			<b>Our Downstream Task</b>		2.12 ↑

Table 1. **Experiment Results.** We tried our method with single archetype generation for MFH and SFH in the region Rancho Palos etc.

Region	Energy Consumption GT[26]( $kWh$ )	PBM [10] Accuracy(%)	MARL with Only Reconstruction Task (%)	MARL Restricted by DTP (%)
Rancho Palo etc.	191779982	71.62	90.36	<b>91.08</b> 18.74 ↑ + 0.72 ↑
Long Beach etc.	104117941	73.10	<b>98.96</b>	25.86 ↑ - 1.37 ↓
Manhattan Beach etc.	121545524	70.51	90.43	<b>92.88</b> 19.92 ↑ + 2.45 ↑
SUM	417443447	71.66	92.52	<b>93.23</b> 20.86 ↑ + 0.70 ↑

Table 2. **Experiment Results on Multiple Regions.** Our estimation compared with PBM using multiple archetypes.

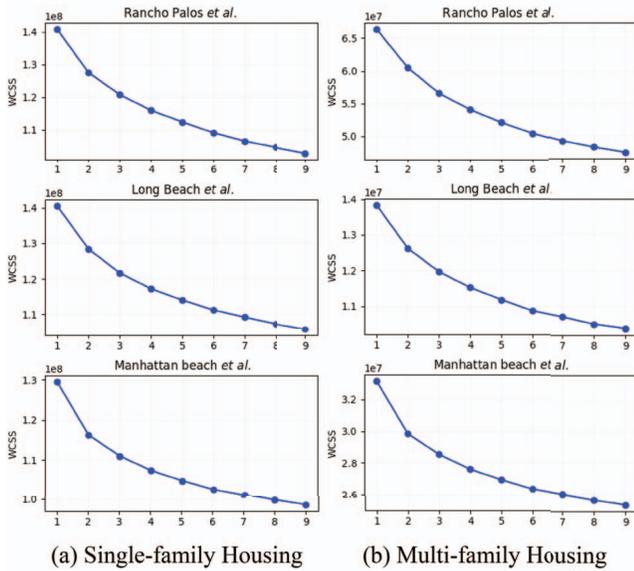


Figure 6. Within-cluster sum of squares for k-means clustering on latent space for different neighborhoods.

the average energy consumption of that region. This average energy consumption is weighted and summed over the building area to get the energy consumption of the whole region.

Table 1 shows the result of this experiment. When there is only the image reconstruction task, the UBEM accuracy from our MARL-provided archetypes is 22 percent higher

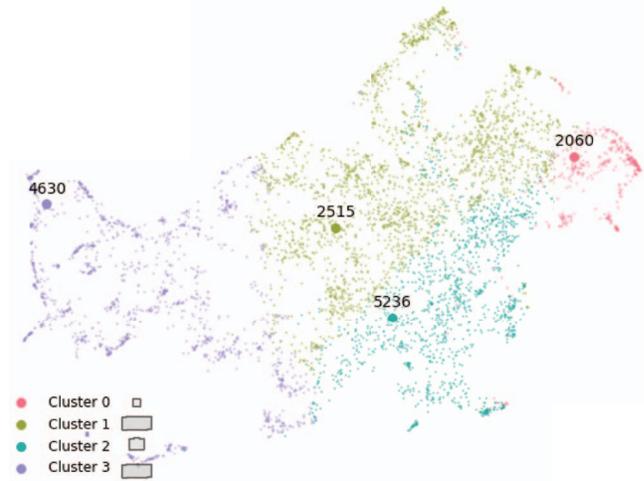


Figure 7. UMAP visualization of k-means clustering on latent space of region (c) with corresponding cluster center

than the one from PBM's, and when the downstream tasks are added during model training, our method is 24.12 percent more accurate than the traditional method.

In the training of our model, our method is not supervised by the ground truth of energy consumption, and the whole process is only supervised by the self-supervision from the building footprints and the labeled supervision from the building meta information. However, it can be seen that when we consider the geometric features and architectural attributes of local buildings in a certain region,

Region	GT[26] ( <i>kWh</i> )	PBM [10]		MARL+DTP (Ours)	
		( <i>kWh</i> )	(%)	( <i>kWh</i> )	(%)
Downey etc.	187349182	144685496	77.23	201129362	<b>92.64</b> 15.42 ↑
Santa Monica etc.	211891201	151819183	71.65	192191917	<b>90.70</b> 19.05 ↑
SUM	399240383	296504678	74.27	393321279	<b>98.52</b> 24.25 ↑

Table 3. **Open Set Experiment Results.** Energy consumption estimation in unseen regions.

our method is able to learn more valuable representations and produce more locally characterized archetypes, which leads to a more accurate simulation of building energy consumption.

#### 4.3.2 Multiple archetype

Based on our trained model, we use k-means clustering to the dimensionally reduced latent space to find representative building footprints prevalent in the selected neighborhoods. To evaluate the clustering results and determine the optimal number of clusters, the within-cluster sum of squares (WCSS) is used as an inertia measure. We select the number of clusters based on the elbow of WCSS, implying the optimal cluster number is within 2 to 5 (Figure 6). With SFH numbers reaching 5537, 5317, and 5510, significantly surpassing the MFH counts of 1422, 2419, and 554 in Regions (b), (c), and (e), respectively, we opt for 4 archetypes among SFH and 2 among MFH for this experiment.

We designate the center of each cluster as the representative building for that particular group. An example of the clustered latent space, along with the representative footprint, is visualized by Uniform Manifold Approximation and Projection (UMAP) in Figure 7.

We further tried our algorithm on three regions, Table 2 shows our results. The archetypes derived from our algorithm still perform significantly better than the conventional archetype in terms of accuracy on UBEM. And the downstream tasks’ restriction further improves the performance of our model in general.

#### 4.4. Open set

Energy modeling for all buildings in a region is costly, not all regions have energy consumption data or simulation for each building, and most regions can only use one archetype for each building category [10], which is designed at a country-wide scale. From this perspective, our approach is valuable and efficient because it can provide locale-specific building types, and does not require all building energy data as labels to oversee the entire model training process. We can train on any piece of area as long as GIS data is available.

We can even directly use our trained model for encoding building footprints in an open set. The experimental results presented in Table 3 show the performance of our

model in two unseen LA regions: (a) Santa Monica etc. and (d) Downey etc. Our experimental results demonstrate that despite not being trained on these regions at all, our model still performs remarkably well in encoding new architectural footprints, leading to excellent representations and archetypes that enable it to succeed in the UBEM tasks.

## 5. Conclusion

In this research, we introduce Multi-scale Archetype Representation Learning (MARL) for the automated creation of building archetypes adaptable to diverse regions. This methodology incorporates building geometry into the building archetype construction process. Utilizing representation learning and downstream task restriction, we extract implicit features from building stock and then employ k-means clustering to identify representative building footprints as archetypes. We further refine our approach by constraining the footprint reconstruction process with energy performance-related metadata such as vintage year and building use type.

As a result, we propose an automated building archetype construction method. We validate the efficacy of our model by benchmarking its energy estimation performance against conventional building archetypes, both for seen and unseen neighborhoods. Our experimental outcomes demonstrate that our approach outperforms conventional models, accentuating the potential of our method in enhancing the precision and adaptability of Urban Building Energy Modeling, and presenting an advancement in the domain of building archetype development and building energy modeling.

By integrating building geometries into the archetype construction process, architects and urban planners can make informed choices about neighborhood configurations, building orientations, and other design elements impacting energy consumption. Our method can also identify areas with suboptimal energy performance, indicating opportunities for retrofitting and design modifications. Furthermore, by focusing on locale-specific geometric nuances, we offer a nuanced approach to building archetypes, fostering designs that harmonize with local characteristics. This promotes energy efficiency and locale-specific design, and provides insights into the wider implications of building morphology on energy and environmental factors.

## References

- [1] Repository of free climate data for building performance simulation. <https://climate.onebuilding.org/>, Accessed July 26, 2023.
- [2] Repository of free climate data for building performance simulation [climate.onebuilding.org](https://climate.onebuilding.org/).
- [3] Usman Ali, Mohammad Haris Shamsi, Cathal Hoare, Eleni Mangina, and James O'Donnell. Review of urban building energy modeling (ubem) approaches, methods and tools using qualitative and quantitative analysis. *Energy and Buildings*, 246:111073, September 2021.
- [4] Carlos Cerezo Davila, Christoph F. Reinhart, and Jamie L. Bemis. Modeling boston: A workflow for the efficient generation and maintenance of urban building energy models from existing geospatial datasets. 117:237–250.
- [5] Carlos Cerezo Davila, Christoph F. Reinhart, and Jamie L. Bemis. Modeling boston: A workflow for the efficient generation and maintenance of urban building energy models from existing geospatial datasets. *Energy*, 117:237–250, 2016.
- [6] Yixing Chen, Tianzhen Hong, Xuan Luo, and Barry Hooper. Development of city buildings dataset for urban building energy modeling. 183:252–265.
- [7] Brian Coffey, Andrew Stone, Paul Ruyssevelt, and Philip Haves. An epidemiological approach to simulation-based analysis of large building stocks. December 2015.
- [8] California County of Los Angeles. Countywide building outlines, 2014. <https://data.lacounty.gov/datasets/lacounty::countywide-building-outlines-2014/about?layer=1>.
- [9] Isabel De Jaeger, Glenn Reynders, Caroline Callebaut, and Dirk Saelens. A building clustering approach for urban energy simulations. *Energy and Buildings*, 208:109671, 2020.
- [10] Department of Energy. Prototype building models, 2021. <https://www.energycodes.gov/prototype-building-models>, Accessed May 25, 2023.
- [11] Michael Deru, Katherine Field, Daniel Studer, Kyle Benne, Brent Griffith, Paul Torcellini, Bing Liu, Mark Halverson, David Winiarski, Michael Rosenberg, Mehry Yazdanian, Jingxin Huang, and Drury Crawley. U.s. department of energy commercial reference building models of the national building stock (nrel/tp-5500-46861). Technical report, National Renewable Energy Lab. (NREL), Golden, CO (United States), 2011.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding.
- [13] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. type: article.
- [14] Geoffrey Hinton, Li Deng, Dong Yu, George E. Dahl, Abdelrahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N. Sainath, and Brian Kingsbury. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. 29(6):82–97. Conference Name: IEEE Signal Processing Magazine.
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models.
- [16] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. volume 25. Curran Associates, Inc.
- [18] Tobias Loga, Britta Stein, and Nikolaus Diefenbach. Tabula building typologies in 20 european countries—making energy-related features of residential building stocks comparable. *Energy and Buildings*, 132:4–12, 2016.
- [19] Adam Nutkiewicz, Booyoung Choi, and Ravi K. Jain. Exploring the influence of urban context on building energy retrofit performance: A hybrid simulation and data-driven approach. *Advances in Applied Energy*, 3:100038, 2021.
- [20] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. *ArXiv:1901.05103*, 2019.
- [21] United Nations Environment Programme. 2022 global status report for buildings and construction: Towards a zero-emission, efficient and resilient buildings and construction sector. Technical report, November 2022.
- [22] Janet Reyna, Eric Wilson, Andrew Parker, Aven Satre-Meloy, Amy Egerter, Carlo Bianchi, Marlena Praprost, Andrew Speake, Lu Liu, Racheal Horsey, Matthew Dahlhausen, Carli Caradonna, and Stacy Rothgeb. U.s. building stock characterization study: A national typology for decarbonizing u.s. buildings. *Renewable Energy*, 2022.
- [23] Janet L. Reyna, Mikhail V. Chester, and Sergio J. Rey. Defining geographical boundaries with social and technical variables to improve urban energy assessments. *Energy*, 112:742–754, 2016.
- [24] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows.
- [25] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. 2018.
- [26] Y. Xu, P. Vahmani, A. Jones, and T. Hong. A multi-scale time-series dataset of anthropogenic heat from buildings in los angeles county (version v1) [data set]. MSD-LIVE Data Repository, 2022. <https://doi.org/10.57931/1892041>.
- [27] Yibo Yang, Stephan Mandt, and Lucas Theis. An introduction to neural data compression.