Space2Vec: Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells

Gengchen Mai

Ph.D. Candidate, STKO Lab, Department of Geography, UC Santa Barbara
http://www.geog.ucsb.edu/~gengchen_mai/

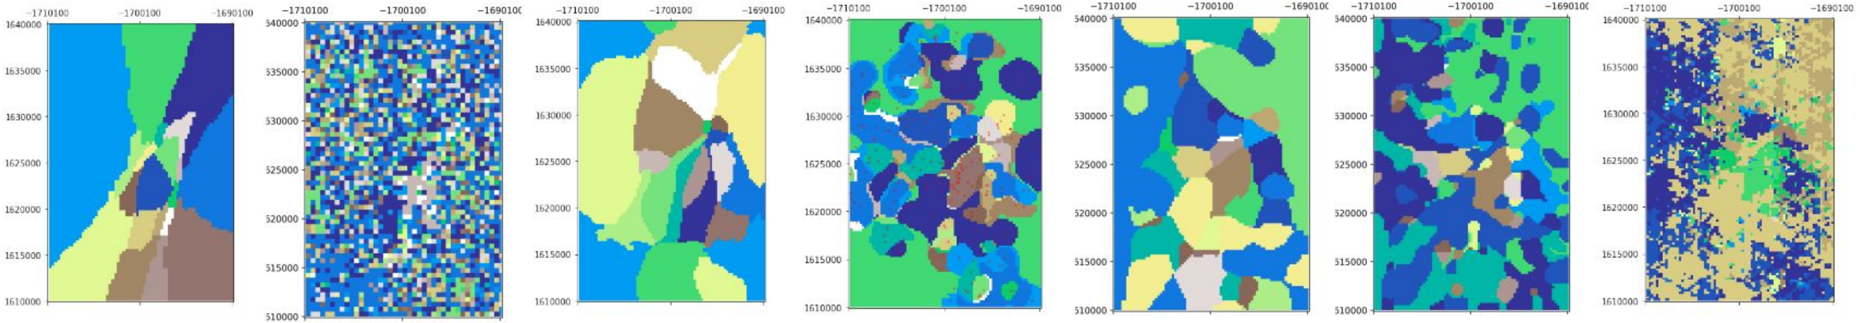

(a) *direct* (b) *tile* (c) *wrap* (d) *rbf* (σ=1k) (e) λ_{min}=1k (f) λ_{min}=500 (g) λ_{min}=50

Embedding clustering of different location encoding models:
(a)-(d) baselines (e)-(f) **Space2Vec**

ICLR 2020 paper: https://arxiv.org/abs/2003.00824
Trans. In GIS paper: https://arxiv.org/abs/2004.14171
GitHub Repo: https://github.com/gengchenmai/space2vec
GitHub Repo: https://github.com/gengchenmai/se-kge

Space2Vec: Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells

Gengchen Mai
AI Resident, Google X Core AI team;
Ph.D. Candidate, STKO Lab, Department of Geography, UC Santa Barbara
http://www.geog.ucsb.edu/~gengchen_mai/

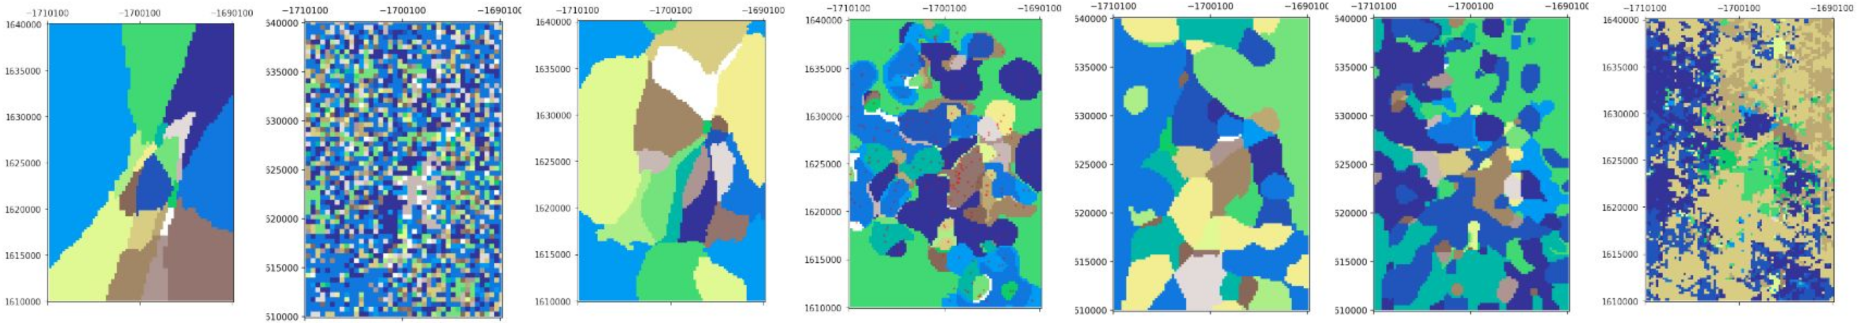

(a) *direct* (b) *tile* (c) *wrap* (d) rbf (σ=1k) (e) λ_{min}=1k (f) λ_{min}=500 (g) λ_{min}=50

Embedding clustering of different location encoding models:
(a)-(d) baselines (e)-(f) **Space2Vec**

ICLR 2020 paper: https://arxiv.org/abs/2003.00824
Trans. In GIS paper: https://arxiv.org/abs/2004.14171

GitHub Repo: https://github.com/gengchenmai/space2vec
GitHub Repo: https://github.com/gengchenmai/se-kge

Speaker Background

- **Education**
  - 5th year PhD in Geographic Information Science/GeoInformatics, UC Santa Barbara
  - BS. in GIS, Wuhan University, China
- **Work Experience**
  - Summer 2020: AI Resident at **Google X, the Moonshot Factory**, Mountainview, CA
  - Summer 2019: Cartographic Engineer at **Apple Map**, Sunnyvale, CA
  - Summer 2018: ML & NLP Research Intern at **Saymosaic Inc**, Palo Alto, CA
  - Summer 2017: ML & Software Development Intern, **Esri Inc**, Redlands, CA
- **Past Achievements**
  - 32 peer-review papers including 10 1st author papers
  - Best Paper Award at AGILE 2019, Best Paper Award at ACM KCAP 2019 (co-author); Spotlight paper at ICLR 2020; Top 10% Most Downloaded Paper at TGIS
- **Research Interests**
  - Spatially Explicit Machine Learning; Geographic Knowledge Graph; GeoAI

Outline

- **Space2Vec** (ICLR 2020 spotlight)
  - A representation learning model called Space2Vec to encode the absolute positions and spatial relationships of places inspired by biological grid cells.
  - **Tasks:** POI Classification; Geo-Aware Fine-Grained Image Classification

- **SE-KGE** (Transactions in GIS 2020)
  - A location-aware knowledge graph embedding model based on Space2Vec
  - **Tasks:** geographic logic query answering; spatial semantic lifting

Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells

Gengchen Mai[1], Krzysztof Janowicz[1], Bo Yan[2], Rui Zhu[1], Ling Cai[1], Ni Lao[3]
[1]STKO Lab, UC Santa Barbara; [2] LinkedIn Corporation; [3] SayMosaic Inc.

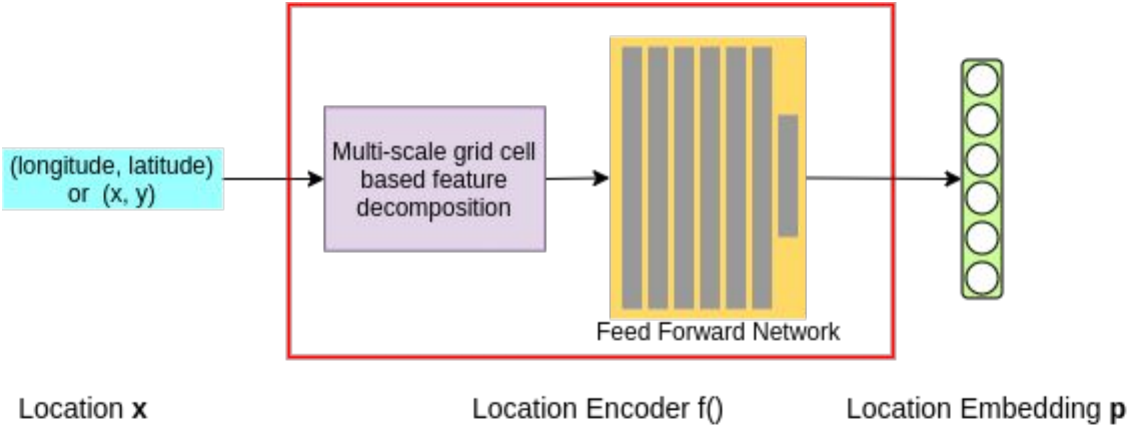

Embedding clustering of different location encoding models:
(a)-(d) baselines (e)-(f) **Space2Vec**

Arxiv paper: https://arxiv.org/abs/2003.00824

GitHub Repo: https://github.com/gengchenmai/space2vec

Unsupervised Text Encoding

Position Encoding: encode word positions with sinusoid functions of different frequencies

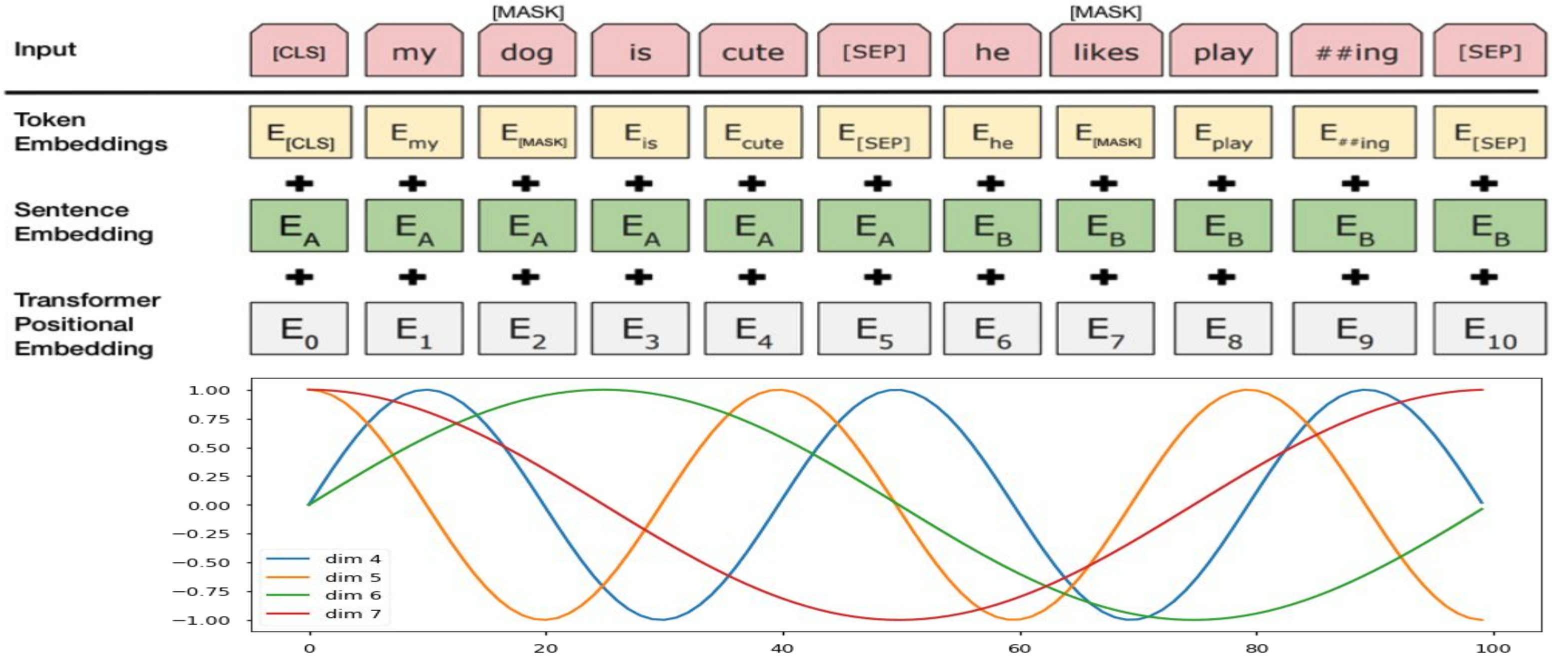

Transformer (Vaswani et al., 2017) BERT (Devlin et al., 2019)

Problem Statement

Distributed representation of point-features in space:

Given a set of points $\mathcal{P} = \{p_i\}$, i.e., Point Of Interests (POIs), in L-D space (L = 2,3), each point $p_i = (\mathbf{x}_i, \mathbf{v}_i)$ is associated with a location $\mathbf{x}_i$ and attributes $\mathbf{v}_i$ (i.e., POI feature such as type, name). We define function

$$f_{\mathcal{P},\theta}(\mathbf{x}) : \mathbb{R}^L \to \mathbb{R}^d \ (L \ll d)$$

which maps any coordinate x in space to a vector representation of d dimension

Unsupervised Location Encoding

1. Radial Basis Function (RBF)

$$K(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\sigma^2}\right)$$

- choosing the correct scale is challenging
- Need to memorize the training samples

2. Tile-based approaches (Berg at al. 2014): discretize the study area into regular grids

- choosing the correct scale is challenging
- does not scale well in terms of memory

3. Directly feed the coordinates into a FFN (inductive single-scale location encoder)

- hard to capture fine grained distributions

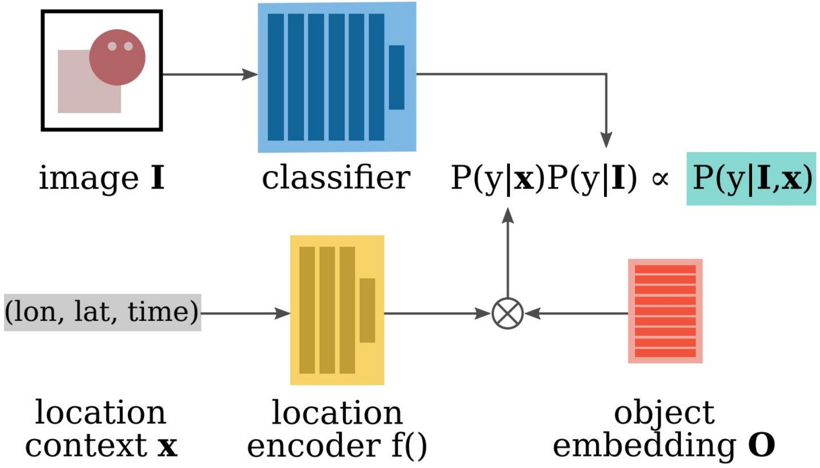

Geo-aware Image Classification (Mac Aodha et al., 2019)

Key challenge for location encoding

- Joint modeling distributions with very different characteristics
- => **multi-scale location representations**

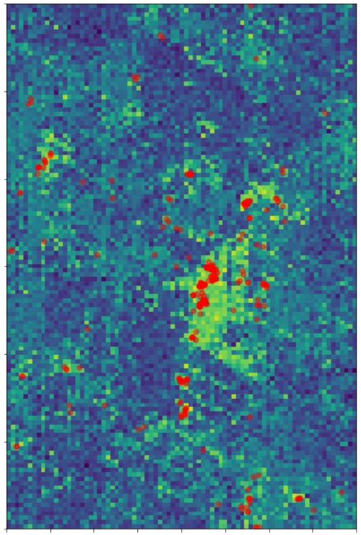

Women’s Clothing
(Clustered Distribution)

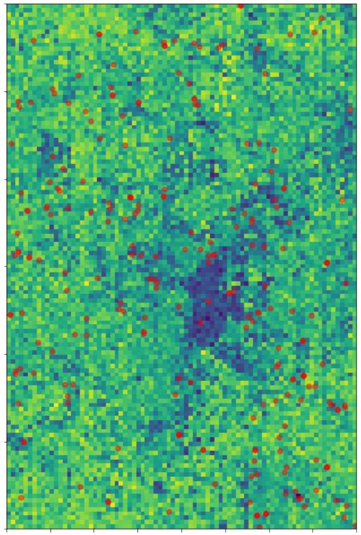

Education
(Even Distribution)

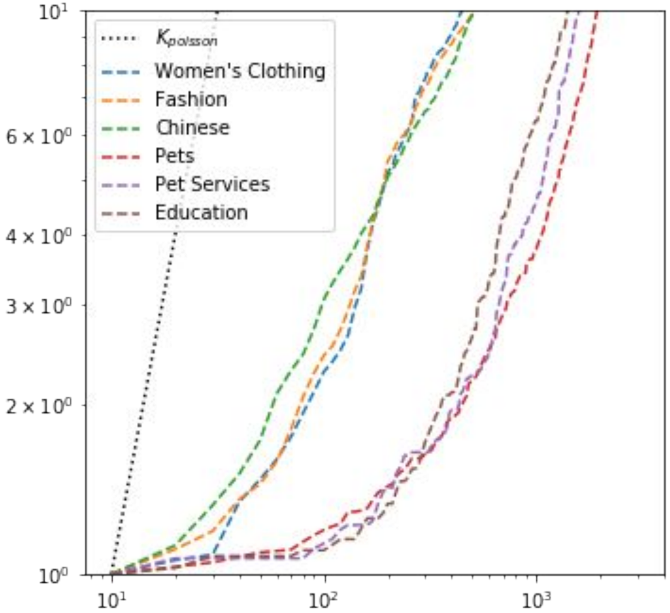

Renormalized Ripley’s K
for different POI types

Grid Cell Based Multi-Scale Location Encoding

By recording the **activities of rat neurons** during navigation within a square region, some neurons in its entorhinal cortex have a **hexagonal firing pattern**. (Hafting et al. 2005; Hafting et al. 2005; Fyhn et al. 2008; Yartsev et al. 2011; Killian et al. 2012)

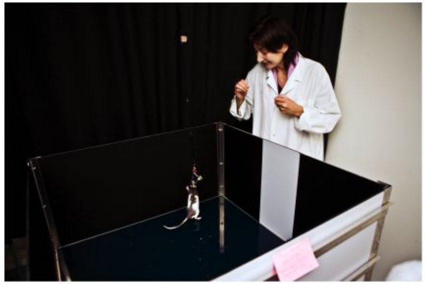

(a)

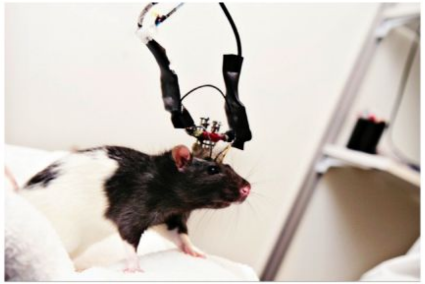

(b)

Grid cells in mammals provide a **multi-scale periodic representation** that functions as a metric for location encoding. (Banino et al., 2018)

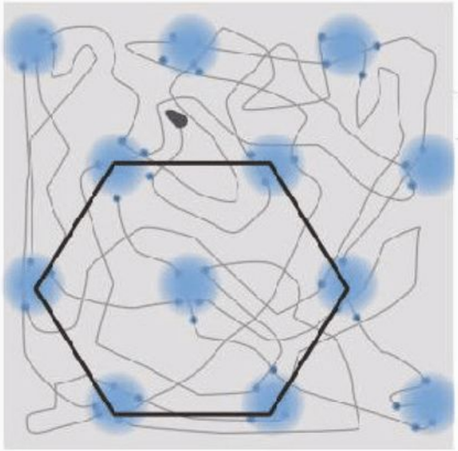

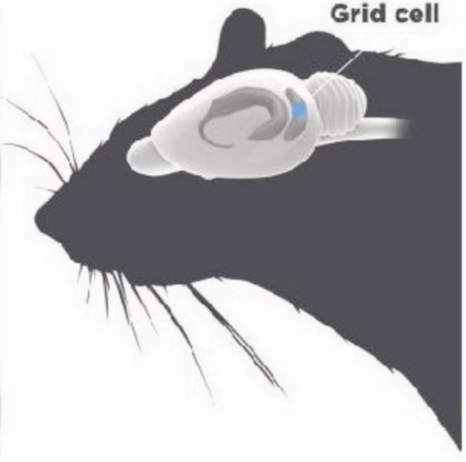

Grid Cell Based Multi-Scale Location Encoding

Grid cell representation can be simulated by summing **three cosine grating functions** oriented 60 degree apart (**a simple Fourier model of the hexagonal lattice**). (Blair et al. 2007)

It is more likely that a 2D location **x** is represented by **a population of neurons**, i.e., grid cells, so that these grid cells form **a vector representation** of this location **x.** (Gao et al. 2019)

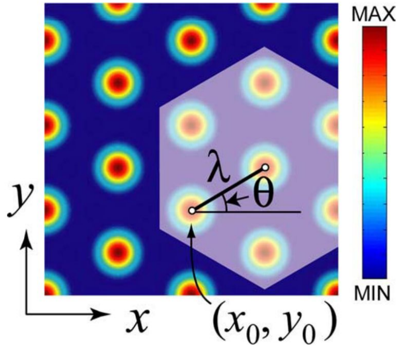

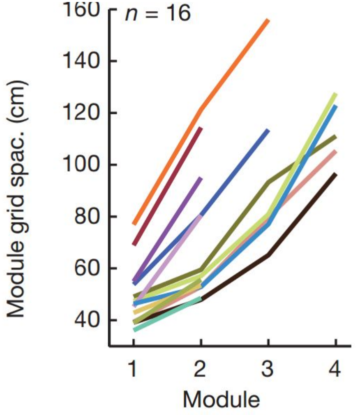

Mean grid spacing for all modules (M1–M4) in all animals (colour-coded)

Stensola et al. (2012)
Gao et al. (2019)

Grid Cell Based Multi-Scale Location Encoding

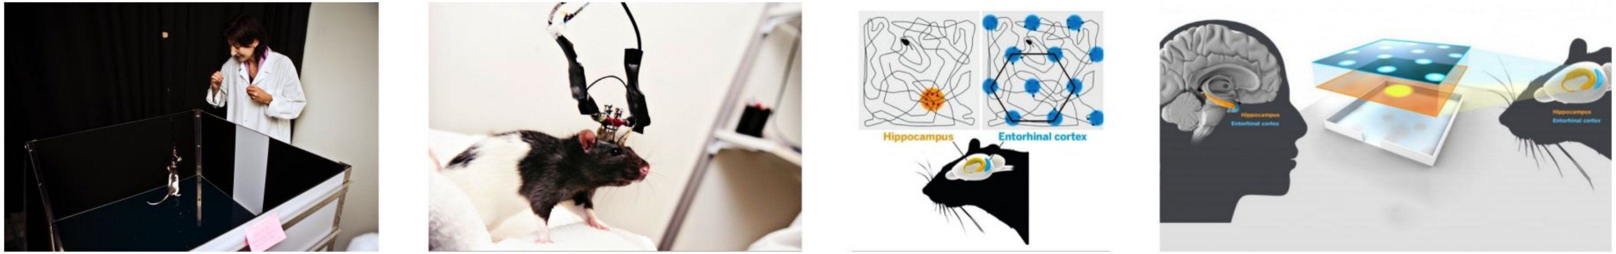

(a) (b) (c) (d)

- **Grid cells** in mammals provide a **multi-scale periodic representation** that functions as a metric for location encoding.

- It can be simulated by summing **three cosine grating functions** oriented 60 degree apart (**a simple Fourier model of the hexagonal lattice**).

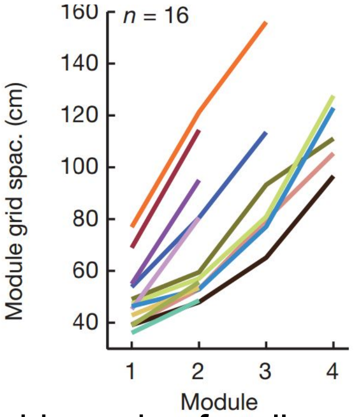

Mean grid spacing for all modules (M1–M4) in all animals (colour-coded)

Point Space Encoder: Space2Vec

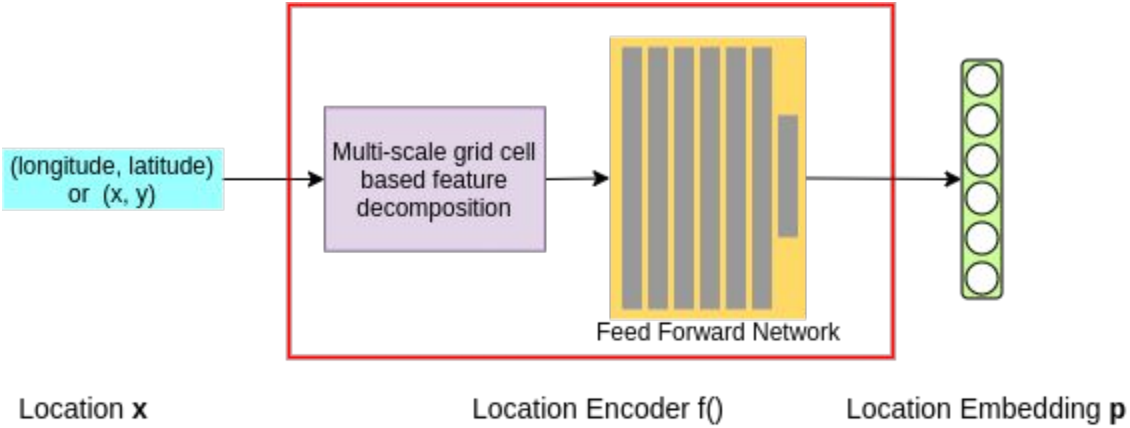

Given a location x:

$$\mathbf{e}[\mathbf{x}] \;=\; Enc^{(x)}_{theory}(\mathbf{x}) = \;\mathbf{NN}(PE^{(t)}(\mathbf{x}))$$

$$PE^{(t)}(\mathbf{x}) \;=\; [PE^{(t)}_0(\mathbf{x}); ...; PE^{(t)}_s(\mathbf{x}); ...; PE^{(t)}_{S-1}(\mathbf{x})]$$

$$PE^{(t)}_s(\mathbf{x}) = [PE^{(t)}_{s,1}(\mathbf{x}); PE^{(t)}_{s,2}(\mathbf{x}); PE^{(t)}_{s,3}(\mathbf{x})]$$

$$PE^{(t)}_{s,j}(\mathbf{x}) = [\cos(\frac{\langle \mathbf{x}, \mathbf{a}_j \rangle}{\lambda_{min} \cdot g^{s/(S-1)}}); \sin(\frac{\langle \mathbf{x}, \mathbf{a}_j \rangle}{\lambda_{min} \cdot g^{s/(S-1)}})] \forall j = 1, 2, 3;$$

Point Feature Encoder

Point feature encoder $Enc^{(v)}()$ encodes such features $\mathbf{v}_i$ nto a feature embedding $\mathbf{e}[\mathbf{v}_i] \in \mathbb{R}^{d^{(v)}}$

$\mathbf{e}[\mathbf{v}_i]$

For example, if each point represents a POI with multiple POI types, the feature embedding can simply be the mean of each POI types' embeddings:

$$\mathbf{e}[\mathbf{v}_i] = \frac{1}{H} \sum_{h=1}^{H} \mathbf{t}_h^{(\gamma)}$$

$\mathbf{t}_h^{(\gamma)}$ indicates the hth POI type embedding of a POI pi with H POI types

POI classification - Location Modeling

Location Decoder $Dec_s()$: Directly reconstructs point feature embedding $\mathbf{e}[\mathbf{v}_i]$ given its space embedding $\mathbf{e}[\mathbf{x}_i]$

$$\mathbf{e}[\mathbf{v}_i]' = Dec_s(\mathbf{x}_i; \theta_{\mathrm{dec}_s}) = \mathbf{NN}_{\mathrm{dec}}(\mathbf{e}[\mathbf{x}_i])$$

For training we use inner product to compare the reconstructed feature embedding $\mathbf{e}[\mathbf{v}_i]'$ against the real feature embeddings $\mathbf{e}[\mathbf{v}_i]$ and other negative points

POI classification - Spatial Context Modeling

Spatial Context Decoder $Dec_c()$: reconstructs the feature embedding $\mathbf{e}[\mathbf{v}_i]$ of the center point p_i based on the space and feature embeddings $\{\mathbf{e}_{i1}, ..., \mathbf{e}_{ij}, ..., \mathbf{e}_{in}\}$ of n nearby points $\{p_{i1}, ..., p_{ij}, ..., p_{in}\}$

Space-Aware Graph Attention Network Model:

$$\mathbf{e}[\mathbf{v}_i]' = Dec_c(\mathbf{x}_i, \{\mathbf{e}_{i1}, ..., \mathbf{e}_{ij}, ..., \mathbf{e}_{in}\}; \theta_{\text{dec}_c}) = g(\frac{1}{K}\sum_{k=1}^{K}\sum_{j=1}^{n} \alpha_{ijk}\mathbf{e}[\mathbf{v}_{ij}])$$

$$\alpha_{ijk} = \frac{exp(\sigma_{ijk})}{\sum_{o=1}^{n} exp(\sigma_{iok})}$$

$$\sigma_{ijk} = LeakyReLU(\mathbf{a}_k^T[\mathbf{e}[\mathbf{v}_i]_{init}; \mathbf{e}[\mathbf{v}_{ij}]; \boxed{\mathbf{e}[\mathbf{x}_i - \mathbf{x}_{ij}]}])$$

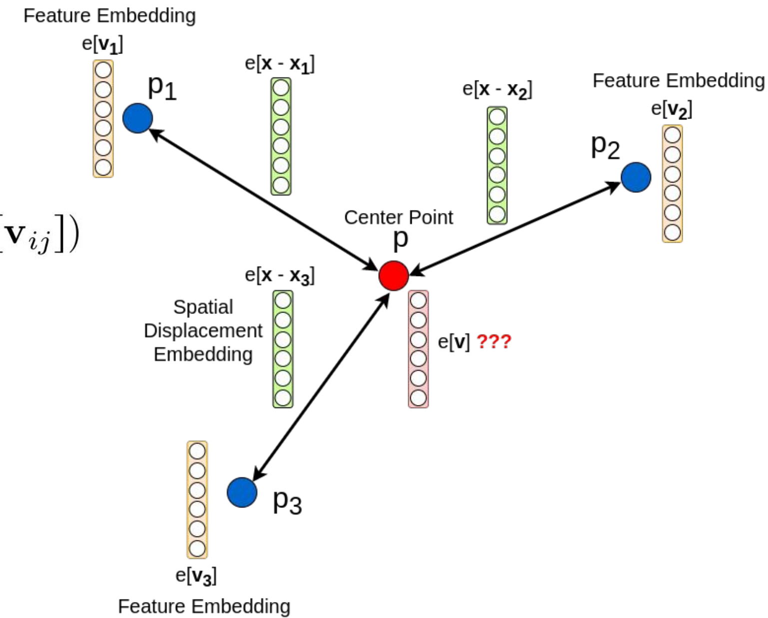

Unsupervised Training

The unsupervised learning task can simply be maximizing the log likelihood of observing the true point p_i at position $\mathbf{x}_i$ among all the points in P

$$\mathcal{L}_{\mathcal{P}}(\theta) = -\sum_{p_i \in \mathcal{P}} \log P(p_i | p_{i1}, ..., p_{ij}, ..., p_{in}) = -\sum_{p_i \in \mathcal{P}} \log \frac{\exp(\mathbf{e}[\mathbf{v}_i]^T \mathbf{e}[\mathbf{v}_i]')}{\sum_{p_o \in \mathcal{P}} \exp(\mathbf{e}[\mathbf{v}_o]^T \mathbf{e}[\mathbf{v}_i]')}$$

Negative Sampling:

$$\mathcal{L}'_{\mathcal{P}}(\theta) = -\sum_{p_i \in \mathcal{P}} \left(\log \sigma(\mathbf{e}[\mathbf{v}_i]^T \mathbf{e}[\mathbf{v}_i]') + \frac{1}{|\mathcal{N}_i|} \sum_{p_o \in \mathcal{N}_i} \log \sigma(-\mathbf{e}[\mathbf{v}_o]^T \mathbf{e}[\mathbf{v}_i]') \right)$$

POI classification - Location Modeling Evaluation

Table 1: The evaluation results of different location models on the validation and test dataset.

	Train	Validation			Testing	
	NLL	NLL	MRR	HIT@5	MRR	HIT@5
random		-	0.052 (0.002)	4.8 (0.5)	0.051 (0.002)	5.0 (0.5)
direct	1.285	1.332	0.089 (0.001)	10.6 (0.2)	0.090 (0.001)	11.3 (0.2)
tile (c=500)	1.118	1.261	0.123 (0.001)	16.8 (0.2)	0.120 (0.001)	17.1 (0.3)
wrap(h=3,o=512)	1.222	1.288	0.112 (0.001)	14.6 (0.1)	0.119 (0.001)	15.8 (0.2)
rbf (σ=1k)	1.209	1.279	0.115 (0.001)	15.2 (0.2)	0.123 (0.001)	16.8 (0.3)
grid (λ_{min}=50)	1.156	1.258	0.128 (0.001)	18.1 (0.3)	0.139 (0.001)	**20.0** (0.2)
hexa (λ_{min}=50)	1.230	1.297	0.107 (0.001)	14.0 (0.2)	0.105 (0.001)	14.5 (0.2)
theorydiag (λ_{min}=50)	1.277	1.324	0.094 (0.001)	12.3 (0.3)	0.094 (0.002)	11.2 (0.3)
theory (λ_{min}=1k)	1.207	1.281	0.123 (0.002)	16.3 (0.5)	0.121 (0.001)	16.2 (0.1)
theory (λ_{min}=500)	1.188	1.269	0.132 (0.001)	17.6 (0.3)	0.129 (0.001)	17.7 (0.2)
theory (λ_{min}=50)	1.098	1.249	**0.137** (0.002)	**19.4** (0.1)	**0.144** (0.001)	**20.0** (0.2)

Multi-scale Analysis of Location Modeling

POI Groups	**Clustered** ($r \leqslant 100m$)	**Middle** ($100m < r < 200m$)	**Even** ($r \geqslant 200m$)
direct	0.080 (-0.047)	0.108 (-0.030)	0.084 (-0.047)
wrap	0.106 (-0.021)	0.126 (-0.012)	0.122 (-0.009)
tile	0.108 (-0.019)	0.135 (-0.003)	0.111 (-0.020)
rbf	0.112 (-0.015)	0.136 (-0.002)	0.119 (-0.012)
theory	0.127 (-)	0.138 (-)	0.131 (-)
# POI	16,016	7,443	3,915
Root Types	Restaurants; Shopping; Food; Nightlife; Automotive; Active Life; Arts & Entertainment; Financial Services	Beauty & Spas; Health & Medical; Local Services; Hotels & Travel; Professional Services; Public Services & Government	Home Services; Event Planning & Services; Pets; Education

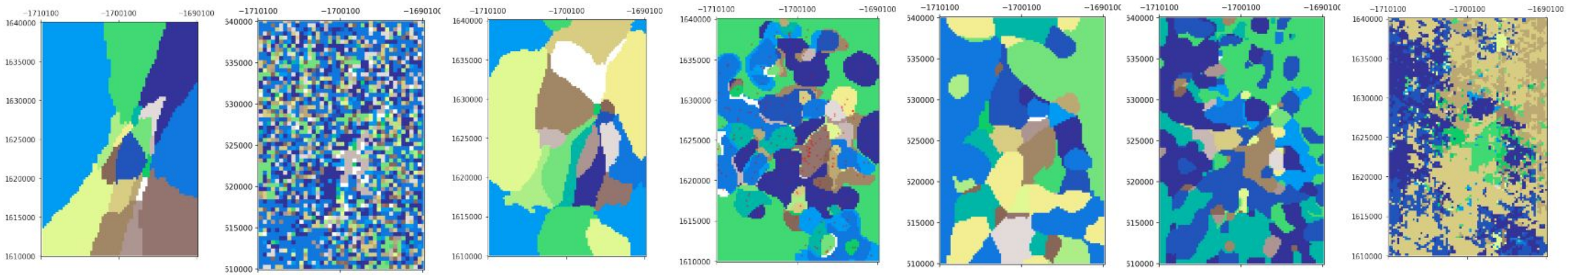

(a) *direct* (b) *tile* (c) *wrap* (d) *rbf* (σ=1k) (e) λ_{min}=1k (f) λ_{min}=500 (g) λ_{min}=50

Spatial Context Modeling Evaluation

Table 3: The evaluation results of different spatial context models on the validation and test dataset. All encoders contains a 1 hidden layer FFN. All grid cell encoders set λ_{min}=10, λ_{max}=10k.

	Train	Validation			Testing	
$Space2Vec$	NLL	NLL	MRR	HIT@5	MRR	HIT@5
$none$	1.163	1.297	0.159 (0.002)	22.4 (0.5)	0.167 (0.006)	23.4 (0.7)
$direct$	1.151	1.282	0.170 (0.002)	24.6 (0.4)	0.175 (0.003)	24.7 (0.5)
$polar$	1.157	1.283	0.176 (0.004)	25.4 (0.4)	0.178 (0.006)	24.9 (0.1)
$tile$ $(c = 50)$	1.163	1.298	0.173 (0.004)	24.0 (0.6)	0.173 (0.001)	23.4 (0.1)
$polar\_tile(S = 64)$	1.161	1.282	0.173 (0.003)	25.0 (0.1)	0.177 (0.001)	24.5 (0.3)
$wrap$ (h=2,o=512)	1.167	1.291	0.159 (0.001)	23.0 (0.1)	0.170 (0.001)	23.9 (0.2)
rbf $(\sigma = 50)$	1.160	1.281	**0.179** (0.002)	25.2 (0.6)	0.172 (0.001)	25.0 (0.1)
$scaled\_rbf$ (σ=40,β=0.1)	1.150	1.272	0.177 (0.002)	**25.7** (0.1)	0.181 (0.001)	25.3 (0.1)
$grid$(λ_{min}=10)	1.172	1.285	0.178 (0.004)	24.9 (0.5)	0.181 (0.001)	25.1 (0.3)
$hexa$ (λ_{min}=10)	1.156	1.289	0.173 (0.002)	24.0 (0.2)	0.183 (0.002)	25.3 (0.2)
$theorydiag$ $(\lambda_{min} = 10)$	1.156	1.287	0.168 (0.001)	24.1 (0.4)	0.174 (0.005)	24.9 (0.1)
$theory$(λ_{min}=200)	1.168	1.295	0.159 (0.001)	23.1 (0.2)	0.170 (0.001)	23.2 (0.2)
$theory$(λ_{min}=50)	1.157	1.275	0.171 (0.001)	24.2 (0.3)	0.173 (0.001)	24.8 (0.4)
$theory$(λ_{min}=10)	1.158	1.280	0.177 (0.003)	25.2 (0.3)	**0.185** (0.002)	**25.7** (0.3)

Geo-Aware Image Classification

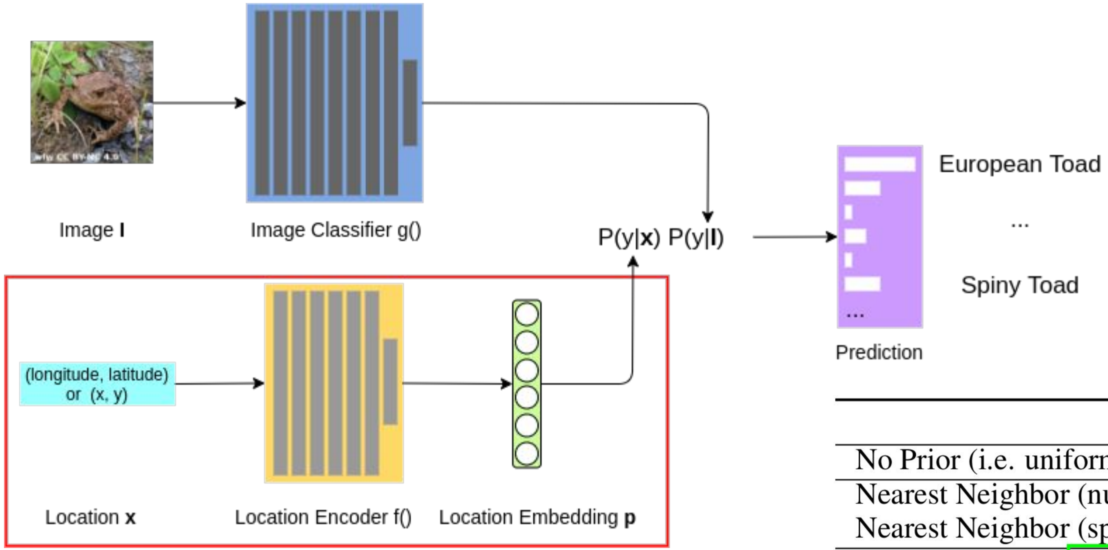

(Mac Aodha et al., 2019)

	BirdSnap†	NABirds†
No Prior (i.e. uniform)	70.07	76.08
Nearest Neighbor (num)	77.76	79.99
Nearest Neighbor (spatial)	77.98	80.79
Adaptive Kernel (Berg et al., 2014)	78.65	81.11
tile (Tang et al., 2015) (location only)	77.19	79.58
wrap (Mac Aodha et al., 2019) (location only)	78.65	81.15
rbf (σ=1k)	78.56	81.13
grid (λ_{min}=0.0001, λ_{max}=360, S = 64)	**79.44**	81.28
theory (λ_{min}=0.0001, λ_{max}=360, S = 64)	79.35	**81.59**

Arxiv paper: https://arxiv.org/abs/2003.00824

GitHub Repo: https://github.com/gengchenmai/space2vec

Conclusion for Space2Vec:

- We introduced an encoder-decoder framework as a general-purpose representation model for space inspired by **biological grid cells' multi-scale periodic representations**.
- We show the effectiveness of Space2Vec on two tasks: **POI classification** and **geo-aware image classification**.
- Our analysis reveals that it is the **ability to integrate representations of different scales** that makes the grid cell models outperform other baselines on these two tasks

SE-KGE: A Location-Aware Knowledge Graph Embedding Model for Geographic Question Answering and Spatial Semantic Lifting

Gengchen Mai[1], Krzysztof Janowicz[1], Ling Cai[1], Rui Zhu[1], Blake Regalia[1], Bo Yan[2], Meilin Shi[1], Ni Lao[3]
[1]STKO Lab, UC Santa Barbara; [2] LinkedIn Corporation; [3] SayMosaic Inc.

Spatial semantic lifting in the SE-KGE embedding space

Arxiv paper: https://arxiv.org/abs/2004.14171

SE-KGE: A Location-Aware KG Embedding Model

A novel KGE model which **directly encodes spatial footprints**, namely **point coordinates** and **bounding boxes**, thereby making them available while learning knowledge graph embeddings.

Encoding spatial footprints of geographic entities:

- **Location encoder** (Mai et al., 2020): the neural network models which encode a pair of coordinates into a high dimensional embedding which can be used in multi downstream tasks

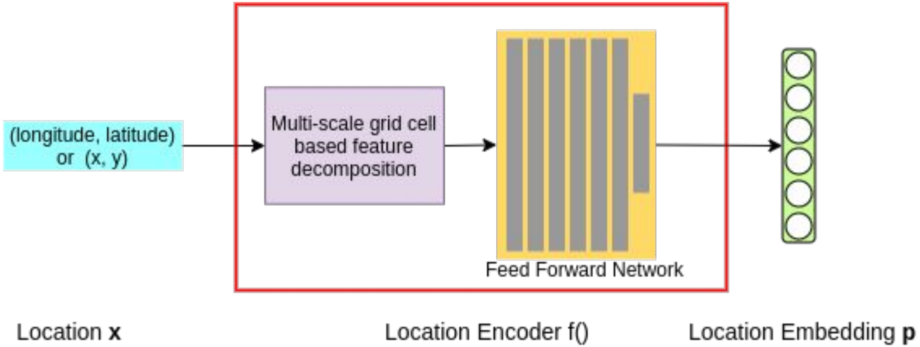

Challenges of SE-KGE

1. Location encoding can handle **point-wise metric relations** (e.g., dbo:nearestCity) and **directional relations** (e.g., dbp:north) in KGs, but it is not easy to encode containment relations (e.g., dbo:isPartOf).
   - Represent geographic entities as **regions** instead of points in the embedding space

2. How to seamlessly handle **geographic** and **non-geographic entities**?

3. How to capture the **spatial** and **other semantic aspects** at the same time?

4. **Spatial Semantic Lifting**: How to design a KGE model so that it can be used to infer new relations between entities in a KG and any arbitrary location in the study area?

Method: GeoKG Definition

Given a geographic knowledge graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

- V : the set of entities/nodes
- E : the set of directed edges
- $\mathcal{V}_{pt} \subseteq \mathcal{V}$: the geographic entity set
- $\mathcal{PT}(\cdot)$: entity $e \in \mathcal{V}_{pt}$ => $\mathcal{PT}(e) = \mathbf{x}$ where $\mathbf{x} \in \mathcal{A} \subseteq \mathbb{R}^2$
- $\mathcal{V}_{pn} \subseteq \mathcal{V}_{pt}$: the set of large-scale geographic entity
- $\mathcal{PN}(\cdot)$: entity $e \in \mathcal{V}_{pn}$ => $\mathcal{PN}(e) = [\mathbf{x}^{min}; \mathbf{x}^{max}] \in \mathbb{R}^4$ where $\mathbf{x}^{min}, \ \mathbf{x}^{max} \in \mathcal{A} \subseteq \mathbb{R}^2$

Method: CQG Definition

Definition 2 (Conjunctive Graph Query (CGQ)). *A query $q \in Q(\mathcal{G})$ that can be written as follows:*

$$
\begin{aligned}
& q = V_?.\exists V_1, V_2, .., V_m : b_1 \wedge b_2 \wedge ... \wedge b_n \\
where \quad & b_i = r_i(e_k, V_l), V_l \in \{V_?, V_1, V_2, .., V_m\}, e_k \in \mathcal{V}, r \in \mathcal{R} \\
or \quad & b_i = r_i(V_k, V_l), V_k, V_l \in \{V_?, V_1, V_2, .., V_m\}, k \neq l, r \in \mathcal{R}
\end{aligned}
$$

- $Q(\mathcal{G})$: a set of all conjunctive graph queries that can be asked over G
- $V_?$: the target variable of query q (target node)
- $V_1, V_2, .., V_m$: existentially quantified bound variables (bound nodes)
- b_i : a basic graph pattern in this CGQ
- e_k : the entity node appeared in the question (anchor node)

The dependency graph of Query q is a **directed acyclic graph** (DAG)

Geographic CGQ: the answer entity is a geographic entity

Method: CQG Example

Which city in Alameda County, California is the assembly place of Chevrolet Eagle and the nearest city to San Francisco Bay?

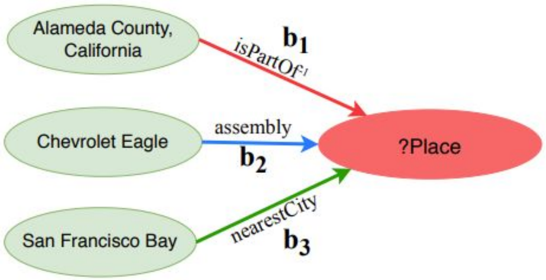

$$
\begin{aligned}
?Place : IsPartOf^{-1}(Alameda\ County, ?Place) \wedge \\
Assembly(Chevrolet\ Eagle, ?Place) \wedge \\
NearestCity(San\ Francisco\ Bay, ?Place)
\end{aligned}
$$

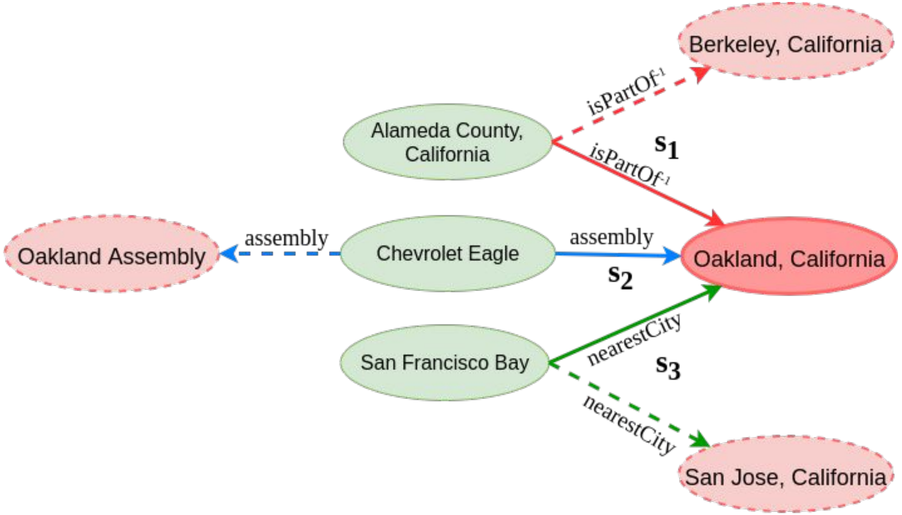

Method: Three Components for GeoQA

There major components of SE-KGE:

- **Entity encoder** $Enc()$
- **Projection operator** $\mathcal{P}()$
- **Intersection operator** $\mathcal{I}()$

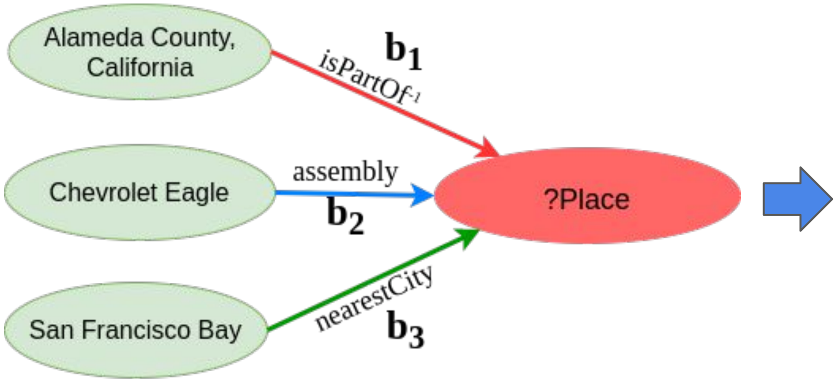

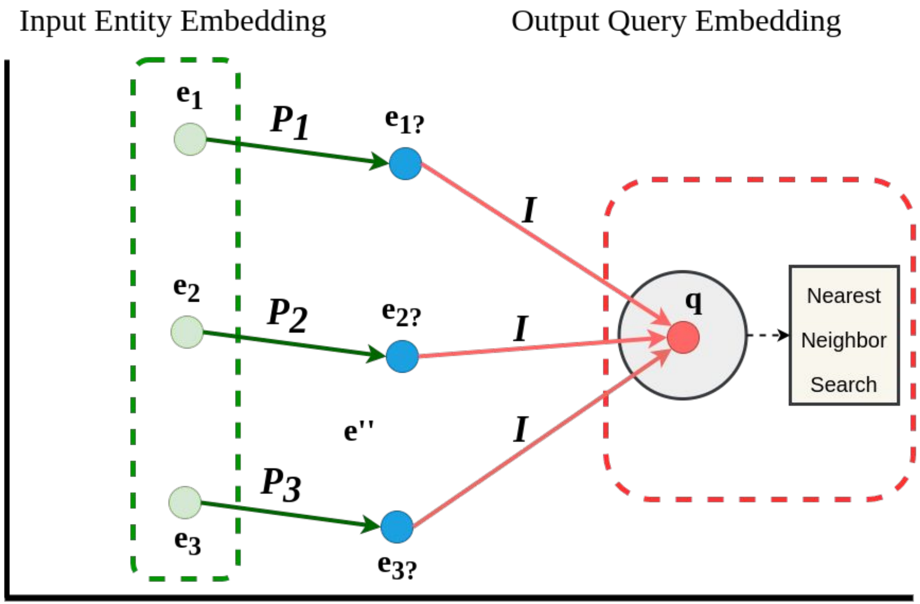

Method: Space Semantic Lifting

Use entity encoder $Enc()$ and projection operator $\mathcal{P}()$ for spatial semantic lifting:

Note that location encoder is one component of entity encoder

Method: Location-Aware Entity Encoder

- Semantic Aspect:

Definition 4 (Entity Feature Encoder: $Enc^{(c)}()$)**.** *Given any entity $e_i \in \mathcal{V}$ with type $c_i = \Gamma(e_i) \in \mathcal{C}$ from $\mathcal{G}$, entity feature encoder $Enc^{(c)}()$ computes the feature embedding $\mathbf{e}_i^{(c)} \in \mathbb{R}^{d^{(c)}}$ which captures the type information of entity e_i by using an embedding lookup approach:*

$$\mathbf{e}_i^{(c)} = Enc^{(c)}(e_i) = \frac{\mathbf{Z}_{c_i}\mathbf{h}_i^{(c)}}{\parallel \mathbf{Z}_{c_i}\mathbf{h}_i^{(c)} \parallel_{L2}} \tag{5}$$

- Space Aspect:

Definition 7 (Entity Space Encoder: $Enc^{(x)}()$)**.** *Given any entity $e_i \in \mathcal{V}$ from $\mathcal{G}$, $Enc^{(x)}()$ computes the space embedding $\mathbf{e}_i^{(x)} = Enc^{(x)}(e_i) \in \mathbb{R}^{d^{(x)}}$ by*

$$\mathbf{e}_i^{(x)} = \begin{cases} LocEnc^{(x)}(\mathbf{x}_i)\ , where\ \mathbf{x}_i = \mathcal{PT}(e_i)\ , & if\ e_i \in \mathcal{V}_{pt} \backslash \mathcal{V}_{pn} \\ LocEnc^{(x)}(\mathbf{x}_i^{(t)})\ , where\ \mathbf{x}_i^{(t)} \sim \mathcal{U}(\mathbf{x}_i^{min}, \mathbf{x}_i^{max})\ ,\ \mathcal{PN}(e_i) = [\mathbf{x}_i^{min}; \mathbf{x}_i^{max}]\ , & if\ e_i \in \mathcal{V}_{pn} \\ \dfrac{\mathbf{Z}_x\mathbf{h}_i^{(x)}}{\parallel \mathbf{Z}_x\mathbf{h}_i^{(x)} \parallel_{L2}}, & if\ e_i \in \mathcal{V} \backslash \mathcal{V}_{pt} \end{cases}$$

Method: Location-Aware Entity Encoder

- Entity Feature Encoder
- Entity Space Encoder

Encoding results are concatenated as the final output

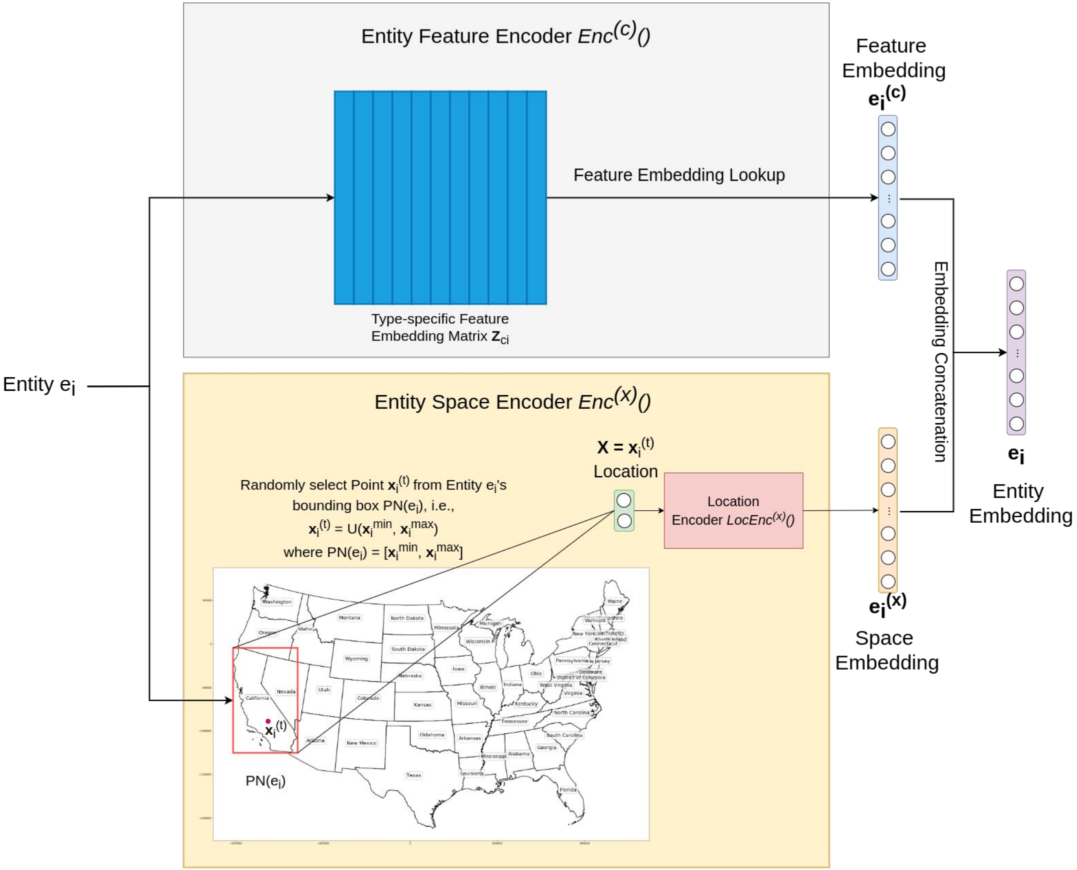

Method: Location-Aware Projection Operator

Definition 8 (Projection Operator $\mathcal{P}()$). *Given a geographic knowledge graph $\mathcal{G}$, a projection operator $\mathcal{P}() : \mathcal{V} \cup \mathcal{A} \times \mathcal{R} \rightarrow \mathbb{R}^d$ maps a pair of (e_i, r), (V_i, r), or $(\mathbf{x}_i, r)$, to an embedding $\mathbf{e}'_i$. According to the input, $\mathcal{P}()$ can be treated as:* **(1) link prediction** $\mathcal{P}^{(e)}(e_i, r)$**:** *given a triple's head entity e_i and relation r, predicting the tail; (2)* **link prediction** $\mathcal{P}^{(e)}(V_i, r)$**:** *given a basic graph pattern $b = r(V_i, V_j)$ and $\mathbf{v}_i$ which is the computed embedding for the existentially quantified bound variable V_i, predicting the embedding for Variable V_j;* **(2) spatial semantic lifting** $\mathcal{P}^{(x)}(\mathbf{x}_i, r)$**:** *given an arbitrary location $\mathbf{x}_i$ and relation r, predicting the most probable linked entity. Formally, $\mathcal{P}()$ is defined as:*

$$\mathbf{e}'_i = \begin{cases} \mathcal{P}^{(e)}(e_i, r) = diag(\mathbf{R}_r^{(c)}, \mathbf{R}_r^{(x)})Enc(e_i) = diag(\mathbf{R}_r^{(c)}, \mathbf{R}_r^{(x)})\mathbf{e}_i & if\ input\ = (e_i, r) \\ \mathcal{P}^{(e)}(V_i, r) = diag(\mathbf{R}_r^{(c)}, \mathbf{R}_r^{(x)})\mathbf{v}_i & if\ input\ = (V_i, r) \\ \mathcal{P}^{(x)}(\mathbf{x}_i, r) = diag(\mathbf{R}_r^{(xc)}, \mathbf{R}_r^{(x)})[LocEnc^{(x)}(\mathbf{x}_i); LocEnc^{(x)}(\mathbf{x}_i)] & if\ input\ = (\mathbf{x}_i, r) \end{cases}$$

Method: Location-Aware Projection Operator

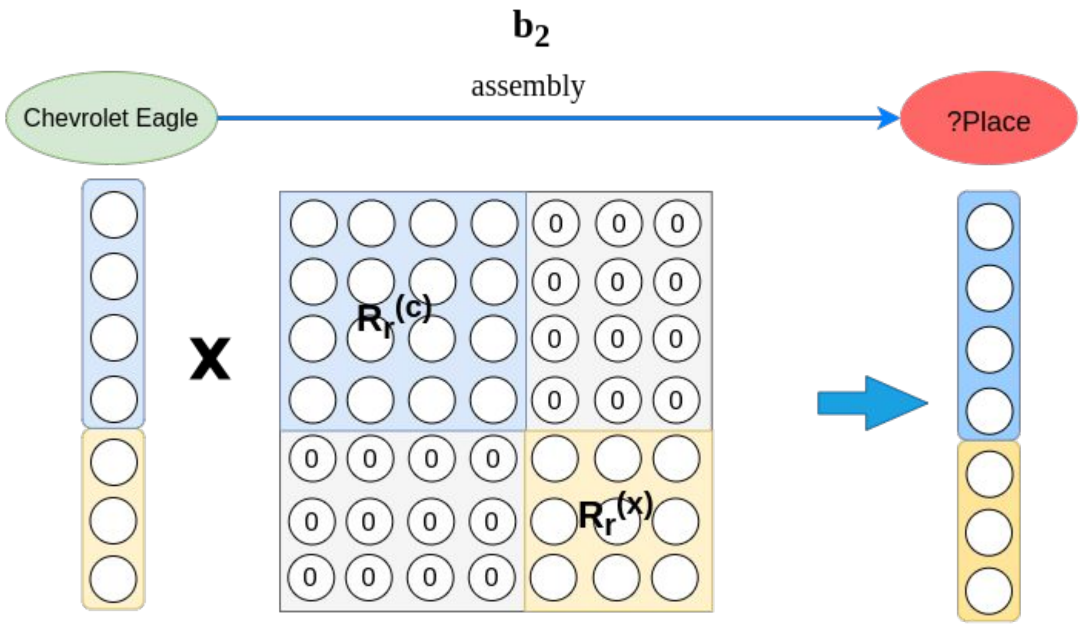

$e_i = [e_i^{(c)} \; ; \; e_i^{(x)}]$ **X** $\text{blockdiag}(R_r^{(c)}, R_r^{(x)})$ → $e_i' = [e_i^{(c)'} \; ; \; e_i^{(x)'}]$

Method: GeoQA and Spatial Semantic Lifting

- GeoQA

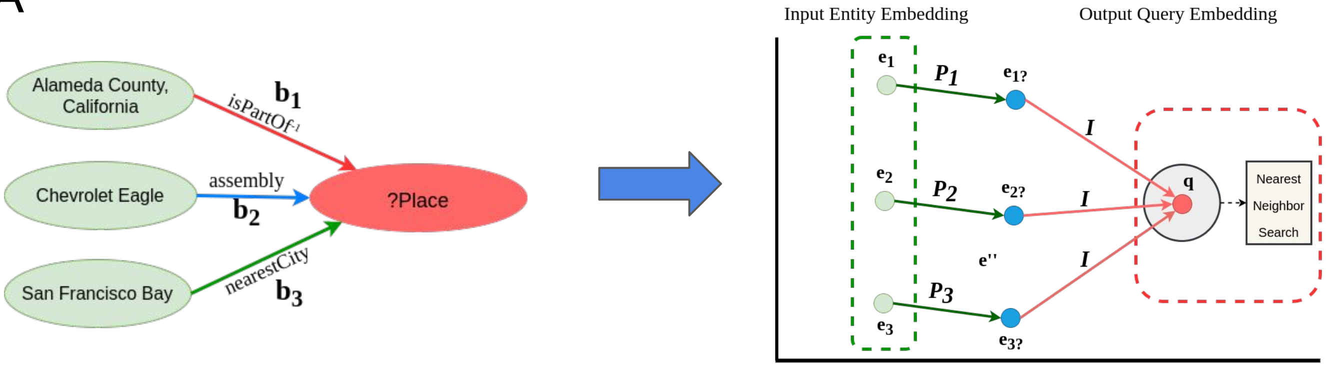

- Spatial Semantic Lifting

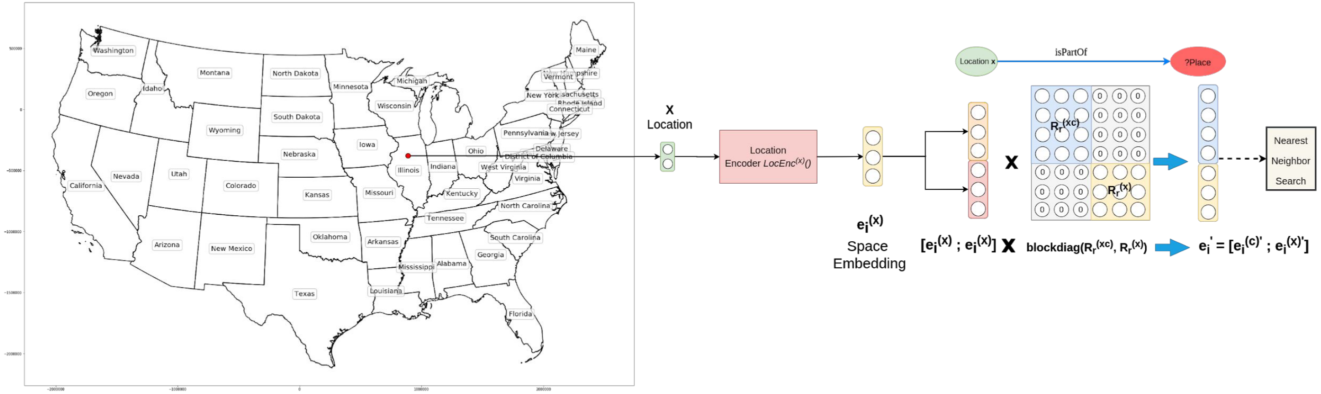

Experiment

Evaluate SE-KGE using the DBGeo dataset which is built based on a subgraph of DBpedia

Table 1: Statistics for our dataset in *DBGeo* (Section 7.1). "XXXX/QT" indicates the number of QA pairs per query type.

		DBGeo Training	Validation	Testing
Knowledge Graph	$\lvert\mathcal{T}\rvert$	214,064	2,378	21,406
	$\lvert\mathcal{R}\rvert$	318	-	-
	$\lvert\mathcal{V}\rvert$	25,980	-	-
	$\lvert\mathcal{V}_{pt}\rvert$	18,323	-	-
	$\lvert\mathcal{V}_{pn}\rvert$	14,769	-	-
Geographic Question Answering	$\lvert Q^{(2)}(\mathcal{G})\rvert$	1,000,000	-	-
	$\lvert Q^{(3)}(\mathcal{G})\rvert$	1,000,000	-	-
	$\lvert Q^{(2)}_{geo}(\mathcal{G})\rvert$	1,000,000	1000/QT	10000/QT
	$\lvert Q^{(3)}_{geo}(\mathcal{G})\rvert$	1,000,000	1000/QT	10000/QT
Spatial Semantic Lifting	$\lvert\mathcal{T}_s \cap \mathcal{T}_o\rvert$	138,193	1,884	17,152
	$\lvert\mathcal{R}_{ssl}\rvert$	227	71	135

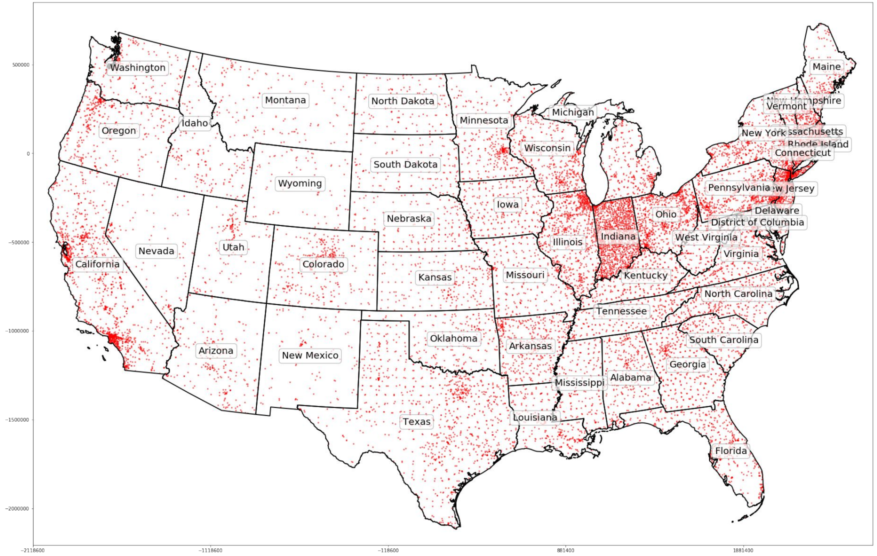

Geographic Question Answering

Table 3: The evaluation of geographic logic query answering on *DBGeo* (using AUC (%) and APR (%) as evaluation metric)

	DAG Type	GQE_{diag}		GQE		CGA		$SE\text{-}KGE_{direct}$		$SE\text{-}KGE_{pt}$		$SE\text{-}KGE_{space}$		$SE\text{-}KGE_{full}$	
		AUC	APR	AUC	APR	AUC	APR	AUC	APR	AUC	APR	AUC	APR	AUC	APR
Valid	2-chain	63.37	64.89	84.23	**88.68**	84.56	86.8	83.12	84.79	**85.97**	84.9	76.81	67.07	85.26	87.25
	2-inter	97.23	97.86	96.00	97.02	98.87	98.58	98.98	98.28	98.95	98.52	85.51	87.13	**99.04**	**98.95**
	Hard-2-inter	70.99	73.55	66.04	73.83	73.43	79.98	73.27	76.36	**74.38**	82.16	63.15	62.91	73.42	**82.52**
	3-chain	61.42	67.94	79.65	79.45	79.11	80.93	77.92	79.26	79.38	83.97	70.09	60.8	**80.9**	**85.02**
	3-inter	98.01	99.21	96.24	98.17	99.18	**99.62**	**99.28**	99.41	99.1	99.56	87.62	89	99.27	99.59
	Hard-3-inter	78.29	85	68.26	77.55	79.59	86.06	79.5	84.28	**80.48**	**87.4**	63.37	67.17	78.86	85.2
	3-inter_chain	90.56	94.08	93.39	91.52	94.59	90.71	95.99	95.11	95.86	94.41	81.16	83.01	**96.7**	**96.79**
	Hard-3-inter_chain	74.19	**83.79**	70.64	74.54	73.97	76.28	74.81	78.9	**76.45**	75.95	65.54	68.21	76.33	83.7
	3-chain_inter	**98.01**	97.45	92.69	93.31	96.72	97.61	97.31	98.67	97.79	**98.76**	83.7	84.42	97.7	98.65
	Hard-3-chain_inter	**83.59**	**88.12**	66.86	74.06	72.12	77.53	73.23	79.24	74.74	80.47	65.13	69.29	74.72	78.11
	Full Valid	81.57	85.19	81.4	84.81	85.21	87.41	85.34	87.43	**86.31**	88.61	74.21	73.9	86.22	**89.58**
Test	2-chain	64.88	65.61	85	87.41	84.91	86.74	83.61	85.97	86.08	88.08	75.46	73.38	**86.35**	**88.12**
	2-inter	96.98	97.99	95.86	97.18	98.79	98.71	98.98	98.94	**98.98**	**99.08**	87.01	85.78	98.93	99.01
	Hard-2-inter	70.39	76.19	64.5	71.86	72.15	79.26	72.04	79.11	**73.72**	**81.78**	61.22	62.97	72.62	81.04
	3-chain	62.3	62.29	79.19	80.19	78.93	80.17	77.53	78.86	79.43	**81.28**	70.55	68.04	**80.49**	80.63
	3-inter	98.09	99.12	96.54	97.94	99.33	99.56	**99.45**	99.47	99.41	**99.63**	88.05	87.63	99.39	99.59
	Hard-3-inter	77.27	83.92	68.69	75.42	78.93	83.52	78.58	84.14	**80.11**	84.87	64.44	64.53	78.76	**84.89**
	3-inter_chain	90.39	91.96	92.54	93.13	93.46	94.36	95.23	95.92	95.02	95.78	81.52	79.61	**95.92**	**96.51**
	Hard-3-inter_chain	72.89	79.12	70.67	75.55	73.47	79.61	73.93	80.21	74.88	79.36	64.99	65.52	**75.36**	**80.72**
	3-chain_inter	97.35	98.27	92.22	94.08	96.55	96.67	97.29	98.39	**97.79**	98.68	85.28	84.08	97.64	**98.75**
	Hard-3-chain_inter	**83.33**	**86.24**	66.77	72.1	72.31	77.89	73.55	77.08	75.19	77.42	65.07	65.41	74.62	77.31
	Full Test	81.39	84.07	81.2	84.49	84.88	87.65	85.02	87.81	**86.06**	88.2	74.36	73.7	86.01	**88.96**

Geographic Question Answering

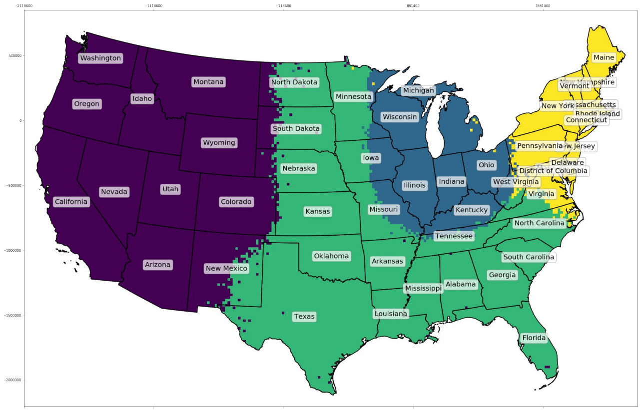

(a) Clustering result of location embeddings produced by the location encoder in SE-KGE$_{space}$

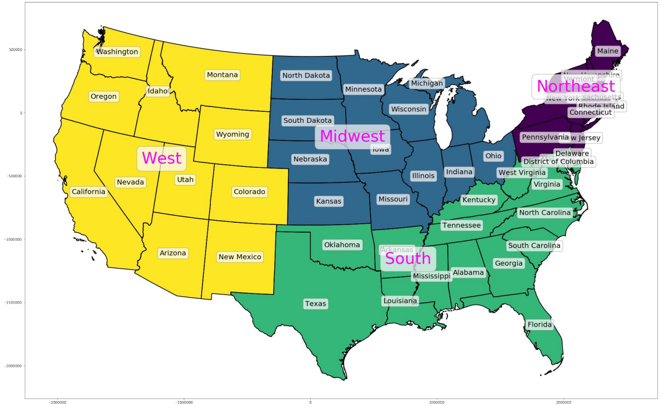

(b) Census Bureau-designated regions of United States

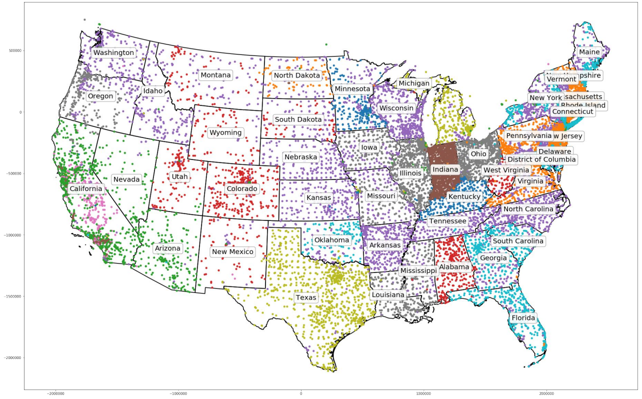

(c) The community detection (Shuffled Louvain) results of KG

Spatial Semantic Lifting

Table 5: The evaluation of spatial semantic lifting on *DBGeo* over all validation/testing triples

	SE-KGE$_{space}$		*SE-KGE$_{ssl}$*		*SE-KGE$_{ssl}$* - *SE-KGE$_{space}$*	
	AUC	APR	AUC	APR	ΔAUC	ΔAPR
Valid	72.85	75.49	**82.74**	**85.51**	9.89	10.02
Test	73.41	75.77	**83.27**	**85.36**	9.86	9.59

Table 6: The evaluation of *SE-KGE$_{ssl}$* and *SE-KGE$'_{space}$* on *DBGeo* for a few selected relation r (using APR (%) as evaluation metric).

	Query Type	*SE-KGE$'_{space}$*	*SE-KGE$_{ssl}$*	ΔAPR
Valid	$state(\mathbf{x}, ?e)$	92.00	**99.94**	7.94
	$nearestCity(\mathbf{x}, ?e)$	84.00	**94.00**	10.00
	$broadcastArea^{-1}(\mathbf{x}, ?e)$	91.60	**95.60**	4.00
	$isPartOf(\mathbf{x}, ?e)$	88.56	**98.88**	10.32
	$locationCity(\mathbf{x}, ?e)$	83.50	**99.00**	15.50
	$residence^{-1}(\mathbf{x}, ?e)$	90.50	**93.50**	3.00
	$hometown^{-1}(\mathbf{x}, ?e)$	61.14	**74.86**	13.71
Test	$state(\mathbf{x}, ?e)$	89.06	**99.97**	10.91
	$nearestCity(\mathbf{x}, ?e)$	87.60	**99.80**	12.20
	$broadcastArea^{-1}(\mathbf{x}, ?e)$	90.81	**96.63**	5.82
	$isPartOf(\mathbf{x}, ?e)$	87.66	**98.87**	11.21
	$locationCity(\mathbf{x}, ?e)$	84.80	**99.10**	14.30
	$residence^{-1}(\mathbf{x}, ?e)$	61.21	**77.68**	16.47
	$hometown^{-1}(\mathbf{x}, ?e)$	61.44	**76.83**	15.39

Conclusion for SE-KGE

- We develop a spatially-explicit knowledge graph embedding model, SE-KGE, which applies a location encoder to incorporate spatial information (coordinates and spatial extents) of geographic entities.
- SE-KGE is extended as end-to-end models for two tasks: geographic question answering and spatial semantic lifting (a new task).
- Evaluation results show that SE-KGE can outperform multiple baselines on two tasks.
- Visualization shows that SE-KGE can successfully capture the spatial proximity information as well as the semantics of relations.

Future work:

- We want to explore a more concise way to encode the spatial footprints of geographic entities in a KG

Reference

1. **Gengchen Mai**, Krzysztof Janowicz, Ling Cai, Rui Zhu, Blake Regalia, Bo Yan, Meilin Shi, Ni Lao. SE-KGE: A Location-Aware Knowledge Graph Embedding Model for Geographic Question Answering and Spatial Semantic Lifting. *Transactions in GIS*. DOI:10.1111/TGIS.12629 [arxiv paper]
2. **Gengchen Mai**, Krzysztof Janowicz, Bo Yan, Rui Zhu, Ling Cai, Ni Lao. Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells, In: *Proceedings of International Conference on Learning Representations (ICLR) 2020*, Apr. 26 - 30, 2020, Addis Ababa, ETHIOPIA . [OpenReview paper] [arxiv paper] [code] [video] [slides] *** Spotlight Paper**
3. **Gengchen Mai**, Krzysztof Janowicz, Bo Yan, Rui Zhu, Ling Cai, Ni Lao. Contextual Graph Attention for Answering Logical Queries over Incomplete Knowledge Graphs, In: *Proceedings of K-CAP 2019*, Nov. 19 - 21, 2019, Marina del Rey, CA, USA. [arxiv]

Reference

1. **Gengchen Mai**, Krzysztof Janowicz, Ling Cai, Rui Zhu, Blake Regalia, Bo Yan, Meilin Shi, Ni Lao. SE-KGE: A Location-Aware Knowledge Graph Embedding Model for Geographic Question Answering and Spatial Semantic Lifting. *Transactions in GIS*. DOI:10.1111/TGIS.12629 [arxiv paper]
2. **Gengchen Mai**, Krzysztof Janowicz, Bo Yan, Rui Zhu, Ling Cai, Ni Lao. Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells, In: *Proceedings of International Conference on Learning Representations (ICLR) 2020*, Apr. 26 - 30, 2020, Addis Ababa, ETHIOPIA . [OpenReview paper] [arxiv paper] [code] [video] [slides] *** Spotlight Paper**
3. **Gengchen Mai**, Krzysztof Janowicz, Bo Yan, Rui Zhu, Ling Cai, Ni Lao. Contextual Graph Attention for Answering Logical Queries over Incomplete Knowledge Graphs, In: *Proceedings of K-CAP 2019*, Nov. 19 - 21, 2019, Marina del Rey, CA, USA. [arxiv]
4. **Gengchen Mai**, Bo Yan, Krzysztof Janowicz, Rui Zhu. Relaxing Unanswerable Geographic Questions Using A Spatially Explicit Knowledge Graph Embedding Model, In: *Proceedings of AGILE 2019*, June 17 - 20, 2019, Limassol, Cyprus. *** 1st Best Full Paper Award**
5. Bo Yan, Krzysztof Janowicz, **Gengchen Mai**, Rui Zhu. A Spatially-Explicit Reinforcement Learning Model for Geographic Knowledge Graph Summarization. *Transactions in GIS*, 23(2019), 620-640. DOI:10.1111/tgis.12547
6. Will Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, and Jure Leskovec. Embedding logical queries on knowledge graphs. In Advances in Neural Information Processing Systems, pp. 2026-2037. 2018.
7. Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. Knowledge Graph Embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering* 29, no. 12 (2017): 2724-2743.