Additive Compositionality in Urban Area Embeddings Based on Human Mobility Patterns

Naoki Tamura
tam@ucl.nuee.nagoya-u.ac.jp
Nagoya University
Japan

Haru Terashima
haru@ucl.nuee.nagoya-u.ac.jp
Nagoya University
Japan

Kazuyuki Shoji
shoji@ucl.nuee.nagoya-u.ac.jp
Nagoya University
Japan

Shin Katayama
shinsan@ucl.nuee.nagoya-u.ac.jp
Nagoya University
Japan

Kenta Urano
vrano@ucl.nuee.nagoya-u.ac.jp
Nagoya University
Japan

Takuro Yonezawa
takuro@nagoya-u.jp
Nagoya University
Japan

Nobuo Kawaguchi
kawaguti@nagoya-u.jp
Nagoya University
Japan

ABSTRACT

Understanding the characteristics of various urban areas is crucial for applications such as urban planning, tourism policies, market analysis, and infection control. Techniques for embedding areas as vectors in a latent space based on human mobility patterns are actively researched. Many of these area embedding methods define areas as points, grids, or polygons on a geospatial plane and then embed them. However, existing methods do not allow for mutual transformation between these forms and sizes after the initial embedding. Additionally, if the characteristics of an area change due to events such as the opening of new buildings, re-embedding is necessary. Meanwhile, the Word2Vec technique, a representative word embedding method, has a property called additive compositionality. This property allows for the arithmetic operation of word meanings through the arithmetic operations of word embeddings. In this paper, we propose a method to apply this property to existing area embedding techniques, leveraging it for practical tasks such as area shape transformation and searching for areas with trends change.

CCS CONCEPTS

• **Computing methodologies → Modeling methodologies**.

KEYWORDS

Urban Data, Representation Learning, Human Mobility

SIGSPATIAL '24, October 29 - November 1, 2024, Atlanta, GA, USA

ACM ISBN 979-8-4007-1107-7/24/10
https://doi.org/https://doi.org/10.1145/3678717.3691279

ACM Reference Format:
Naoki Tamura, Haru Terashima, Kazuyuki Shoji, Shin Katayama, Kenta Urano, Takuro Yonezawa, and Nobuo Kawaguchi. 2024. Additive Compositionality in Urban Area Embeddings Based on Human Mobility Patterns. In *Proceedings of The 32nd ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '24)*. ACM, New York, NY, USA, 4 pages. https://doi.org/https://doi.org/10.1145/3678717.3691279

1 INTRODUCTION

Understanding the functions of areas in a city is crucial for a wide range of applications, including urban planning[7], market analysis[5], and tourism policies[12]. Area embedding methods extract features related to each area and embed them into low-dimensional vectors in a latent space[4, 6, 9–11, 13, 14]. This allows for quantitative evaluation of urban functions and tasks such as area search and classification and mobility prediction, utilizing area information in downstream models.

Existing area embedding methods first determine the area shapes and time periods, then embed the data associated with each area for specific periods into a latent space. The spatial definitions of areas, such as points, grids, or polygons, vary in size and shape depending on the dataset and analysis target. However, it is inefficient to relearn embeddings every time the required area shapes change for different applications and datasets. If we can obtain transformed area embeddings using pre-trained embeddings without additional training, we can analyze areas in the required shapes as needed, making the process highly efficient. Additionally, even for the same area, its characteristics can change over different time periods due to events such as the opening of new buildings[3]. These changes are reflected in the area embeddings, but this aspect has not been analyzed so far. For example, searching for areas where specific types of buildings have been replaced, as shown in Figure 1, would be useful for analyzing urban dynamics.

We utilize a property called additive compositionality of area embeddings to enable post-hoc transformations of area shapes and the search for areas with changing trends. Additive compositionality in word embedding methods allow for arithmetic operations on word

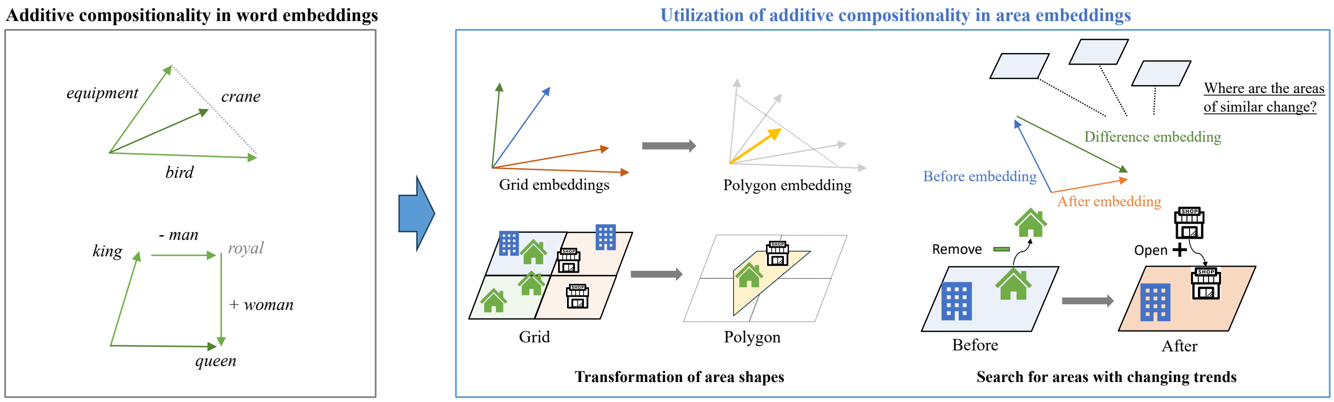

Figure 1: Utilization of additive compositionality in area embeddings.

meanings through arithmetic operations on embeddings[8]. In this paper, we verify the additive compositionality of area embeddings using Area2Vec[10], a Word2Vec-based method for embedding stay features. Furthermore, we demonstrate that the additive compositionality of area embeddings can be applied to tasks involving transforming area shapes and searching areas where trends have changed due to events.

2 METHODOLOGY

2.1 The hypothesis additive compositionality in area embedding.

First, we hypothesize additive compositionality in area embeddings. According to Naito et al., the embedding of a polysemous word w with k meanings $w_1, w_2, ..., w_k$ are approximated by the weighted average of the embedding of its meanings. This is known as OR compositionality in additive compositionality. In other words, if v_w is the embedding of the polysemous word and $v_{w_1}, v_{w_1}, ...v_{w_k}$ are the embedding of its meanings, the following relationship holds.

$$\boldsymbol{v_w} \approx \sum_{i=1}^{k} \frac{p(w_i)}{p(w)} \boldsymbol{v_{w_i}} \tag{1}$$

$p(w)$, $p(w_i)$ represent the occurrence frequencies of the polysemous word and its senses in the corpus, respectively. In other words, the embedding of the polysemous word can be approximated by the weighted average of the embedding of its senses, with weights corresponding to their occurrence frequencies.

Here, we consider a hierarchical grid to examine the additive compositionality in area embeddings. First, suppose a grid super-grid G can be divided into multiple grids $g_1, g_2, ..., g_k$. The area of the super-grid is equal to the sum of the areas of the sub-grids. In this context, the set of instance d_G for the super-grid is

$$d_G = d_{g_1} \bigcup d_{g_2} \bigcup ... \bigcup d_{g_k} \tag{2}$$

. Here, an instance $i \in d_G$ in the super-grid belongs to one of $d_{g_1}, d_{g_2}, ..., d_{g_k}$. This is analogous to the relationship between a polysemous word and its senses, where G represents the polysemous word and $g_1, g_2, ..., g_k$ represent the senses. Thus, the following relationship holds.

$$\boldsymbol{v_G} \approx \sum_{i=1}^{k} \frac{c_{g_i}}{c_G} \boldsymbol{v_{g_i}} \tag{3}$$

c_G and c_{g_i} represent the number of instances corresponding to the super-grid and sub-grids. In other words, the embedding of the super-grid can be approximated by the weighted average of the sub-grid embeddings, weighted by the number of instances in each sub-grid. We call this the hypothesis of additive compositionality in area embeddings. This hypothesis will be tested in Section 5.

2.2 Transformation of area shapes

Assuming that Equation (3) holds, we consider a method for interchanging embeddings of different area shapes. As shown in Figure 1, we obtain the embedding of a polygon that spans multiple grids from the embeddings of these grids. Assuming that the stay density within the area is uniform, the number of instances in the intersection of the polygon and each grid is proportional to its area. Therefore, the number of instances in the intersection of a polygon A with area S_A and a grid g with area S_g is estimated as follows.

$$\widehat{c_{g \cap A}} \approx c_g \frac{S_{A \cap g}}{S_g} \tag{4}$$

c_g represents the number of instances in grid g. Therefore, using Equation (3), the embedding of the polygon A is approximated as the weighted average of the estimated number of instances in the intersections.

$$\boldsymbol{v_A} \approx \sum_{i=1}^{k} \frac{\widehat{c_{g_i \cap A}}}{\widehat{c_A}} \boldsymbol{v_{g_i}}, \quad \widehat{c_A} = \sum_{i=1}^{k} \widehat{c_{g_i \cap A}} \tag{5}$$

In this way, using Equation (3) or Equation (5), we obtain embeddings with transformed area shapes and sizes post-hoc. Similar transformations can be applied to other shapes, such as converting from points to grids, although decomposing into finer shapes is more challenging.

2.3 Search for areas with changing trends

Assuming that Equation (3) holds, we consider the embeddings before and after an event, such as the opening of a new building, or over time, when the data within a certain area changes. Let the set of instances before the event be d_{bef} and the set of instances after the event be d_{aft}. The instance groups change from d_{bef} to d_{aft}, where d_{pos} represents the instances that increased and d_{neg} represents the instances that decreased. In this case, the following relationship holds:

$$d_{aft} \bigcup d_{neg} = d_{bef} \bigcup d_{pos} \tag{6}$$

$$c_{aft} + c_{neg} = c_{bef} + c_{pos} \tag{7}$$

Here, c_{aft}, c_{bef}, c_{pos}, and c_{neg} represent the number of instances in each set. Applying Equation (3) to Equation (6), we can write it as Equation (8):

$$\frac{c_{aft}\boldsymbol{v_{aft}} + c_{neg}\boldsymbol{v_{neg}}}{c_{aft} + c_{neg}} = \frac{c_{bef}\boldsymbol{v_{bef}} + c_{pos}\boldsymbol{v_{pos}}}{c_{bef} + c_{pos}} \tag{8}$$

Furthermore, by using Equation (7) for simplification, we finally obtain the following equations:

$$c_{pos}\boldsymbol{v_{pos}} - c_{neg}\boldsymbol{v_{neg}} = c_{aft}\boldsymbol{v_{aft}} - c_{bef}\boldsymbol{v_{bef}} \tag{9}$$

$$\boldsymbol{v_{aft}} = \frac{c_{bef}\boldsymbol{v_{bef}} + c_{pos}\boldsymbol{v_{pos}} - c_{neg}\boldsymbol{v_{neg}}}{c_{bef} + c_{pos} - c_{neg}} \tag{10}$$

Equation (9) and (10) represent the relationship between v_{pos}, v_{neg} and v_{bef}, v_{aft}. Therefore, by calculating the left-hand side for d_{pos}, d_{neg} corresponding to a specific area change, and precomputing the right-hand side for all combinations of d_{bef}, d_{aft}, we can use the left-hand side as a query to search the right-hand side. This allows us to find areas where specific area changes have occurred. A noteworthy point about Equations (9) and (10) is that the difference in instance groups can be represented by the difference in embeddings. In other words, a decrease in area trends can be represented by a negative embedding.

3 EXPERIMENT

3.1 Experimental setup

The study area includes the entire city of Nagoya, Japan. The data used for the embeddings is a GPS dataset collected from an app installed on users' smartphones, provided by BlogWatcher Co.[1] with prior consent. The experiments used stay data from Nagoya City from 2023/01/01 to 2023/12/31. Additionally, to verify area changes, we obtained information on store openings and closures from a survey website [2]. This data includes the category of opened or closed stores, event dates, latitude, and longitude. It is inferred that trends in the target area changed during the periods before and after these events.

3.2 Transformation of area shapes

Here, we experiment with transforming 50m, 250m, and 1km grids to parent grids using Equation (3) and converting between districts and 50m grids using Equation (5). We embed the area shapes before and after transformation simultaneously with Area2Vec and evaluate the effectiveness by comparing the embeddings transformed by the proposed method with the actual learned embeddings. The metrics are the Mean Euclidean Distance (MED) and Mean Cosine Similarity (MCS), which calculate and average the distance between two embeddings based on [13], as well as Accuracy, which indicates the proportion of times the two embeddings belong to the same cluster. The results are shown in Table 1.

Table 1: Metrics of area shape transformation.

	MED	MCS	6 cluster Accuracy	12 cluster Accuracy
50m grid → 250m grid	1.1562	0.9427	0.8174	0.7543
250m grid → 1km grid	0.5172	0.9842	0.8886	0.8727
50m grid → district	1.1680	0.9296	0.7933	0.7370
district → 50m grid	1.8386	0.8674	0.7187	0.6422

Table 1 shows that the accuracy of the grid size transformation improves from 250m to 1km compared to 50m to 250m. The results also indicate that the 50m -> district transformation is more accurate than district -> 50m. This is likely because larger areas have more data associated with each embedding, resulting in more definitive embeddings. Additionally, according to Naito et al[8]., the embeddings of polysemous words tend to have smaller norms and cluster near the origin, as they represent averages of the embeddings of their meanings. Larger areas encompass more diverse trends, leading to more averaged embeddings, which is why *MED* of 250m -> 50m is extremely small.

3.3 Search for areas with changing trends

First, to identify areas where trends have changed, we obtained the months and grids where new constructions occurred in Nagoya City in 2023 from [2]. As a result, we identified 509 combinations of 50m grids and months where new buildings were constructed. For these 509 events, we collected stay data for the month before and the month after the event, resulting in d_{bef} and d_{aft}. We calculated the difference in the frequency of the 168-dimensional stay tokens between d_{bef} and d_{aft}, aggregating the increasing tokens as d_{pos} and the decreasing tokens as d_{neg}. We then mixed the d_{pos} and d_{neg} obtained from each of the 509 events with the 50m grid data corresponding to each month of 2023 and embedded them together. For each area where an event occurred, we designated the embedding for the month after the event as v_{aft} and the embedding for the month before the event as v_{bef} and verified Equation (10). The results are shown in Table 2.

In the experiment, we compared Equation (10) with cases where either v_{pos} or v_{neg} was excluded. The results showed that the post-event embedding was closer to the embedding obtained by considering both v_{pos} and v_{neg} rather than just one of them. This confirms the validity of Equation (10).

Finally, using Equation (9), we search for areas with specific trend changes. We first extracted the trend changes caused by two events, "taverns opened(+taverns)" and "workplaces built(+workplaces)," as search queries from the [2] dataset. Using the embeddings of these d_{pos} and d_{neg} and their combinations, we searched all 50m grids in

Table 2: Metrics of temporal additive compositionality.

	MED	MCS	6 cluster Accuracy	12 cluster Accuracy
aft = bef + pos - neg (Equation (10))	0.9852	0.6951	0.9150	0.8664
aft = bef + pos	0.9518	1.2532	0.8219	0.7146
aft = bef - neg	0.9027	2.1568	0.7357	0.6434
aft = bef	0.8597	2.1962	0.6856	0.5578

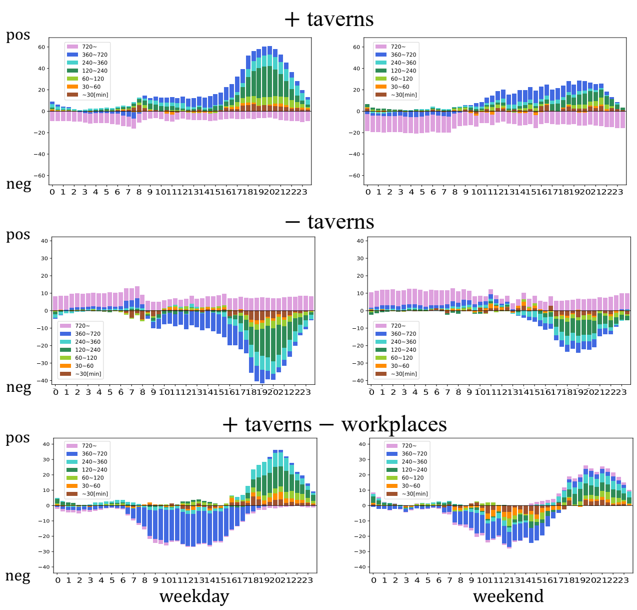

Figure 2: Search results for each query (100 areas)

Nagoya City in 2023 for grids with similar changes from one month to the next. Figure 2 shows the aggregated results of searching for 100 specific combinations of grids and months for each query.

Figure 2 shows that we can search not only for positive trend change queries (+ taverns) but also for negative trend changes (- taverns) by applying a negative factor. Additionally, by combining these queries (+ taverns - workplaces), we can identify areas where specific increases and decreases in trends occurred simultaneously. This demonstrates that by using the additive compositionality of area embeddings, we can search for both positive and negative trend changes. This demonstrates the ability to search for areas with specific trend changes from among numerous areas.

4 CONCLUSION

In this paper, we verified the additive compositionality of area embeddings and evaluated it through experiments involving two practical tasks. The results demonstrated that while it is limited to shape transformations from fine-grained to coarse-grained areas, it is feasible and can be used to search for areas with specific directional changes over time. We conducted experiments using Area2Vec, which embeds stay information with a Word2Vec-based architecture. However, as suggested by [8], additive compositionality may also hold for Attention-based embedding methods such as BERT. In future work, we aim to explore the additive compositionality in various area embedding methods, including BERT-based and AutoEncoder-based methods, and investigate the potential applications of our approach.

ACKNOWLEDGMENTS

This research was supported in part by JST CREST (JPMJCR21F2, JPMJCR22M4), NICT (222C01, 22609), NEDO (JPNP23003, JPJ012495), JSPS KAKENHI (22H03580, 22K18422) and JST ACT-X. We would like to thank Blogwatcher Inc. for providing the valuable data used in this paper.

REFERENCES

[1] [n. d.]. Blogwatcher, Inc. https://www.blogwatcher.co.jp/.
[2] [n. d.]. KaitenHeiten.com. https://kaiten-heiten.com/.
[3] Yin-Hao Chiu, I-Ting Chuang, and Chi-Yao Tsai. 2021. Analyzing correlation of urban functionality and spatial configuration. *Environment and Planning B: Urban Analytics and City Science* 48, 5 (2021), 1296–1313.
[4] Christophe Haikal, Pegah Alizadeh, Christophe Rodrigues, and Bi Chongke. 2022. Place embedding across cities in location-based social networks. In *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing (SAC '22)*. Association for Computing Machinery, 539–546.
[5] Zhigang Han, Caihui Cui, Chenghong Miao, Haiying Wang, and Xiang Chen. 2019. Identifying Spatial Patterns of Retail Stores in Road Network Structure. *Sustainability* 11, 17 (2019).
[6] Yan Lin, Huaiyu Wan, Shengnan Guo, and Youfang Lin. 2021. Pre-training Context and Time Aware Location Embeddings from Spatial-Temporal Trajectories for User Next Location Prediction. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 5 (May 2021), 4241–4248.
[7] Xiaohui Mou, Fei Cai, Xin Zhang, Jie Chen, and Rongrong Zhu. 2020. Urban Function Identification Based on POI and Taxi Trajectory Data. In *Proceedings of the 3rd International Conference on Big Data Research (ICBDR '19)*. 152–156.
[8] Masahiro Naito, Sho Yokoi, Geewook Kim, and Hidetoshi Shimodaira. 2021. Revisiting Additive Compositionality: AND, OR and NOT Operations with Word Embeddings. *CoRR* abs/2105.08585 (2021). arXiv:2105.08585
[9] Toru Shimizu, Takahiro Yabe, and Kota Tsubouchi. 2020. Enabling Finer Grained Place Embeddings using Spatial Hierarchy from Human Mobility Trajectories. In *Proceedings of the 28th International Conference on Advances in Geographic Information Systems (SIGSPATIAL '20)*. Association for Computing Machinery, 187–190.
[10] Kazuyuki Shoji, Shunsuke Aoki, Takuro Yonezawa, and Nobuo Kawaguchi. 2024. Area Modeling using Stay Information for Large-Scale Users and Analysis for Influence of COVID-19. arXiv:2401.10648 [cs.LG]
[11] Hongjian Wang and Zhenhui Li. 2017. Region Representation Learning via Mobility Flow. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM '17)*. Association for Computing Machinery, 237–246.
[12] Yang Xu, Jiaying Xue, Sangwon Park, and Yang Yue. 2021. Towards a multidimensional view of tourist mobility patterns in cities: A mobile phone data perspective. *Computers Environment and Urban Systems* 86 (01 2021).
[13] Takahiro Yabe, Kota Tsubouchi, Toru Shimizu, Yoshihide Sekimoto, and Satish V. Ukkusuri. 2019. City2City: Translating Place Representations across Cities. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '19)*. Association for Computing Machinery, 412–415.
[14] Zijun Yao, Yanjie Fu, Bin Liu, Wangsu Hu, and Hui Xiong. 2018. Representing Urban Functions through Zone Embedding with Human Mobility Patterns. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 3919–3925.