

Reachability Embeddings: Scalable Self-Supervised Representation Learning from Mobility Trajectories for Multimodal Geospatial Computer Vision

Swetava Ganguli*, C. V. Krishnakumar Iyer, Vipul Pandey
Apple, Cupertino, CA
{swetava, cvk, vipul}@apple.com

Abstract—Self-supervised representation learning techniques utilize large datasets without semantic annotations to learn meaningful, universal features that can be conveniently transferred to solve a wide variety of downstream supervised tasks. In this paper, we propose a self-supervised method for learning representations of geographic locations from unlabeled GPS trajectories to solve downstream geospatial computer vision tasks. Tiles resulting from a raster representation of the earth’s surface are modeled as nodes on a graph or pixels of an image. GPS trajectories are modeled as allowed Markovian paths on these nodes. A scalable and distributed algorithm is presented to compute image-like representations, called *reachability summaries*, of the spatial connectivity patterns between tiles and their neighbors implied by the observed Markovian paths. A convolutional, contractive autoencoder is trained to learn compressed representations, called *reachability embeddings*, of reachability summaries for every tile. Reachability embeddings serve as task-agnostic, feature representations of geographic locations. Using reachability embeddings as pixel representations for five different downstream geospatial tasks, cast as supervised semantic segmentation problems, we quantitatively demonstrate that reachability embeddings are semantically meaningful representations and result in 4–23% gain in performance, as measured using area under the precision-recall curve (AUPRC) metric, when compared to baseline models that use pixel representations that do not account for the spatial connectivity between tiles. Reachability embeddings transform sequential, spatiotemporal mobility data into semantically meaningful tensor representations that can be combined with other sources of imagery and are designed to facilitate multimodal learning in geospatial computer vision.

Index Terms—Self-Supervised Learning, Representation Learning, Markov Chains, Multimodal Machine Learning

I. INTRODUCTION

A. Background

Graphs are natural data structures for applications in diverse domains [1] e.g., recommendation systems, communication, social, and biological networks. Geospatial datasets (e.g., road networks, point clouds, 3D object meshes) can organically be represented as graphs with natural definitions of nodes and edges. Machine learning algorithms for graph analysis require feature vector representations of nodes, edges, substructures, or the whole graph. Instead of hand-engineering task-specific and domain-specific features, recent methods [2] have focused on automatically learning low-dimensional, feature vector representations of graphs (*graph embeddings*) and their components

(e.g., *node embeddings*). Reachability of one node from another is a fundamental concept in graph theory.

In parallel, *self-supervised learning* (SSL) has been an area of active research and has achieved promising performance on both natural language processing (NLP) [3]–[5] and computer vision tasks [6], [7]. Bypassing the need for large, clean, labeled datasets which are expensive to produce in time and money, SSL often uses predefined *pretext tasks* to derive supervision signals directly from unlabeled data by making neural networks predict withheld parts or properties of inputs. SSL aims to learn semantically meaningful, task-agnostic representations of data that can be used as inputs by downstream (usually supervised) task-specific models. SSL has been used to learn context-independent [5] and contextual [3], [4], task-agnostic word embeddings for NLP applications. Most popular SSL techniques for learning visual representations can be classified into two types: *generative* approaches that learn representations under the *pretext* of generating images by modeling the data distribution [8]–[10], and *discriminative* approaches that use *pretext tasks*, designed to efficiently produce labels for inputs (e.g., based on heuristics [11]–[13] or contrastive learning [6], [7]), coupled with a supervised objective.

Geospatial computer vision research has largely focused on active (e.g., SAR) and passive (e.g., optical) imagery data [14]–[16]. This paper proposes a novel technique to make GPS trajectories, which are sequential, spatiotemporal datasets with rich geospatial connectivity information, amenable to computer vision-based analysis by itself or in combination with other data modalities in a multimodal setting. GPS and mobility datasets have been used for diverse applications [17] including transportation modeling [18], public safety and health policy [19], and building mapping services [20], [21].

B. Present Work and Contributions

We present a novel computer vision-based, self-supervised method for learning task-agnostic feature representations of geographic locations from GPS trajectories by modeling tiles resulting from a raster representation of the earth’s surface as nodes of a graph (termed *earth surface graph* (ESG)) and modeling the *observed* GPS trajectories as evidence of *allowed* Markovian paths on this graph. Equivalently, the proposed method learns node embeddings, for all nodes in the ESG, from observed paths on the ESG. The proposed method has

*Corresponding author. Alternative email: swetava@cs.stanford.edu.

two stages. For each node in the ESG, the first stage (Sections III-A–III-D) generates an *image-like* representation, termed *reachability summary*, of the connectivity pattern of the node with its neighbors inferred using observed trajectories passing through the node during a predefined observation time interval, Δt . A scalable and distributed algorithm is proposed for the first stage. The second stage (Section III-E) generates d_R -dimensional, compressed representations from reachability summaries, called *reachability embeddings*, using a fully-convolutional, contractive autoencoder. In Appendix Section E, reachability embeddings are theoretically motivated and interpreted using the Chapman-Kolmogorov Equations (CKE) for Markov chains.

Experiments in Section V evaluate the impact of reachability embeddings on solving downstream geospatial computer vision tasks using supervised semantic segmentation. The model inputs are combinations of image-like derivatives of road network and GPS trajectory data (Section IV) with pixels corresponding to nodes (convenient by design) of the ESG. For five different downstream tasks, we show that using reachability embeddings as input features outperforms baselines (models with input features that do not account for spatial connectivity between nodes) resulting in 4–23% gain in AUPRC (area under precision-recall curve) while requiring upto 67% less trajectory data than the baselines. Thus, reachability embeddings are semantically meaningful representations of geographic locations and denser, more informative representations derived from trajectories compared to those not accounting for spatial connectivity between nodes. Experiments in the latter part of Section V demonstrate that representing mobility data as reachability embeddings facilitates multimodal modeling using computer vision techniques. Pixel-wise embeddings yield multi-channel image-like tensors. Thus, by design, pixel-wise alignment of data modalities implies spatial alignment and early fusion is feasible via concatenation of image-like tensor representations of various data modalities along the channel dimension.

Presence of geospatial features (e.g., traffic control devices, physical traffic restrictions, traffic infrastructure) at a location, ℓ , imposes constraints on traffic flow that manifests in the (i) frequency of, (ii) directionality (from or to ℓ) of, (iii) time taken during, and (iv) distance traveled during observed transitions between ℓ and its neighboring locations across various motion modalities (e.g., walking, driving, biking). Reachability summaries, and consequently reachability embeddings, are designed to be *contextual* representations (analogous to advantage of contextual [4] over fixed [5] word embeddings in NLP) that encode these motion characteristics and traffic flow patterns occurring during Δt . Events that significantly change traffic patterns like road construction, road closures, seasonal effects (e.g., winter vs. summer), etc. change the computed embeddings making them a good input feature candidate for geospatial feature or change detection models.

The following are the main contributions of this paper: (1) A novel, two-stage, computer vision-based, self-supervised procedure to learn low dimensional, geospatial feature representations, called reachability embeddings, from GPS trajectories

is proposed; (2) A scalable and distributed algorithm that can incorporate distance traveled and time taken during a trajectory for computing these embeddings is presented; (3) Enhanced performance on five downstream geospatial feature detection tasks and the ease of multimodal modeling with reachability embeddings is demonstrated; (4) A theoretical interpretation of reachability embeddings using CKE is provided.

While the presentation in this paper is targeted towards geospatial and remote sensing applications, the concept of generating node embeddings based on reachability can be extended to graphs in other domains. For example, (i) graph traversal algorithms can be used to assemble node reachability summaries; (ii) sampled random walk paths, used in [22]–[24], can be used analogous to trajectories to compute reachability summaries for nodes. The generated summaries can be subsequently encoded into reachability embeddings.

II. NOTATION AND PRELIMINARIES

Figure 1 provides a reference schematic for notation and terminology introduced in this and the next section.

Definition II.1. A *zoom- q tile* is the cell or tile resulting from a $2^q \times 2^q$ raster representation of the spherical Mercator projection of the earth’s surface [25]. With the northwest corner tile as origin, every tile, s , is assigned 2D, non-negative integer coordinates, $(\mathbf{x}_s, \mathbf{y}_s)$, which increase in the \mathbf{x} -direction towards the east (right) and in the \mathbf{y} -direction towards the south (down).

Definition II.2. A *GPS trajectory*, $\mathbf{T}_i \in \mathcal{T}(t_0, \Delta t)$, or simply *trajectory*, is a discrete, sequential, and chronological representation of the spatiotemporal movement of an object [18] consisting of sequence of ordered pairs, $\mathbf{T}_i = (\mathbf{p}_1^i, \dots, \mathbf{p}_{n_i}^i)$, where pair, $\mathbf{p}_k^i = (\mathbf{z}_k^i, \mathbf{t}_k^i)$, for $k = 1, \dots, n_i$. \mathbf{p}_k^i is called a *GPS record* or simply a *record*. $\mathcal{T}(t_0, \Delta t)$ is the set of $|\mathcal{T}(t_0, \Delta t)| = M$ trajectories occurring within observation time interval Δt starting at t_0 , n_i is length of trajectory \mathbf{T}_i , and \mathbf{z}_k^i is the zoom-24 tile¹ where the object with trajectory \mathbf{T}_i was present at time \mathbf{t}_k^i . Note: $\mathbf{t}_1^i < \dots < \mathbf{t}_{n_i}^i$ for $i = 1, \dots, M$.

Definition II.3. The *implied allowed transitions multiset*, \mathbf{A}_i^m , associated with trajectory \mathbf{T}_i is the unordered multiset of all ${}^{n_i}P_2 = n_i(n_i - 1)$ ordered pairs of zoom-24 tiles, $(\mathbf{z}_k, \mathbf{z}_l)$, that can be constructed from trajectory \mathbf{T}_i such that $l \geq k$. The unordered set consisting of only the unique ordered pairs is called the *implied allowed transitions set*, \mathbf{A}_i^s .

Ordered pairs generated from the trajectories are modeled as Markovian state transitions over the countably finite state space of zoom-24 tiles. Hence, the trajectory index is dropped from the the ordered pairs in def. (II.3).

Definition II.4. The *earth surface graph* (ESG), $\mathbf{G}_{ES}(\mathbf{V}_{ES}, \mathbf{E}_{ES})$, is an inferred, weighted directed graph representation of the earth’s surface where zoom-24

¹Corresponds to spatial resolution $\approx 2.38\text{m}$ at the equator. Finite area tiles mitigate handling real-valued latitude-longitude GPS location pairs. The tile represents all such pairs within itself.

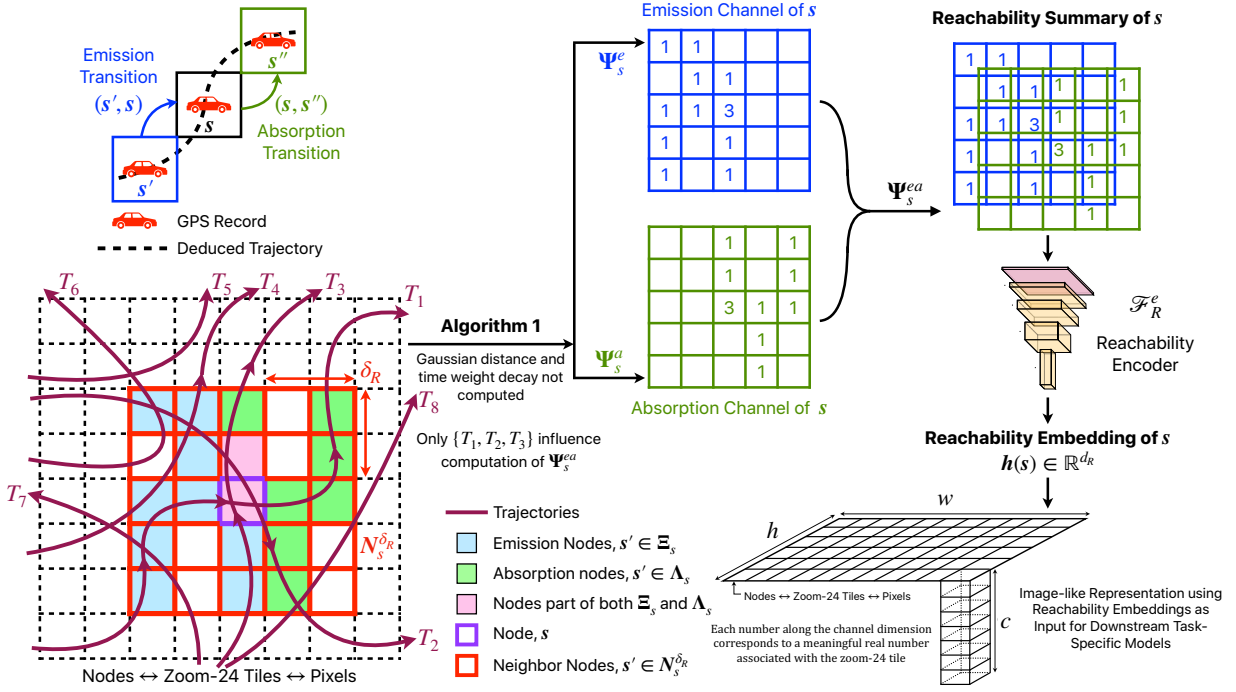


Fig. 1: Schematic describing the algorithm for generation of reachability embeddings from mobility trajectories.

tiles form the nodes, V_{ES} , and the set union of A_i^s of all $T_i \in \mathcal{T}(t_0, \Delta t)$ forms the edge set, E_{ES} . For nodes, $s, s' \in V_{ES}$, $c_i^{(s, s')}$ is the number of occurrences of the transition (s, s') in A_i^m . The directed edge, (s, s') , in E_{ES} is assigned the weight, $w_{ES}^{(s, s')} = \sum_{i=1}^M c_i^{(s, s')}$.

Definitions (II.2), (II.3), and (II.4) present the equivalences used in this paper: (i) zoom-24 tiles \leftrightarrow Markovian state space $V_{ES} \leftrightarrow$ nodes of G_{ES} , (ii) GPS trajectory \leftrightarrow Markovian trajectory in $V_{ES} \leftrightarrow$ path on G_{ES} .

Definition II.5. Given two nodes, $s_1, s_2 \in V_{ES}$, s_2 is *reachable* from s_1 if $\exists (s_1, s_2) \in E_{ES}$.

While $q = 24$ in this paper, the algorithms presented in Section III, are generic and treat q as a parameter.

III. METHODOLOGY

A. Reachability Embeddings

The proposed method learns a mapping function, $f_R : V_{ES} \rightarrow \mathbb{R}^{d_R}$, from nodes in the ESG, V_{ES} , to real valued, d_R -dimensional feature representations we call reachability embeddings. Each record in a trajectory is associated with a motion modality (e.g., driving, walking, biking). A pre-trained, LSTM-based, neural motion modality filter is used to keep only the records corresponding to a chosen modality. Henceforth, we will assume that all records in all trajectories in $\mathcal{T}(t_0, \Delta t)$ have the same motion modality.

Given a node, $s \in V_{ES}$, two sets of nodes can naturally be defined relative to s based on observed reachability patterns in trajectories in $\mathcal{T}(t_0, \Delta t)$. The first set, called the *emission set*

for node s , Ξ_s , consists of nodes from which transitions to s are observed. Such transitions are termed *emission transitions*. The second set, called the *absorption set* for node s , Λ_s , consists of nodes to which transitions from s are observed. Such transitions are termed *absorption transitions*. Thus,

$$\Xi_s := \{s' \mid s' \in V_{ES} \text{ and } (s', s) \in E_{ES}\}, \quad (1a)$$

$$\Lambda_s := \{s' \mid s' \in V_{ES} \text{ and } (s, s') \in E_{ES}\}. \quad (1b)$$

Owing to transitions from s to itself, $s \in \Xi_s, \Lambda_s$. With the intuition that transitions occurring over large spatial distances should not influence the learned representations of nodes, transitions occurring only within a specified buffer, termed the *reachability buffer*, $\delta_R \in \mathbb{Z}^+$, are considered valid. The set of nodes within δ_R of node s , called the *reachable neighborhood*, $N_s^{\delta_R}$, are

$$N_s^{\delta_R} := \{s' \mid s' \in V_{ES}, |x_s - x_{s'}| \leq \delta_R, |y_s - y_{s'}| \leq \delta_R\}. \quad (2)$$

Combining definitions (1a), (1b), and (2), the set of valid emission and absorption transitions with respect to node s , denoted as $\Xi_s^{\delta_R}$ and $\Lambda_s^{\delta_R}$ respectively, are

$$\Xi_s^{\delta_R} = \Xi_s \cap N_s^{\delta_R} \quad \text{and} \quad \Lambda_s^{\delta_R} = \Lambda_s \cap N_s^{\delta_R}. \quad (3)$$

The Cartesian grid structure of nodes in V_{ES} can be exploited to create an *image-like* representation of the reachability pattern for all nodes, s , implied by their $\Xi_s^{\delta_R}$ and $\Lambda_s^{\delta_R}$. With $L = 2\delta_R + 1$, for each $s \in V_{ES}$, construct two zero-initialized $L \times L$ square matrices, Ψ_s^e (*emission channel*) and Ψ_s^a (*absorption channel*), with the top-left corner entries assigned the index $(0, 0)$, first index increasing rightward, and

second index increasing downward. For every $s' \in \Xi_s^{\delta_R}$ and $s'' \in \Lambda_s^{\delta_R}$, set

$$\Psi_s^e[x_{s'} - x_s + \delta_R, y_{s'} - y_s + \delta_R] = w_{ES}^{(s',s)}, \quad (4a)$$

$$\Psi_s^a[x_{s''} - x_s + \delta_R, y_{s''} - y_s + \delta_R] = w_{ES}^{(s,s'')}. \quad (4b)$$

The (δ_R, δ_R) -indexed entry equals $w_{ES}^{(s,s)}$ in both, Ψ_s^e and Ψ_s^a . Matrices Ψ_s^e and Ψ_s^a , have three important properties: (i) they preserve the relative spatial proximity of nodes in $N_s^{\delta_R}$ with respect to s ; (ii) Ψ_s^e and Ψ_s^a are the spatial-relation preserving matrix representations of the column and row corresponding to s of the adjacency matrix of nodes in $N_s^{\delta_R}$; (iii) normalization by matrix sums yields the Markovian transition probabilities to s from all $s' \in N_s^{\delta_R}$ and from s to all $s'' \in N_s^{\delta_R}$, respectively. Matrices, Ψ_s^e and Ψ_s^a , are treated as two channels and stacked along the channel dimension to create a $L \times L \times 2$ -dimensional, image-like entity, called the *reachability summary*, Ψ_s^{ea} , for node s . A contractive [26], fully-convolutional autoencoder, \mathcal{F}_R , is trained to generate compressed, robust, low-dimensional representations of Ψ_s^{ea} , called the *reachability embedding*, $h(s) \in \mathbb{R}^{d_R}$, for node s using contractive reconstruction of Ψ_s^{ea} as the self-supervision task for representation learning. Reachability embeddings, produced for every node $s \in V_{ES}$ by employing the encoder component of \mathcal{F}_R , captures the spatial connectivity of node s to nodes in $N_s^{\delta_R}$ as evidenced by observed trajectories (modeled as Markovian) in $\mathcal{T}(t_0, \Delta t)$.

B. Distributed data-parallel algorithm for generating reachability summaries

Computing Ψ_s^{ea} for all $s \in V_{ES}$ requires analysing all M trajectories in $\mathcal{T}(t_0, \Delta t)$ and computing upto 2^{2q} matrices of size $L \times L \times 2$, where $q = 24$. In practice, however, given a choice of motion modality, only a tractable subset (e.g., no driving or walking on water bodies) of the zoom-24 tiles (called *active tiles* and denoted as \overline{V}_{ES}) have trajectories that pass through them. The embedding vector is set to zero for non-active tiles. Algorithm 1 proposes an efficient method to compute $\Psi_s^{ea} \forall s \in \overline{V}_{ES}$ from $\mathcal{T}(t_0, \Delta t)$ assuming all its trajectories are of the same, user-chosen motion modality. The key idea in Algorithm 1 is to compute $w_{ES}^{(s_1, s_2)}$, $s_1, s_2 \in \overline{V}_{ES}$ for all unique (s_1, s_2) pairs (emission and absorption transitions) independently in parallel and then assemble the dense tensor, Ψ_s^{ea} , for each $s \in \overline{V}_{ES}$. For node s , let $S_s^a : r_s(s') \rightarrow c_s^a(s')$ be a map constructed from observed absorption transitions (s, s') for $s' \in N_s^{\delta_R}$. Here, $r_s(s')$ is the row-major index of s' in Ψ_s^a and $c_s^a(s')$ is the frequency of observing the transition (s, s') , which equals $w_{ES}^{(s,s')}$ when Algorithm 1 finishes. Analogously, define $S_s^e : r_s(s') \rightarrow c_s^e(s')$ with $c_s^e(s')$ being frequency of observing transition (s', s) which equals $w_{ES}^{(s',s)}$ when algorithm 1 finishes. The reachability map, $S : s \rightarrow (S_s^e, S_s^a)$, is a map from each node $s \in \overline{V}_{ES}$ to the tuple of the node's emission and absorption transition maps, S_s^e and S_s^a . S_s^e and S_s^a are sparse representations of the information required to assemble Ψ_s^e and Ψ_s^a (i.e., Ψ_s^{ea}).

C. Capturing more than spatial connectivity in $\Psi_s^e, \Psi_s^a, \Psi_s^{ea}$

The emission (Ψ_s^e) and absorption (Ψ_s^a) channels, as defined in Section III-A, capture only the spatial connectivity between two nodes. In Algorithm 1, if the transition (s, s') , with $s' \in N_s^{\delta_R}$, is observed in trajectory T_i , $\Psi_s^a(s')$ and $\Psi_{s'}^e(s)$ are both incremented by 1. These increments, denoted as $\Delta\Psi_s^a(s')$ and $\Delta\Psi_{s'}^e(s)$, are independent of the distance covered, $\Delta_{d,i}^{(s,s')}$, or time taken, $\Delta_{t,i}^{(s,s')} = t_{s'}^i - t_s^i$, to reach s' from s in trajectory T_i . Note that time taken and distance covered depend on the trajectory if multiple paths between the two nodes exist. Gaussian weight decay is used to incorporate this information into $\Psi_s^a(s')$ and $\Psi_{s'}^e(s)$. The Gaussian with parameters μ and σ is $G(\mu, \sigma) = (1/\sqrt{2\pi}\sigma)e^{-\mu^2/2\sigma^2}$. Given σ_d and σ_t , two user-supplied hyperparameters that control the strengths of the distance and time weight decays, respectively, the weighted increment² in $\Psi_s^a(s')$ and $\Psi_{s'}^e(s)$ associated with the transition (s, s') in trajectory T_i , is $\Delta\Psi_s^a(s') = \Delta\Psi_{s'}^e(s) := \Delta w(s, s'; \sigma_d, \sigma_t) = G(\Delta_{d,i}^{(s,s')}, \sigma_d) \cdot G(\Delta_{t,i}^{(s,s')}, \sigma_t)$. A modified `analyzeTrajectories` procedure for Algorithm 1 implementing the equation above is presented in Algorithm 3 in Appendix Section B.

D. Scalability of Algorithm 1

Algorithm 1 modified using Algorithm 3 is implemented in a distributed data-parallel fashion using the Scala-based, Spark Dataset API. To demonstrate scalability, the publicly available T-Drive dataset [27] is suitably pre-processed (detailed in Appendix Section C) to obtain three datasets of 2000, 8000, and 64000 trips on which strong scaling [28] analysis is performed. The Spark implementation is executed on these datasets by varying the number of CPU cores from 10 to 2000. The mean value of the runtime and strong scaling efficiency³ [28] obtained over seven runs is plotted in Figures 2a and 2b, respectively, as the number of CPU cores are varied for the three datasets. Typical characteristics of parallel applications such as increased efficiency on larger number of cores for bigger problem sizes, efficiency drop for smaller problem sizes as number of cores increases, and communication latency dominating compute time are observed.

Another scalability test is to increase the number of trajectories processed on a fixed number of cores. Plotting the runtime of Algorithm 1 as the number of trajectories are varied from 1000 to 64000 and the number of cores are varied from 10 to 100 in Figure 2c, it is empirically observed that the algorithm run-time scales sub-linearly with the number of trajectories with a power law exponent of roughly 1/5. While Figure 2 and the observations in this sub-section are stated specifically for the T-Drive dataset, similar observations are empirically observed on much larger proprietary datasets (described in section IV). Scalability of Algorithm 1 is especially crucial in

²The weight-decays can be thought of as the *attention* paid by s to a neighboring node $s' \in N_s^{\delta_R}$ analogous to the soft-attention mechanism.

³For a given dataset size, if t_c^{id} is the ideal time required on c CPU cores, $t_{c_0}^{\text{m}}$ is the time measured when executing algorithm on c_0 cores, and e_{ss} is strong scaling efficiency, then $t_c^{\text{id}} = t_{c_0}^{\text{m}}(c_0/c)$ and $e_{ss} = t_c^{\text{id}}/t_{c_0}^{\text{m}}$.

Algorithm 1 Generate Ψ_s^{ea} for all nodes $s \in \overline{V}_{ES}$

Input: $\mathcal{T}_{t_0}^{t_0+\Delta t} = (T_1 \dots T_M)$, $|T_i| = n_i \forall i \in [1, M]$, $T_i = (p_1^i, \dots, p_{n_i}^i)$ where $p_k^i = (z_k^i, t_k^i)$

Output: Map $\Psi : s \rightarrow \Psi_s^{ea}$

Parameters: δ_R . Define $L = 2\delta_R + 1$. Let $\bar{n} = \max\{n_1, \dots, n_M\}$

procedure ReachabilitySummaryGenerator()

Initialize map $\Psi = \emptyset$

$S \leftarrow \text{analyzeTrajectories}(\mathcal{T}_{t_0}^{t_0+\Delta t})$

for node $s \in S$ **do**

Arrays Ψ_s^e, Ψ_s^a . Zero-initialized of length L^2

$S_s^e, S_s^a \leftarrow S[s]$

$\Psi_s^e[\kappa] = S_s^e[\kappa]$ **for** all keys κ of map S_s^e

$\Psi_s^a[\kappa] = S_s^a[\kappa]$ **for** all keys κ of map S_s^a

Reshape Ψ_s^e and Ψ_s^a to $L \times L$ arrays

$\Psi_s^{ea} \leftarrow$ Stack channels Ψ_s^e, Ψ_s^a

$\Psi[s] \leftarrow \Psi_s^{ea}$

function analyzeTrajectories($\mathcal{T}_{t_0}^{t_0+\Delta t}$) $\triangleright \mathcal{O}(M\bar{n}^2)$

Initialize map $S = \emptyset$

for trajectory $T_i \in \mathcal{T}_{t_0}^{t_0+\Delta t}$ **do**

for $k \leftarrow 1$ to n_i **do**

for $l \leftarrow k$ to n_i **do**

if $z_l^i \in N_{z_k^i}^{\delta_R}$ **then**

if $z_k^i \notin S$, $S[z_k^i] = (S_{z_k^i}^e = \emptyset, S_{z_k^i}^a = \emptyset)$

if $z_l^i \notin S$, $S[z_l^i] = (S_{z_l^i}^e = \emptyset, S_{z_l^i}^a = \emptyset)$

$r_{z_k^i}(z_l^i) \leftarrow \text{getIndex}(z_k^i, z_l^i)$

$r_{z_l^i}(z_k^i) \leftarrow \text{getIndex}(z_l^i, z_k^i)$

$c_{z_k^i}^a(z_l^i) = c_{z_l^i}^e(z_k^i) = 1.0 \quad \triangleright$ Refined in III(C)

$S_{z_k^i}^e[r_{z_k^i}(z_l^i)] += c_{z_k^i}^a(z_l^i)$

$S_{z_l^i}^e[r_{z_l^i}(z_k^i)] += c_{z_l^i}^e(z_k^i)$

return S

function getIndex(s, s') $\triangleright \mathcal{O}(1)$

return $L(y_{s'} - y_s + \delta_R) + (x_{s'} - x_s + \delta_R)$

a setting where embeddings need to be generated periodically for maintaining and updating geospatial maps given the large number of zoom-24 tiles that are possible on the surface of the earth. Contrast this to the vocabulary sizes ($\mathcal{O}(10^6)$) used in NLP approaches like word2vec [5]. From a machine learning systems perspective [20], scalable generation of reachability embeddings makes it convenient for feature stores [29] to store generated embeddings as reusable features for training and deployment of multiple downstream task-specific models.

E. Encoding reachability summaries

The neural architecture of the custom-designed, fully-convolutional, contractive autoencoder, \mathcal{F}_R , is shown in Figure 6 of Appendix Section D. During training, the encoder, \mathcal{F}_R^e , maps Ψ_s^{ea} to a compressed representation, $\mathbf{h}(s) = \mathcal{F}_R^e(\Psi_s^{ea}) \in \mathbb{R}^{d_R}$, which is then used by the decoder, \mathcal{F}_R^d , to reconstruct Ψ_s^{ea} . The objective function to be minimized, \mathcal{L} , consists of the reconstruction loss, \mathcal{L}_{rec} , regularized by the contractive [26] loss, \mathcal{L}_{con} , shown in equations (5a), (5b), and (5c).

$$\mathcal{L}_{rec}(s; \theta) = \|\Psi_s^{ea} - \mathcal{F}_R^d(\mathcal{F}_R^e(\Psi_s^{ea}))\|_2^2 \quad (5a)$$

$$\mathcal{L}_{con}(s; \theta) = \sum_{i=0}^{d_R-1} \|\nabla_{\Psi_s^{ea}} \mathbf{h}_i(s)\|_2^2 \quad (5b)$$

$$\mathcal{L}(\theta) = \mathbb{E}_{s \sim \overline{V}_{ES}} [\mathcal{L}_{rec}(s; \theta) + \lambda \mathcal{L}_{con}(s; \theta)] \quad (5c)$$

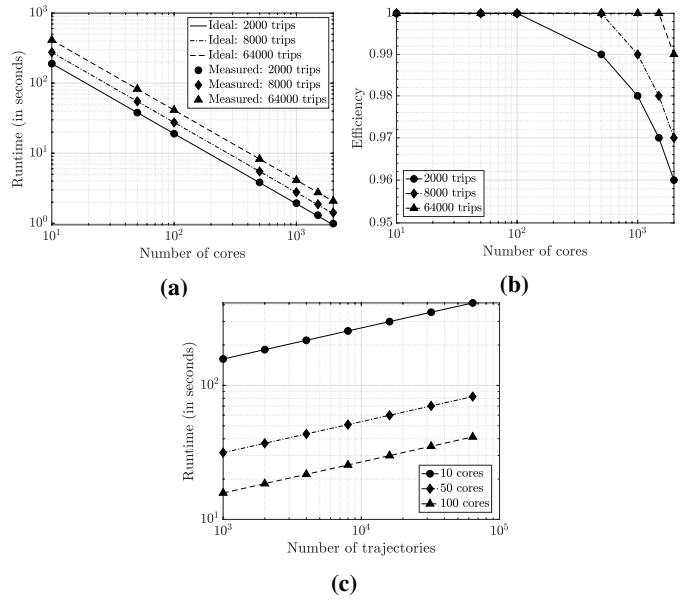


Fig. 2: Strong scaling analysis of (a) runtime, (b) efficiency, and (c) scaling with number of trajectories of Algorithm 1.

Here, θ are the parameters of \mathcal{F}_R and λ is a hyperparameter. A ReLU nonlinearity forces all embeddings to be positive real vectors. Since there is no upper bound on the pixel values in Ψ_s^{ea} , the input to \mathcal{F}_R is log-normalized and the reconstruction output exponentiated (similar to [21]) before calculating the losses (5a)–(5c) to train the neural network stably and avoid spurious effects from the large dynamic range of Ψ_s^{ea} . The training set is composed of zoom-24 tiles from a variety of geographically diverse regions. Once trained, \mathcal{F}_R^e is used to generate the embeddings, $\mathbf{h}(s)$, for all $s \in \overline{V}_{ES}$. As shown in [26], contractive regularization forces $\mathbf{h}(s)$ to be invariant to small perturbations of Ψ_s^{ea} (locally contractive in neighborhood of Ψ_s^{ea}) thereby yielding robust and sparse compressed representations of Ψ_s^{ea} . Presence of \mathcal{L}_{con} yields only a few non-zero elements in the embeddings for most nodes. From the perspective of energy-based models, \mathcal{L}_{con} helps minimize the volume of low-energy regions to help learn semantically meaningful representations.

IV. DATASET DESCRIPTION

Small, publicly available GPS trajectory datasets (e.g., [27], [30], [31], [32]) (i) have varying sampling rates with incomplete trajectories, (ii) do not cover the entire road network in the geographical areas they are obtained from, (iii) have inconsistent observation times and intervals, and (iv) lack geographical diversity. Most importantly, in order to test the utility of learned unsupervised representations like reachability embeddings on downstream supervised tasks, we need labeled data for these tasks. To the best of our knowledge, there is no sufficiently large, geographically diverse, publicly available GPS trajectory dataset covering a sufficiently large observation time interval which can be matched to labeled datasets collected for various geospatial tasks in the same geographical regions.

To overcome this difficulty, we use *probe data* [33], which is a privacy-preserving, structured, sequential, and proprietary dataset⁴. Labels are collected across a variety of geographic regions, in the form of *Well-Known Text* (WKT) polygons⁵ [34], for training supervised semantic segmentation models for five downstream geospatial tasks described in Section V.

There are multiple ways, called *local aggregate representations* (LAR), simpler to implement than reachability embeddings, to create semantically meaningful, multi-channel, *image-like* representations of GPS trajectory datasets making them amenable for use by downstream computer vision models for geospatial tasks. The pixel values for a given zoom-24 tile are obtained by analyzing records observed *only in that tile* (local) to yield one (single-channel) or a vector of aggregate counts (multi-channel). Three such representations are count-based raster map (CRM), heading count-based raster map (HCRM), and speed count (SC).

Given the Cartesian grid structure of V_{ES} , associating c meaningful real numbers (represented as c channels) to each tile (pixel) within any set of contiguous $h \times w$ tiles (pixels) forms a $h \times w \times c$ -dimensional image-like representation of the geographic location represented by the $h \times w$ tiles, as shown in Figure 1. A *count-based raster map* (CRM) is a single-channel representation ($c = 1$) where the value of each pixel is the number of occurrences of GPS records in the zoom-24 tile corresponding to the pixel counted over all trajectories in $\mathcal{T}(t_0, \Delta t)$. Instead, if the occurrences of records per pixel is bucketed based on the direction in which the record was heading into 12 buckets of 30° , and each bucket is represented as an individual channel ($c = 12$), we obtain the *heading count-based raster map* (HCRM) representation. Similarly, if the occurrences of records per pixel are bucketed based on the speed of the GPS records into 14 buckets of 5 miles per hour starting from 0, and each bucket is represented as an individual channel ($c = 14$), we obtain the *speed-count* (SC) representation.

If pixel values are set to the reachability embeddings of the corresponding zoom-24 tile, we obtain the d_R -channel reachability embeddings representation. Reachability embeddings may be viewed as a mapping of a geographic location to an *embedding field* with $h \times w \times d_R$ parameters based on observed mobility in that geographic location over a specified observation interval Δt . In contrast to LAR, reachability embeddings are *global* pixel-level representations — representation of tile s depends on all activity in $N_s^{\delta_R}$. Trajectory-independent, road geometry related information is supplied using the *road network presence* (RNP) channel ($c = 1$) which is a binary representation of the road network as a raster image where each pixel is assigned the value 1 if the corresponding zoom-24 tile has a road segment present and the value 0 otherwise. In this paper, $h = w = 256$.

⁴Probe data is similar to publicly available GPS trajectory datasets such as [27], [30]. Section titled “Probe data and privacy” in [33] has more details.

⁵Geospatial feature coordinates are suitably buffered to obtain WKT polygons. Assigning value 1 to pixels inside and 0 to those outside the polygons yields labels for semantic segmentation.

Analogous to spectral imaging for remote sensing, where channels corresponding to different spectral bands capture additional information of a geographic location to complement their RGB (visible spectrum) counterparts, CRM, HCRM, SC, and reachability-based channels are multi-channel, image-like representations of information of the location deduced from trajectory data. The RNP channel can be considered as an *ontological* representation (where traffic *should* exist) of G_{ES} while trajectory datasets are the *epistemological* representation (where traffic *actually* exists) of G_{ES} .

V. EXPERIMENTS AND RESULTS

The impact of reachability embeddings is evaluated on the performance of supervised, pixel-wise, semantic segmentation models for five downstream geospatial tasks, viz. (i) detection of overpasses, (ii) detection of pedestrian crosswalks, (iii) detection of driving access (entry/exit) points, (iv) detection of locations with traffic lights, and (v) detection of locations with stop signs. For all tasks, the UNet [35] architecture is trained to minimize the pixel-wise binary cross-entropy between the predicted segmentation map and labels. A 60-20-20 randomly chosen training-validation-testing set is created and fixed for all subsequent experimentation.

Since reachability embeddings encode spatial connectivity (frequency and directionality of transitions) and speed (distance covered and time taken), baseline models using a combination (best model obtained after ablation study) of LAR-based channels as inputs are compared to the same model trained by replacing all LAR-based channels with reachability embeddings-based channels. Direct comparison between (i) models using inputs accounting for spatial connectivity versus inputs that do not, and (ii) the proposed self-supervised, learned representations versus explicit encoding of different aspects of trajectory data, demonstrates that learning more informative representations like reachability embeddings lead to better performance on downstream supervised tasks while using lesser trajectory data (i.e., reduced observation intervals, Δt). Probe data-based trajectory sets filtered for driving ($\mathcal{T}^d(t_0, \Delta t)$) and walking ($\mathcal{T}^w(t_0, \Delta t)$) motion modalities are obtained. Reachability embeddings (RE) are generated for both walking and driving modalities and denoted by WRE and DRE, respectively. CRM is also computed for both walking and driving modalities and denoted by WCRM and DCRM, respectively. HCRM and SC are calculated only for driving modality. For the models built to evaluate the predictive power of reachability embeddings on downstream tasks (results in table II), only trajectory-based representations and/or RNP are used as inputs. No inputs from other remote sensing modalities (e.g., satellite or SAR imagery) are used since the goal of these experiments is to assess image-like representations of trajectory data for solving geospatial tasks as opposed to building the best model (which may combine multimodal inputs) for a given task. The last part of this section (results in table III) deals with multimodal modeling with reachability embeddings.

TABLE I: Combination of LAR-based channels identified from ablation study for the best baseline model for the downstream tasks. Corresponding inputs for reachability embeddings (RE) based models.

Application	Baseline Models					Models w/ RE		
	RNP	DCRM	WCRM	HCRM	SC	RNP	DRE	WRE
Overpass	×	×	×	✓	×	×	✓	×
Crosswalk	×	✓	✓	×	×	×	✓	✓
Access Point	✓	×	×	×	×	✓	✓	×
Traffic Lights	✓	✓	×	✓	✓	✓	✓	×
Stop Signs	✓	✓	✓	✓	✓	✓	✓	✓

A. Baselines for Each Downstream Task

An ablation study is performed over all 31 possible combinations of RNP, DCRM, WCRM, HCRM, and SC to identify the best combination of LAR-based input channels to build the baseline model for each task. Each input combination is concatenated along the channel dimension. Inputs are log-normalized so as to avoid spurious effects of the large dynamic range owing to the lack of an upper bound on the pixel-values. Hyperparameters are tuned on the validation set. Once trained, segmentation map predictions from all 31 models for each task is obtained for all examples of the test set. For every task, using ground truth labels available for the test set, the mean precision-recall curve (PRC) and the area under the curve (AUPRC) over all test examples is obtained for each of the 31 models by varying the confidence threshold of the binary segmentation outputs. The combination of LAR-based input representations that yield the highest AUPRC for each task is chosen as the baseline model for that task and is documented in Table I.

B. Experiments with Reachability Embeddings

With the goal of assessing the importance of self-supervised representations that account for various aspects of spatial connectivity (transition frequency, time, and distance), reachability-based models, documented in Table I, are built that directly correspond to the baseline models by simply replacing any combination of DCRM, HCRM, and SC channels with DRE and replacing WCRM with WRE. Two variants each of the reachability-based models and baseline models are compared: reachability-based models that use DRE and WRE computed from trajectories observed within $[t_0, t_0 + \Delta t]$ with varying embedding sizes (a) $d_R = 8$ and (b) $d_R = 16$; baseline models using LAR-based inputs computed using trajectories observed within (c) $[t_0, t_0 + \Delta t]$ and (d) $[t_0, t_0 + 3\Delta t]$. Since LAR are frequency counts of GPS records which can be noisy, increasing the observation time interval is a proxy-method to increase the signal-to-noise ratio (SNR) of the LAR-based channels. Increasing d_R allows the contractive autoencoder, \mathcal{F}_R , to retain more semantically meaningful information that may be helpful for the reconstruction task. Since larger d_R results in lowering the maximum possible batch-size (governed by GPU memory) for training models, we refrain from increasing d_R beyond 16 given the practical constraints of heterogeneous compute clusters having GPUs with varied memory.

C. Impact of Reachability Embeddings

Table II shows the quantitative comparison of the performance for all 5 downstream tasks, quantified using AUPRC on the test set after model training converges (100 epochs), between the two variants each of baseline and reachability-based models. Figure 3 shows the corresponding precision-recall curves. Three key observations emerge: (i) increasing the observation interval (increasing SNR) for computing LAR and increasing d_R increases AUPRC — for most values of recall, precision increases due to reduction in false positives; (ii) reachability-based models outperform the baseline models, including those using LAR computed by observing 3 times more trajectories: the AUPRC gain by the inferior reachability-based model ($d_R = 8$) over the superior baseline model (observation interval $3\Delta t$) varies from 1.6% for stop signs detection task to 18.4% for the access point detection task; (iii) for the same observation time interval, Δt , simply replacing the LAR-based inputs by reachability embeddings ($d_R = 16$) results in a AUPRC gain that varies from 4.1% for the stop signs detection task to 23.3% for the access point detection task. These observations conclusively demonstrate that reachability embeddings are more informative, denser representations of trajectory data requiring lesser trajectories (upto 67% less) to compute. Thus, reachability embeddings may be used to compute semantically meaningful representations of trajectories in geographical areas with less traffic or to build computer vision-based models for low-resource geospatial tasks.

D. Visual Evaluation

Good learned representations ensure that similar entities have similar representations [36]. Since reachability embeddings are compressed representations of connectivity patterns, geographic locations with similar traffic patterns must have similar reachability embeddings. A visual demonstration of the semantics captured by reachability embeddings is shown by the UMAP projection of embeddings computed in a large geographical area in Figure 4a. Black points denoting nodes on highways (mostly straight roads with few connections to neighbors) are clustered and cleanly separated from red points that denote residential roads. Visualizations of labels and predictions from reachability-based semantic segmentation models for three of the tasks in table II is shown in Figure 5.

E. Multimodal Modeling with Reachability Embeddings

Representation, alignment, and fusion are three key challenges in multimodal machine learning [37]. The main advantage of reachability embeddings is the ability to combine mobility, satellite imagery, and imagery-like representations of graph data (e.g., RNP) for building multimodal models that either complement each other (when all data modalities are present) or supplement the incompleteness or corruption in a given modality. We choose the overpass, crosswalk, and access point detection tasks for this demonstration. In satellite imagery, tree canopies and buildings, among others, often occlude crosswalks and access points while some overpasses are hard to distinguish from junctions. We show that having mobility data

TABLE II: Comparison of AUPRC (\uparrow is better) obtained from both variants of LAR and reachability-based models for 5 downstream geospatial tasks. Percentage gain compared to first row shown in brackets.

Input Channels	Observation Interval	Overpass Detection	Crosswalk Detection	Access Point Detection	Traffic Lights Detection	Stop Signs Detection
LAR Baseline	Δt	0.782	0.922	0.663	0.890	0.921
LAR Baseline	$3\Delta t$	0.785 (+0.4%)	0.923 (+0.1%)	0.684 (+3.0%)	0.925 (+3.8%)	0.924 (+0.3%)
Reachability, $d_R = 8$	Δt	0.899 (+15.0%)	0.959 (+4.0%)	0.843 (+21.4%)	0.974 (+8.6%)	0.938 (+1.9%)
Reachability, $d_R = 16$	Δt	0.925 (+18.3%)	0.976 (+5.9%)	0.864 (+23.3%)	0.985 (+9.6%)	0.959 (+4.1%)




Fig. 3: Precision-Recall curves obtained from both variants of LAR and reachability-based models for the 5 downstream tasks.

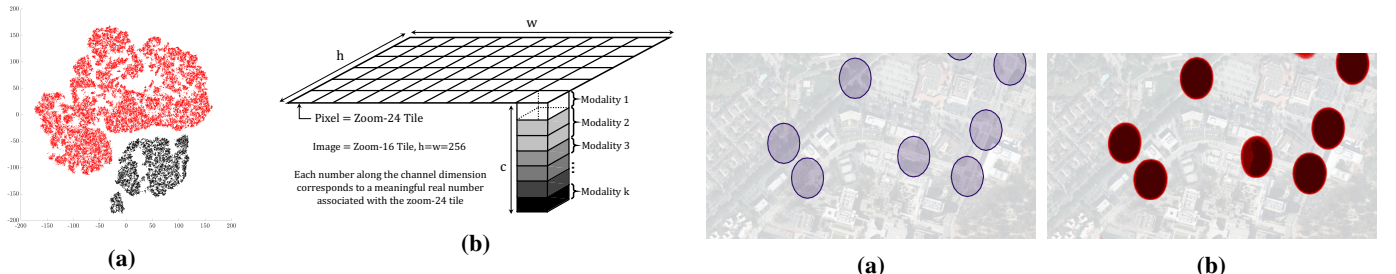


Fig. 4: (a) UMAP of embeddings in a geographic location: highways nodes (black), otherwise (red). (b) Alignment and fusion of data modalities.

can help the model achieve better performance. We compare the AUPRC of models using (i) only reachability embeddings, (ii) only satellite imagery, (iii) using satellite imagery and RNP, and (iv) using satellite imagery, RNP, and reachability embeddings. The design of reachability embeddings implies that aligning modalities at the pixel-level results in spatial alignment since pixels are zoom-24 tiles. The data modalities are concatenated along the channel dimension (called early fusion, shown in figure 4b). Two neural architectures, UNet [35] and SegNet [38], are compared for semantic segmentation. As can be seen in table III, multimodal inputs combining mobility data, satellite imagery, and road network graph improves model performance by 2–4% compared to unimodal inputs, irrespective of neural architecture, which is facilitated by reachability embeddings.

VI. RELATED WORK

A. Graph Node Embeddings

While a number of deep learning methods have been proposed for generating node embeddings of a known graph [2], they can be grouped into three distinct categories: (a) techniques that efficiently sample random walk paths from a given graph

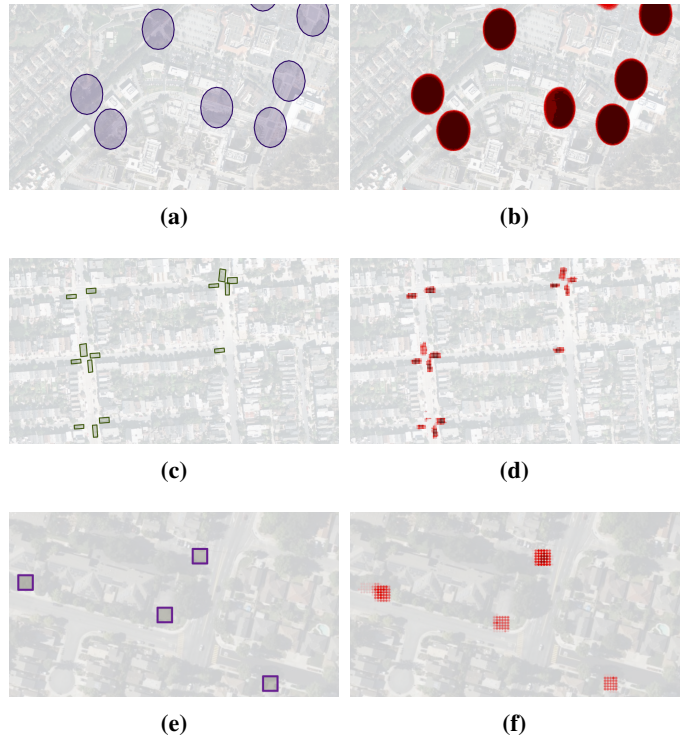


Fig. 5: Examples of labels and semantic segmentation model predictions from reachability-based models on the test set for 3 downstream tasks. Subfigures (a) and (b) show labels and predictions for traffic lights, (c) and (d) show labels and predictions for stop signs, and (e) and (f) show labels and predictions for access points. Satellite imagery is superposed for visual reference only.

TABLE III: Reachability embeddings (RE, $d_R = 16$) facilitate multimodal modeling improving model performance (AUPRC, \uparrow is better) for overpass (Task 1), crosswalk (Task 2), and access point (Task 3) detection.

Input Modalities	UNet [35]			SegNet [38]		
	Task 1	Task 2	Task 3	Task 1	Task 2	Task 3
RE	0.925	0.976	0.864	0.917	0.981	0.881
Sat. Img.	0.923	0.968	0.856	0.901	0.970	0.874
Sat. Img., RNP	0.925	0.972	0.856	0.905	0.974	0.875
Sat. Img., RNP, RE	0.949	0.988	0.895	0.939	0.992	0.905

and generate node embeddings by processing the paths as sequences using NLP-inspired methods like skip-gram [5] such as [22]–[24] or explicitly using recurrent neural networks (RNN) based on LSTMs or GRUs such as [39]; (b) techniques that generate node embeddings by compressing matrix representations derived from graphs (e.g., adjacency matrix, positive point-wise mutual information matrix) such as [40], [41]; (c) techniques that generate node embeddings directly from the graph using graph neural networks (GNNs) and its variants [42] (e.g., graph convolutional network (GCN)). In contrast to techniques in category (a), (b), and (c), reachability embeddings *infer* an implicit graphical structure for the earth’s surface (earth surface graph) by treating geographic locations as nodes and observed GPS trajectories as allowed Markovian paths on this graph. Edges of the earth surface graph can be inferred using the observed paths. Reachability summaries are image-like representations for nodes deduced from observed paths and embedding generation is posed as an image reconstruction task using convolutional neural architectures as opposed to NLP-inspired models and RNN-based architectures used in techniques in category (a). While reachability summaries are matrix representations, they stand in contrast to techniques in category (b) since summaries are derived from observed trajectories and not from the graph or random walk paths sampled from a graph.

B. Self-Supervised Learning for Computer Vision

SSL has been successfully applied to learn effective visual representations [43], [44]. Generative methods learn representations by modeling the data distribution [8]–[10] or reconstructing the input or feature [45], [46]. Examples of heuristic, discriminative pretext tasks used to learn representations include context prediction [11], solving jigsaw puzzles of image patches [13], and predicting image rotations [12]. Among discriminative methods, contrastive methods currently achieve state-of-the-art performance in SSL [6], [7]. Reachability summary generation and its contractive reconstruction can together be viewed as the generative pretext task that encodes the co-occurrence relationships inherent in geospatial transitions (resulting from interaction of traffic and local transport infrastructure) to obtain reachability embeddings.

C. GPS Trajectory Embeddings

Most existing methods using GPS record or trajectory representations process trajectories similar to sequences with

NLP-inspired methods like skip-gram or RNNs e.g., location similarity prediction [47], motion modality classification [48]–[50], demographic attribute prediction [51], and living pattern recognition [52]. To the best of our knowledge, no prior work uses the Markovian concept of reachability and a computer vision-based SSL pretext task to learn self-supervised, contextual representations of geographic locations.

VII. CONCLUSIONS

In this paper, we propose reachability embeddings, a novel, computer vision-based, self-supervised method to learn representations of geographic locations (zoom-24 tiles or nodes of the ESG) from observed GPS trajectories modeled as Markovian paths. Reachability summaries for each node are image-like representations that capture the inferred connectivity pattern based on evidence from observed trajectories (mobility data) of the inflow (emission transitions) and outflow (absorption transitions) of traffic at the node. Summaries are compressed to a vector representation for each node, called reachability embeddings, using the encoder of a contractive, fully-convolutional autoencoder trained to reconstruct reachability summaries. Reachability summary generation and its contractive reconstruction can together be viewed as the pretext task to obtain reachability embeddings. A theoretical interpretation of reachability embeddings is provided using the Chapman-Kolmogorov equations. The contractive regularization incentivizes robustness and invariance of embeddings to small perturbations in the summary leading embeddings of similar locations being close to each other in the learned, low-dimensional manifold. The proposed scalable, distributed algorithm (Algorithm 1) to generate reachability summaries, encoding spatial connectivity along with distance travelled and time taken during node transitions, shows good strong-scaling performance. Experiments in Section V confirm that reachability embeddings are more informative and denser (uses up to 67% less trajectory data compared to LAR) representations of geographic locations derived entirely from trajectory data leading to gains of 4–23% in AUPRC on five different downstream supervised prediction tasks. Reachability embeddings are demonstrated to facilitate multimodal learning in geospatial computer vision with spatiotemporal mobility data as one of the data modalities.

Comparing versions of embeddings computed by varying t_0 , Δt , or both can be used for detecting changes in geospatial features over time, identifying locations with timed turn restrictions or road closures, identifying seasonal patterns, etc. As evidenced in this work, self-supervision is a promising approach to analyze spatiotemporal trajectory datasets.

REFERENCES

- [1] J. Leskovec *et al.*, *Mining of massive data sets*. Cambridge University Press, 2020.
- [2] H. Cai *et al.*, “A comprehensive survey of graph embedding: Problems, techniques, and applications,” *arXiv:1709.07604*, 2018.
- [3] T. B. Brown *et al.*, “Language models are few-shot learners,” *arXiv preprint arXiv:2005.14165*, 2020.
- [4] J. Devlin *et al.*, “BERT: Pre-training of deep bidirectional transformers for language understanding,” *arXiv:1810.04805*, 2018.

- [5] T. Mikolov *et al.*, “Efficient estimation of word representations in vector space,” *arXiv:1301.3781*, 2013.
- [6] T. Chen *et al.*, “A simple framework for contrastive learning of visual representations,” in *ICML*, 2020.
- [7] J. Grill *et al.*, “Bootstrap your own latent: A new approach to self-supervised learning,” in *NeurIPS*, 2020.
- [8] J. Donahue *et al.*, “Large scale adversarial representation learning,” *arXiv:1907.02544*, 2019.
- [9] D. P. Kingma *et al.*, “Auto-encoding variational Bayes,” *arXiv:1312.6114*, 2013.
- [10] I. J. Goodfellow *et al.*, “Generative adversarial networks,” *arXiv:1406.2661*, 2014.
- [11] C. Doersch *et al.*, “Unsupervised visual representation learning by context prediction,” *arXiv:1505.05192*, 2015.
- [12] S. Gidaris *et al.*, “Unsupervised representation learning by predicting image rotations,” in *ICLR*, 2018.
- [13] M. Noroozi *et al.*, “Unsupervised learning of visual representations by solving jigsaw puzzles,” *ECCV*, 2016.
- [14] S. Ganguli *et al.*, “GeoGAN: A conditional GAN with reconstruction and style loss to generate standard layer of maps from satellite images,” *arXiv:1902.05611*, 2019.
- [15] L. Ma *et al.*, “Deep learning in remote sensing applications: A meta-analysis and review,” *ISPRS JPRS*, vol. 152, pp. 166–177, 2019.
- [16] A. Perez *et al.*, “Semi-supervised multitask learning on multispectral satellite images using Wasserstein Generative Adversarial Networks (GANs) for predicting poverty,” *arXiv:1902.11110*, 2019.
- [17] D. Wang *et al.*, “Big trajectory data mining: A survey of methods and applications,” *Sensors*, vol. 20, no. 16, 2020.
- [18] Y. Zheng, “Trajectory data mining: An overview,” *ACM TIST*, vol. 6, no. 3, pp. 1–41, 2015.
- [19] C. D. Smith *et al.*, “Incorporating GIS and technology in response to COVID-19,” *Prev. Chronic Dis.*, vol. 17, 2020.
- [20] C. V. K. Iyer *et al.*, “Trinity: A no-code AI platform for complex spatial datasets,” *arXiv:2106.11756*, 2021.
- [21] X. Xiao *et al.*, “VAE-Info-cGAN: Generating synthetic images by combining pixel-level and feature-level geospatial conditional inputs,” *arXiv:2012.04196*, 2020.
- [22] Y. Dong *et al.*, “Scalable representation learning for heterogeneous networks,” in *ACM SIGKDD*, 2017.
- [23] A. Grover *et al.*, “node2vec: Scalable feature learning for networks,” in *ACM SIGKDD*, 2016.
- [24] B. Perozzi *et al.*, “Deepwalk: Online learning of social representations,” in *ACM SIGKDD*, 2014.
- [25] EPSG Geodetic Parameter Registry, “EPSG:3857,” <https://epsg.io/3857>.
- [26] S. Rifai *et al.*, “Contractive auto-encoders: Explicit invariance during feature extraction,” in *ICML*, 2011.
- [27] J. Yuan *et al.*, “T-Drive: Driving directions based on taxi trajectories,” in *ACM SIGSPATIAL*, 2010.
- [28] A. Grama *et al.*, *Introduction to parallel computing*. Pearson Education, 2003.
- [29] FeatureStore.org, “What is a feature store?” <https://www.featurestore.org/what-is-a-feature-store>.
- [30] L. Moreira-Matias *et al.*, “Predicting taxi-passenger demand using streaming data,” *IEEE Trans. Intelligent Transportation Sys.*, vol. 14, pp. 1393–1402, 2013.
- [31] OpenStreetMap, “Public GPS Traces,” <https://www.openstreetmap.org/traces>.
- [32] Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma, “Mining interesting locations and travel sequences from GPS trajectories,” in *ACM WWW*, 2009.
- [33] M. Panzarino, “Apple is rebuilding Maps from the ground up,” *TechCrunch*, 2018.
- [34] Wikipedia, “Well-known text,” <https://w.wiki/4Aj8>.
- [35] O. Ronneberger *et al.*, “U-Net: CNNs for biomedical image segmentation,” *arXiv:1505.04597*, 2015.
- [36] Y. Bengio *et al.*, “Representation learning: A review and new perspectives,” *arXiv:1206.5538*, 2014.
- [37] T. Baltrušaitis *et al.*, “Multimodal machine learning: A survey and taxonomy,” *IEEE PAMI*, vol. 41, no. 2, 2018.
- [38] V. Badrinarayanan *et al.*, “Segnet: A deep convolutional encoder-decoder architecture for image segmentation,” *IEEE PAMI*, vol. 39, no. 12, 2017.
- [39] Z. Liu *et al.*, “Semantic proximity search on heterogeneous graph by proximity embedding,” in *AAAI*, 2017.
- [40] S. Cao *et al.*, “Deep neural networks for learning graph representations,” in *AAAI*, 2016.
- [41] D. Wang *et al.*, “Structural deep network embedding,” in *ACM SIGKDD*, 2016.
- [42] J. Zhou *et al.*, “Graph neural networks: A review of methods and applications,” *arXiv:1812.08434*, 2018.
- [43] L. Jing *et al.*, “Self-supervised visual feature learning with deep neural networks,” *arXiv:1902.06162*, 2017.
- [44] X. Liu *et al.*, “Self-supervised learning: Generative or contrastive,” *arXiv:2006.08218*, 2020.
- [45] S. Gidaris *et al.*, “Learning representations by predicting bags of visual words,” in *CVPR*, 2020.
- [46] P. Vincent *et al.*, “Extracting and composing robust features with denoising autoencoders,” in *ICML*, 2008.
- [47] A. Crivellari *et al.*, “From motion activity to geo-embeddings: Generating and exploring vector representations of locations, traces and visitors through large-scale mobility data,” *ISPRS IJGI*, vol. 8, no. 3, p. 134, 2019.
- [48] S. Dabiri *et al.*, “Semi-supervised deep learning approach for transportation mode identification using GPS trajectory data,” *IEEE TKDE*, vol. 32, no. 5, pp. 1010–1023, 2019.
- [49] X. Jiang *et al.*, “TrajectoryNet: An embedded GPS trajectory representation for point-based classification using RNNs,” *arXiv:1705.02636*, 2017.
- [50] Q. Gao *et al.*, “Identifying human mobility via trajectory embeddings,” in *IJCAI*, 2017.
- [51] A. Solomon *et al.*, “Predict demographic information using word2vec on spatial trajectories,” in *UMAP*, 2018.
- [52] H. Cao *et al.*, “Habit2vec: Trajectory semantic embedding for living pattern recognition in population,” *IEEE Trans. Mob. Comp.*, vol. 19, no. 5, pp. 1096–1108, 2019.
- [53] Wikipedia, “Haversine formula,” <http://w.wiki/4Aj7>.
- [54] D. Stürzaker *et al.*, *Probability and random processes*. Clarendon Press, 1992.

APPENDIX

A. Distributed, data-parallel implementation of Algorithm 1

Algorithm 2 presents the pseudo-code of a distributed, data-parallel implementation of Algorithm 1 written using the Scala Dataset API in Apache Spark⁶ syntax. Apache Spark is an open-source, fault-tolerant, distributed, and data parallel cluster computing framework with support for distributed in-memory computation. Adopting an object-oriented style, Algorithm 2 uses the `Pair` and `IdxPair` case classes defined as follows. A transition between two nodes is represented by the `Pair` object that has four parameters: (i) `node1` and `node2` denote the nodes involved; (ii) `flg` denotes the type of transition from the vantage point of `node1` and is set to `flg='a'` for the transition `node1→node2` while it is set to `flg='e'` for the `node2→node1` transition; (iii) `cnt`, initialized to 1.0, that counts the number of occurrences of the transition. An `IdxPair` object also represents a transition between two nodes, $s, s' \in \overline{V}_{ES}$, and has four parameters: (i) `node` is one of `s` or `s'`; (ii) `rmIdx` is the row-major index of `s'` in Ψ_s^a if `node=s` while it is the row-major index of `s` in $\Psi_{s'}^e$ if `node=s'`; (iii) `flg` is a flag set to `'a'` if `node=s` and to `'e'` if `node=s'`; (iv) `cnt` is the frequency of the observed transition.

Line 2 of Algorithm 2 creates valid absorption transitions as `Pair` objects from T_i . Line 3 computes the corresponding emission transition for every absorption transition. All `Pair` objects representing emission and absorption transitions are converted to corresponding `IdxPair` objects in line 4. Lines 5 and 6 count the observed frequency of both absorption and

⁶<https://spark.apache.org/>

Algorithm 2 Distributed, data-parallel implementation of Algorithm 1 for generating $\Psi_s^{ea} \forall s \in \overline{V}_{ES}$

Input: $\mathcal{T}_{t_0+\Delta t} = (\mathbf{T}_1 \dots \mathbf{T}_M)$, $|\mathbf{T}_i| = n_i \forall i \in [1, M]$, $\mathbf{T}_i = (p_1^i, \dots, p_{n_i}^i)$ where $p_k^i = (z_k^i, t_k^i)$
Output: $\Psi_s^e, \Psi_s^a \forall s \in \overline{V}_{ES}$
Parameters: δ_R
Initialize: Ingest $\mathcal{T}_{t_0+\Delta t}$ as `trj`
1: **procedure** ReachabilitySummaryGenerator()
2: `result = trj.flatMap(t \Rightarrow createAPairs(t))`
3: `.flatMap(p \Rightarrow createEPairs(p))`
4: `.map(p \Rightarrow getIdxForDestinationTile(p))`
5: `.map(p \Rightarrow (p.cTile, p.dIdx, p.flg) \rightarrow p.cnt)`
6: `.reduceByKey((v1, v2) \Rightarrow v1 + v2)`
7: `.map((k, v) \Rightarrow k.cTile \rightarrow (k.dIdx, k.flg, v))`
8: `.groupByKey()`
9: `.map((k, v) \Rightarrow Create $\Psi_k^e, \Psi_k^a, \Psi_k^{ea}$ for tile k)`
10: `.saveToDistributedStorage($\Psi_k^{ea} \forall k \in \overline{V}_{ES}$)`
11: **function** createAPairs(\mathbf{T}_i)
12: **for** $k \leftarrow 1$ to n_i **do**
13: **for** $l \leftarrow k$ to n_i **do**
14: **if** $z_l^i \in N_{z_k^i}^{\delta_R}$ **then**
15: **return** **Pair** ($z_k^i, z_l^i, 'a', 1.0$)
16: **function** createEPairs(p: **Pair**)
17: **return** (p, **Pair** (p.node2, p.node1, 'e', p.cnt))
18: **function** getIdxForDestinationTile(p: **Pair**)
19: $s, s' \leftarrow$ p.node1, p.node2
20: `rmIdx \leftarrow $L(x_{s'} - x_s + \delta_R) + (y_{s'} - y_s + \delta_R)$`
21: **return** **IdxPair** (p.node1, rmIdx, p.flg, p.cnt)

emission transitions between all pairs of nodes in \overline{V}_{ES} from all trajectories in $\mathcal{T}(t_0, \Delta t)$. All associated transitions are gathered in line 8 to compute the absorption and emission channels which are concatenated in line 9 to obtain the reachability summaries for all nodes in \overline{V}_{ES} . Line 10 saves the output into distributed storage.

B. Modified analyzeTrajectories procedure

Algorithm 3 presents the modified version of `analyzeTrajectories` procedure when distance and time weighting, discussed in Section III-C, are used. Note that the same transition, (s, s') , can have different counts within a single trajectory if the transitions take two different possible paths during which the distance covered and time taken to cover the distance are different. Line 5 in Algorithm 3 uses the haversine formula [53] to calculate the great-circle distance between the centroids of two zoom-24 tiles. One can analogize the weight decays for nodes $s' \in N_s^{\delta_R}$ in the reachable neighborhood of node s to the soft-attention mechanism in NLP where the magnitude of the weight decay can be thought of as the ‘‘attention’’ paid by node s to the node s' in the calculation of the reachability summary, Ψ_s^{ea} .

C. T-Drive Dataset Pre-processing

To demonstrate scalability of Algorithm 1 in Section III-D, the publicly available T-Drive dataset [27] from Microsoft Research, containing location coordinates (latitude-longitude pairs) and timestamps for 7 days of 10,357 taxis in Beijing, is used. To transform these into a trajectory dataset, $\mathcal{T}(t_0, \Delta t)$, the data is pre-processed so that the daily chronological sequence of

Algorithm 3 Encode distance & time weighting in Ψ_s^{ea}

Input: $\mathbf{T}_i \in \mathcal{T}_{t_0+\Delta t}^i$, $|\mathbf{T}_i| = n_i$, $\mathbf{T}_i = (p_1^i, \dots, p_{n_i}^i)$ where $p_k^i = (z_k^i, t_k^i)$
Output: Map \mathbf{S}
Parameters: $\delta_R, \sigma_d, \sigma_t$
Initialize: $\mathbf{T}_i = (\overline{p}_1^i, \dots, \overline{p}_{n_i}^i)$ such that the 3-tuple $\overline{p}_k^i = (z_k^i, t_k^i, d_k^i = 0) \forall \mathbf{T}_i \in \mathcal{T}_{t_0+\Delta t}^i$. Let $\overline{n} = \max\{n_1, \dots, n_M\}$
1: **function** analyzeTrajectories($\mathcal{T}_{t_0+\Delta t}^i$) $\triangleright \mathcal{O}(M\overline{n}^2)$
2: Initialize map $\mathbf{S} = \emptyset$
3: **for** trajectory $\mathbf{T}_i \in \mathcal{T}_{t_0+\Delta t}^i$ **do**
4: **for** $k \leftarrow 2$ to n_i **do** $\triangleright \mathcal{O}(n_i)$
5: $\Delta d = \text{haversineDistance}(z_{(k-1)}^i, z_k^i)$
6: $d_k^i \leftarrow d_{(k-1)}^i + \Delta d$
7: **for** $k \leftarrow 1$ to n_i **do** $\triangleright \mathcal{O}(n_i^2)$
8: **for** $l \leftarrow k$ to n_i **do**
9: **if** $z_l^i \in N_{z_k^i}^{\delta_R}$ **then**
10: **if** $z_k^i \notin \mathbf{S}$, $\mathbf{S}[z_k^i] = (\mathbf{S}_{z_k^i}^e = \emptyset, \mathbf{S}_{z_k^i}^a = \emptyset)$
11: **if** $z_l^i \notin \mathbf{S}$, $\mathbf{S}[z_l^i] = (\mathbf{S}_{z_l^i}^e = \emptyset, \mathbf{S}_{z_l^i}^a = \emptyset)$
12: $r_{z_k^i}(z_l^i) \leftarrow \text{getIndex}(z_k^i, z_l^i)$
13: $r_{z_l^i}(z_k^i) \leftarrow \text{getIndex}(z_l^i, z_k^i)$
14: $\Delta_{d, z_k^i, z_l^i} \leftarrow (d_l^i - d_k^i)$
15: $\Delta_{t, z_k^i, z_l^i} \leftarrow (t_l^i - t_k^i)$
16: $\Delta \mathbf{w}(z_k^i, z_l^i; \sigma_d, \sigma_t) \leftarrow$ Section (III-C)
17: $c_{z_k^i}^a(z_l^i) = \Delta \mathbf{w}(z_k^i, z_l^i; \sigma_d, \sigma_t)$
18: $c_{z_l^i}^e(z_k^i) = \Delta \mathbf{w}(z_l^i, z_k^i; \sigma_d, \sigma_t)$
19: $\mathbf{S}_{z_k^i}^a[r_{z_k^i}(z_l^i)] += c_{z_k^i}^a(z_l^i)$
20: $\mathbf{S}_{z_l^i}^e[r_{z_l^i}(z_k^i)] += c_{z_l^i}^e(z_k^i)$
21: **return** \mathbf{S}

coordinate-timestamp pairs for each taxi is considered as one contiguous trajectory leading to $|\mathcal{T}(t_0, \Delta t)| = M = 68,851$, where $t_0 =$ February 02, 2008 and $\Delta t = 7$ days. Coordinates are converted to zoom-24 tiles to match the format required for consumption by Algorithm 1. Three datasets of 2000, 8000, and 64000 trips are randomly sampled from this processed dataset for the scaling analysis in Section III-D.

D. The Contractive Autoencoder

The neural architecture of the custom-designed contractive autoencoder, \mathcal{F}_R , is presented in Figure 6. The encoder, \mathcal{F}_R^e , has 5.2 million trainable parameters and the decoder, \mathcal{F}_R^d , has 5.8 million parameters, respectively. Using cross-validation, λ in equation (5c) is found to be 0.5. The autoencoder is implemented and trained using the Python API for TensorFlow. Distributed training of the model uses 8 NVIDIA Tesla V100 GPUs with a batch-size of 256 examples per GPU. For predicting reachability embeddings, distributed prediction on 120 NVIDIA Tesla V100 GPUs with a batch-size of 2000 examples per GPU is used.

E. Interpreting Reachability Embeddings using the Chapman-Kolmogorov Equations

The Markovian transitions in \mathbf{V}_{ES} define a *transition probability matrix* [54], $\mathbf{P}_{ES} = [\mathbb{P}_{ES}^{(s, s')}] \in [0, 1]^{|V_{ES}| \times |V_{ES}|}$,

$$\mathbf{Z} = \sum_{s'' \in V_{ES}} \mathbf{w}_{ES}^{(s, s'')} ; \mathbb{P}_{ES}^{(s, s')} = \mathbf{w}_{ES}^{(s, s')} / \mathbf{Z} \forall s, s' \in V_{ES}.$$

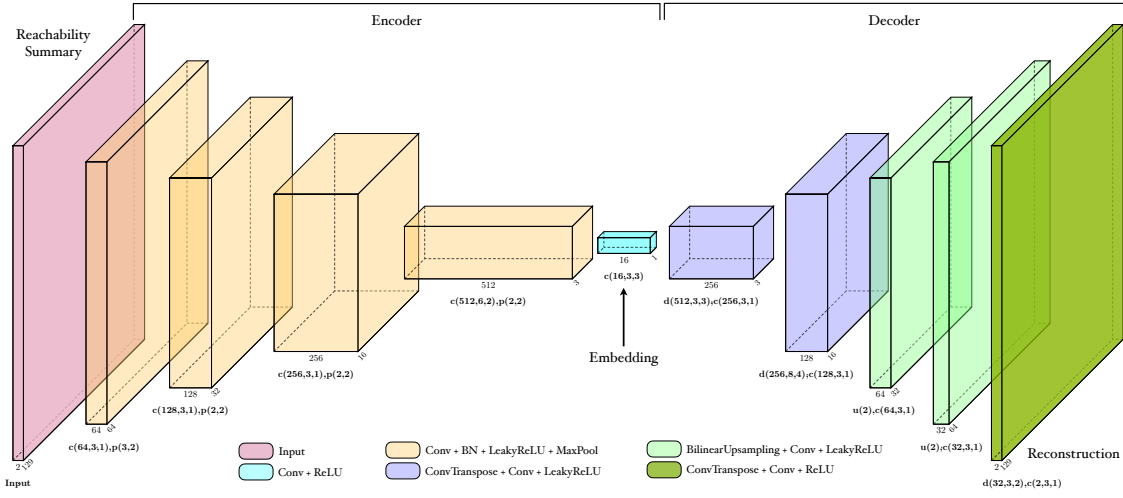


Fig. 6: Contractive autoencoder architecture. Blocks describe outputs from preceding layer detailed in block caption. Layers denoted using $c(f, k, s)$, $p(k, s)$, $d(f, k, s)$, and $u(m)$, where, c is convolution, p is max-pooling, d is transposed convolution, $u(m)$ is bilinear upsampling with multiplicity factor m , f is number of filters, k is kernel size, and s is stride.

Matrices, Ψ_s^e and Ψ_s^a , defined in equations (4a) and (4b), are closely related to \mathbf{P}_{ES} . For any square matrix, X , let $\text{sum}(X)$ denote the sum of all the matrix elements and $\text{scale}(X) = X/\text{sum}(X)$. Then, $\text{scale}(\Psi_s^e)$ and $\text{scale}(\Psi_s^a)$ are spatial proximity-preserving rearrangements in matrix form of the column and row vectors of \mathbf{P}_{ES} corresponding to node s , respectively. The Chapman-Kolmogorov equations (CKE) [54] are a fundamental identity of Markov chains and the inspiration for the construction of reachability summaries described in section III-A. While CKE is a more general result, in this paper we consider the specific case of *discrete* and *homogeneous* [54] Markov chains.

Consider the state space V_{ES} . All states are discrete and all Markovian transitions in the state space are modeled as homogeneous. For any node, s , we wish to learn a representation of s that captures its contribution to connecting any two arbitrarily chosen nodes $s', s'' \in V_{ES}$. Specifically, we wish to encode the information captured by the two-step transition probability from s' to s'' passing through s into the representation for s , for all choices of $s', s'' \in V_{ES}$ as evidenced by transitions deduced from observed trajectories in the set $\mathcal{T}(t_0, \Delta t)$. This is in contrast to algorithms like [22]–[24] which synthetically construct random walks on a graph. Setting $k = 1, n = 2$, for chosen state $s \in V_{ES}$ and $\forall s', s'' \in V_{ES}$, it follows from equation CKE that

$$\left(\mathbf{P}_{ES}^{(s', s'')}\right)^2 = \mathbf{P}_{ES}^{(s', s)} \mathbf{P}_{ES}^{(s, s'')} + \sum_{z \in V_{ES} \setminus \{s\}} \mathbf{P}_{ES}^{(s', z)} \mathbf{P}_{ES}^{(z, s'')} \quad (6)$$

where, the first term on the right-hand side, henceforth denoted as $C_s^{(s', s'')}$, measures the contribution of state s to the total two-step transition probability from s' to s'' . Note that only the subset of trajectories in $\mathcal{T}(t_0, \Delta t)$ that pass through s , denoted as $\mathcal{T}(t_0, \Delta t)(s)$, contribute to $C_s^{(s', s'')}$. Summing $C_s^{(s', s'')}$ over all possible $s', s'' \in V_{ES}$, which is the total contribution of s

to all transitions in V_{ES} , denoted as C_s , and if only transitions from $N_s^{\delta R}$ are considered valid, we obtain

$$C_s = \sum_{s', s'' \in V_{ES}} C_s^{(s', s'')} = \sum_{s' \in \Xi_s^{\delta R}} \mathbf{P}_{ES}^{(s', s)} \sum_{s'' \in \Lambda_s^{\delta R}} \mathbf{P}_{ES}^{(s, s'')}.$$

If the total number of Markovian transitions in $\mathcal{T}(t_0, \Delta t)(s)$ that start at s is ν_s^a , and that end at s is ν_s^e , we have

$$\nu_s^a \nu_s^e C_s = \sum_{s' \in \Xi_s^{\delta R}} \left(\nu_s^e \mathbf{P}_{ES}^{(s', s)}\right) \sum_{s'' \in \Lambda_s^{\delta R}} \left(\nu_s^a \mathbf{P}_{ES}^{(s, s'')}\right) = \psi_s^{e\top} \psi_s^a.$$

Here, ψ_s^e and ψ_s^a are L^2 -dimensional column vectors of the emission and absorption probabilities of s corresponding to the L^2 states in $N_s^{\delta R}$. When ψ_s^e and ψ_s^a are reshaped to spatial-relation-preserving $L \times L$ -dimensional matrices based on the relative positions of the nodes in $N_s^{\delta R}$ on G_{ES} , they yield Ψ_s^e and Ψ_s^a , respectively. Thus, Ψ_s^e and Ψ_s^a exactly capture the contribution of s to connecting nodes in $N_s^{\delta R}$ where it acts as the intermediary, thereby capturing essential spatial connectivity information. The CKE (equation 6) are reinterpreted using the lens of reachability-based emission and absorption channels by noting

$$\sum_{s', s'' \in V_{ES}} \left(\mathbf{P}_{ES}^{(s', s'')}\right)^2 = \sum_{z \in V_{ES}} \psi_z^{e\top} \psi_z^a. \quad (7)$$

⁷Analogous to the role layover airports or connection hubs play (e.g., Dubai for Emirates) in 1-stop flights. The key idea is to use the information of the set of flights flying *into* the airport and the information of the set of flights flying *out of* the airport to learn the representation of the layover airport.