

Towards a Semantic Indoor Trajectory Model

Alexandros Kontarinis
ETIS UMR 8051, University of
Paris-Seine, University of
Cergy-Pontoise, ENSEA, CNRS /
DAVID Lab, University of Versailles
Saint-Quentin, University of
Paris-Saclay
alexandros.kontarinis@ensea.fr

Karine Zeitouni
DAVID Lab, University of Versailles
Saint-Quentin, University of
Paris-Saclay
Versailles, France
karine.zeitouni@uvsq.fr

Claudia Marinica
ETIS UMR 8051, University of
Paris-Seine, University of
Cergy-Pontoise, ENSEA, CNRS
Cergy-Pontoise, France
claudia.marinica@u-cergy.fr

Dan Vodislav
ETIS UMR 8051, University of
Paris-Seine, University of
Cergy-Pontoise, ENSEA, CNRS
Cergy-Pontoise, France
dan.vodislav@u-cergy.fr

Dimitris Kotzinos
ETIS UMR 8051, University of
Paris-Seine, University of
Cergy-Pontoise, ENSEA, CNRS
Cergy-Pontoise, France
Dimitrios.Kotzinos@u-cergy.fr

ABSTRACT

In this paper we present a Semantic Indoor Trajectory Model aimed at supporting the design and implementation of context-aware mobility data mining and statistical analytics methods. Motivated by a compelling museum case study, and by what we perceive as a lack in indoor trajectory research, we are interested in combining aspects of state-of-the-art semantic outdoor trajectory models, with a semantically-enabled hierarchical symbolic representation of the indoor space, which abides by OGC's IndoorGML standard. We drive the discussion on those modeling issues and details that have been overlooked so far or where our approach deviates from typical practices. We illustrate the modeling part with instantiations from the Louvre Museum in an effort to provide a pragmatic view of what a Semantic Indoor Trajectory Model ought to represent and ideally also how.

1 INTRODUCTION

It has long been of paramount importance for museums to “know” their visitors, meaning to study and understand their motivations, expectations, engagement, and satisfaction. In this regard, multimedia guides offering Location-Based Services (e.g. way-finding, contextualized content delivery) are becoming an invaluable tool for museums, since they provide them with access to an unprecedented wealth of visitor movement data. Similar opportunities have appeared in other domains of indoor human mobility such as retail stores, arenas, hospitals, airports, universities [16].

So far, trajectory-based human mobility data analytics research has solely focused on outdoor trajectories, driven by the fact that Geographic Information Science (GIS) has traditionally only supported outdoor spatial information. This type of research differs considerably in indoor environments, mainly due to interior architectural components constraining (or otherwise affecting) the way people can move. For example, an indoor trajectory model has to consider multiple ways of entering a room, floor changes, specific locations of entrance/exit to/from the building, sensor coverage gaps and/or sensor detection area overlaps, movement

data of varying spatial granularity, and other peculiarities. In addition, indoor trajectory analytics may gain from avoiding cumbersome calculations over geometric representations of space and objects within it, that are typical of outdoor environments. Instead, operations such as intersection, containment, and proximity can be simplified in order to prioritize the non-geometric aspects of movement [15], instead of metric aspects often focused on Euclidean distances from potential targets. In fact, reasoning about space without precise quantitative information has been at the core of Qualitative Spatial Relations research [9].

Moreover, in order to reason about movement in informationally rich domains, a trajectory model must also account for multiple types of contextual and semantic information. As identified by [22] and further explored in [4, 5], there are three fundamental sets pertinent to movement, representing the “where” (set of locations), “when” (set of instants or intervals), and “what” (set of objects) of spatiotemporal data. This is true across applications. Distinguishing between semantics of time, semantics of places, and semantics of moving objects, in addition to the semantics of movement itself could empower a synergistic interplay between different types of semantics. Such semantic information can be derived either from the moving object's environment or from external data sources. It can then be used to add a meaningful dimension to “raw” trajectories. Unfortunately, semantic trajectory models have - to a large extent - targeted outdoor settings.

This has resulted in an emphasis on the enrichment of GPS data, the identification of stops and moves, the identification of transportation means, and other conceptual modeling issues that are either not interesting or not transferable in indoor settings. On the other hand, the adoption of some modeling approaches, such as the segmentation of trajectories into episodes and the use of semantic annotations, seems to be promising.

In this paper, we present a new model for spatiotemporal indoor trajectories enriched with semantic annotations. The proposed model makes use of an indoor space modeling framework, instead of assuming 2D coordinate data as is the norm. To this end, on the one hand, we identify certain limitations of state-of-the-art conceptual semantic (outdoor) trajectory models and propose ways to overcome them, and on the other hand we discuss different indoor space modeling approaches and the choices that we made. Equally important, the new model is developed in order to support mining and analysis tasks.

The rest of this paper is divided as follows: Section 2 presents an overview of the related work and its limitations with regards to indoors. Section 3 introduces our trajectory model. Section 4 introduces the Louvre case study and the corresponding model instantiation. Finally, Section 5 concludes with the key issues addressed in our model and a brief description of the types of analytical tasks that it supports.

2 RELATED WORK AND BACKGROUND

In this section, we describe the state-of-the-art in modeling indoor spaces and (outdoor) semantic trajectories.

2.1 Indoor Space Models

In order to represent movement phenomena in terms of trajectories, first a formal spatial model is needed to provide an abstraction of their physical environment. Every trajectory model (TM) proposed in the literature, either explicitly or more usually implicitly, uses a certain model of location and therefore space. In this regard, a fundamental distinction exists between quantitative and qualitative spatial representation approaches. The former are preferable when precise spatial information is important, while the latter when it is unnecessary or unavailable [9].

A qualitative spatial representation formalism, coupled with qualitative relations between spatial objects and qualitative reasoning about spatial knowledge, constitutes what is known as Qualitative Spatial Reasoning (QSR) [23]. Two of the most widespread qualitative spatial calculi are RCC (Region Connection Calculus) [10] and n-intersection [13]. In specific, RCC-8 and 4-intersection (as well as other variants) result in the definition of eight binary topological relations: “disjoint”, “touch” (“meet”), “overlap”, “contains” “insideOf”, “covers”, “coveredBy”, “equal” [14]. From a more applied perspective, most indoor spatial data models can be classified into geometric ones and symbolic ones [1]. The former focus on representing the geometry of indoor features using primitives such as points, lines, areas, and volumes. The latter focus on representing the ontological aspects of spatial units and the topological relationships between them, maintaining a more abstract view of indoor space [2]. Hybrid models represent both symbolic concepts and geometric properties.

Furthermore, a line of research works on indoor space modeling (e.g. [6]) has culminated into the development of IndoorGML [19], an OGC standard aimed at representing and allowing the exchange of geoinformation for indoor navigational systems. IndoorGML’s core module considers an indoor space as a set of non-overlapping cells that represent its smallest organizational/structural units: $S = \{c_1, c_2, \dots, c_n\}$, $c_i \cap c_j = \emptyset$. Technically, IndoorGML describes a hybrid indoor space model and not a TM, but it can be used in support of one. More specifically, the cell space and the topological relationships between its objects are represented by one or more Node-Relation Graphs (NRGs). In particular, the Poincaré duality provides the means of mapping the physical indoor space (embedded in a 2D/3D Euclidean primal space) into an adjacency NRG (in the corresponding dual space). Therefore, a cell (e.g. room) becomes a node and a cell boundary (e.g. a thin wall) becomes an edge. The respective formal terminology is summarized in Table 1. If cell boundary semantics are also taken into account (e.g. doors vs. walls, ramps) then a connectivity and/or an accessibility NRG may be derived as well. Connectivity suggests that there exists an opening in the common boundary of two cells. Accessibility additionally suggests that the opening is traversable by the moving object (MO).




Figure 1: A 2-level hierarchical graph representing the central part of the 1st floor of the Louvre’s Denon Wing.

Moreover, IndoorGML’s Multi-Layered Space Model (MLSM) is the description of multiple interpretations of the same physical indoor space, through the instantiation of multiple cell decompositions and corresponding NRGs. Each NRG is treated as a separate graph layer. Nodes belonging to different layers are connected via inter-layer “joint” edges. While intra-layer edges represent either adjacency, connectivity, or accessibility relations between non-overlapping cells, joint edges represent potential locations where a physical object might actually reside. Therefore, given that a physical object may be in only one cell of each layer at any given point in time (called the “active” state), joint edges express all the valid active state combinations (called “overall” states) and are derived by pairwise cell intersection. Equivalently, a joint edge represents any of the eight binary topological relationships derived by the n-intersection model [13], except for “disjoint” and “meet”. In Figure 1 for example, if a visitor is inside the hall represented as node 5 in layer $i + 1$, then the joint edges suggest that he can only be in either 5a, 5b, or 5c in layer i .

The MLSM can be used to represent spatial hierarchies but it is unclear how its flexible cell subdivision mechanism is ought to be used: each node may be split independently of the rest which favors ad-hoc hierarchical modeling approaches. For instance, in the Louvre example of Figure 1, we may want to split nodes 4 and 5 into smaller cells to take advantage of more precise localization data available there. It is however unclear, whether or not we should also split 1,2,3 correspondingly, or whether or not we should split 4 and 5 in the same layer (as depicted). These indoor space modeling issues have been identified in [17] and [11], but the former only provides some general partitioning criteria (e.g. splitting cells that have multiple properties or that are too big), while the latter categorizes such criteria (geometry-driven, topology-driven, semantics-driven, navigation-driven) but is more interested in furnished 3D indoor spaces, rather than 2D multi-floor spaces. However, such space modeling issues will eventually affect the spatial granularity of the symbolic TM.

N-intersection	Primal Space (2D)	Dual Space (NRG)	Dual Space (Navigation)
(spatial) region ¹	cell/"cellspace"	node	state
(region) boundary	(cell/"cellspace") boundary	(intra-layer) edge	transition
"overlap" / "coveredBy" / "inside" / "covers" / "contains" / "equal"	binary topological relationship (between cells/"cellspaces")	(inter-layer) joint edge	valid active state combination / valid overall state

Table 1: Closely related terms, often used interchangeably under the context of indoor space modeling and IndoorGML.

2.2 Semantic Trajectory Models

In the last decade, accounting for the semantics of movement has received a lot of attention in the trajectory data modeling and analytics literature. Pivotal to this has been the proposal to view a trajectory as "the user-defined record of spatiotemporal evolution of the position of a MO, during a given time interval of its lifespan, and in order to achieve a certain goal": $[t_{begin}, t_{end}] \rightarrow space$ [24]. In the same work, a purposefully generic way of semantically segmenting a trajectory into stops and moves was also established, leaving its implementation to be specified at the application level. For example, [3] adopted the conceptual TM of [24] and defined stops based on temporal stay value thresholds. Similarly, [7] adopted the conceptual TM of [24] and associated stops with important visited places, before extending it with fundamental data mining concepts in order to support frequent/sequential patterns and association rules.

More recently, in [4, 5], the authors propose a general conceptual modeling framework aimed at connecting the analysis of movement data with its spatiotemporal context, which is defined as the physical space and time where movement takes place together with the objects and events that co-exist in it. Their framework exhaustively categorizes the types of information that can be represented by movement data. First, it breaks movement down to its most essential elements: the set of locations S (space), the set of time units T (time instants or intervals), and the set of objects O (physical and abstract entities). Their elements may have properties represented as spatial, temporal, or thematic attribute values, which in turn may involve other elements of S , T , O . The framework does not address semantic modeling, apart from proposing dynamic thematic attributes, said to represent any attribute available in the movement data or "any other existing or conceivable thing", which can be thought of as the equivalent of semantic annotations used in other semantic TMs.

SeMiTri [25] is an application-independent framework for the semantic enrichment of raw GPS trajectories in the form of annotations based on spatial and temporal properties of raw data streams. The enrichment happens, either at a low level via the notion of a "semantic place" $sp_i \in P = P_{region} \cup P_{line} \cup P_{point}$, which represents a meaningful geographic object (with a Region Of Interest (ROI), a Line Of Interest (LOI), or a Point Of Interest (POI) as its extent), or at a high level via the notion of an "episode", the abstraction of a subsequence of the spatiotemporal trajectory's points that are highly correlated with respect to some identifiable spatiotemporal feature (e.g. velocity, time interval).

The conceptual semantic TM proposed by [21] is similarly structured as a sequence of potentially annotated timestamped coordinate positions or episodes. An annotation is defined as any additional data (captured or inferred) that enrich the knowledge about a trajectory or any part thereof. It can be an attribute value, a link to an object, or a complex value composed of both. Also, an "episode" is defined verbatim from [25] as "a maximal subsequence of a semantic trajectory, such that all its spatiotemporal positions comply with a given predicate, bearing on the spatiotemporal

characteristics of the positions". Lastly, temporal gaps in the movement track greater than the sampling rate of raw data, are said to be either accidental ("holes") or intentional ("semantic gaps"), in which case their list makes part of the main TM.

Finally, CONSTANT [8] is a conceptual semantic TM that resembles the TM of [21], but supports more strictly defined types of trajectory semantics. A trajectory T is defined as an ordered list of timestamped (x, y) coordinate points. Enriched with contextual information, a semantic trajectory is defined as the tuple (tid, oid, S, g, d) , where oid is the MO identifier, S is a list of semantic subtrajectories, g is the general goal of the trajectory (i.e. the reason/objective of the movement), and d is the device that generated the trajectory. g is required and S must contain at least one semantic subtrajectory, which means that a semantic trajectory must have exactly one goal and at least one meaningful part. Moreover, a semantic subtrajectory $s \subset T$ is defined as a list of consecutive semantic points, that corresponds to at least a goal, or a means of transportation, or a behavior, if not to multiple ones. Lastly, a semantic point $p \subset s$ is defined as a coordinate point, annotated with a set of environments related to where it was collected and/or with a set of places where it is located.

More generally, in the earlier semantic TM literature, semantics were largely exhausted in the names and types of the geographic places of interest related to the MO's physical stops. Efforts have since been undertaken to integrate movement ontologies, linked open data, information extracted from social network platforms, or complementary case-specific datasets, with spatiotemporal trajectory data. But they have largely concerned outdoor contexts, as made evident by the terminology (e.g. "traveling objects" [24]) and definitions introduced. On the contrary, a Semantic Indoor Trajectory Model (SITM) needs to at least consider the building's topology and space semantics. The interior of buildings is typically divided into clearly delimited spatial entities such as rooms, halls, floors. Thus, its physical segmentation already holds a considerable amount of semantic information.

3 SEMANTIC INDOOR TRAJECTORY MODEL

In this section, we define a semantic indoor trajectory model (SITM) aimed at supporting:

- all types of indoor settings;
- both human and inanimate moving objects (from hereon both referred to as MOs);
- mining and analysis applications using both statistical and reasoning approaches in order to provide insight both at the individual and collective level.

3.1 Model Components

The proposed SITM mainly consists of a semantically enriched sequence of an individual MO's spatiotemporal presence, but also makes use of a semantically enriched representation of indoor space.

The semantically enriched representation of indoor space that we propose is represented as a layered multigraph. Its nodes symbolically represent indoor spatial regions, and its edges represent topological relationship information between those regions. Static semantic information about the regions is represented through node classes and attributes as well as node-edge grouping into layers. The proposed representation is compatible with OGC’s IndoorGML standard and can be viewed as an extension of it. It is described in Subsection 3.2.

The semantically enriched representation of an individual MO’s trajectory that we propose is a couple consisting of a trace of consecutive presence intervals inside the indoor regions represented by the graph’s nodes, and a set of semantic annotations describing the trajectory in its entirety. It is semantically enriched and uses the above indoor space representation. It is described in Subsection 3.3.

3.2 Indoor Space Modeling

Based on the modeling framework provided by the IndoorGML standard and in particular its Multi-Layered Space Model (MLSM), we represent a 2D multiple floor (i.e 2.5D) indoor space as a layered multigraph $G = (V, E)$ where

$$V = \bigcup_{i=0}^m V_i$$

and

$$E = \bigcup_{i=0}^m E_i^{acc} \cup E^{top}$$

G comprises $m + 1$ different layers of nodes and edges, each constituting an accessibility Node-Relation Graph (NRG):

$$G_i = (V_i, E_i^{acc}) (0 \leq i \leq m)$$

that corresponds to a different decomposition of the indoor space. On the one hand, node $v \in V_i$ represents a cell belonging to the i -th layer and an edge $e \in E_i^{acc} \subseteq V_i \times V_i$ represents the accessibility between two cells of the i -th layer. On the other hand, a joint edge $e' \in E^{top} \subseteq V_i \times V_j$ represents a binary topological relationship between two cells of different layers ($i \neq j$). Figure 2 illustrates an example of such an indoor space graph representation consisting of five hierarchical layers (detailed in Section 4), but in general G need not be strictly hierarchical.

In the proposed indoor space model, we adopt IndoorGML’s implicit assumption that each node belongs to only one layer: $\bigcap_{i=0}^m V_i = \emptyset$. If a node is relevant to multiple layers then it is essentially replicated in each one and all the copies are connected to each other via “equal” joint edges. Moreover, assuming that cells represent the physical reality of planar space (instead of a conceptual space) and that same-layer cells do not overlap at all, an intra-layer edge $e \in E_i^{acc}$ actually presupposes the “meet” relation between its two cells, because they need to share a common surface for the MO to be able to physically transition between them. At the same time, as explained in Section 2, a joint edge $e' \in E^{top}$ signifies that either one of the “overlap”, “contains”, “insideOf”, “covers”, “coveredBy”, or “equal” topological relations holds between the two cells that it connects. Thus, intra-layer edges and inter-layer edges are always of a different type, and therefore G can be considered as an edge-coloured multigraph which can be mapped to a multilayer network [18].

An important modeling decision is whether G is directed or not. Although IndoorGML does not explicitly assume either case,

it considers undirected edges in all of its examples. As far as intra-layer edges go, we can think of “adjacency” and “connectivity” as being symmetric relations. However, “accessibility” is not symmetric since often indoor movement is only unidirectionally possible due to technical, safety or other limitations. In Figure 1 for example, room 4 (“Salle des Etats”) houses the “Mona Lisa” and accommodates a vast number of visitors on a daily basis. To facilitate their flow, entering it from room 2 is often prohibited by the museum personnel while exiting it that way is allowed. Therefore, we assume directed accessibility NRGs. As far as joint edges go, while “overlap” and “equal” can be thought of as symmetric binary relations, “contains” and “covers” can not. Therefore, we also assume directed joint edges (Figure 2). If we wanted to simply model intersection non-emptiness, instead of the specific nature of the relation, then undirected joint edges would suffice.

In our model, we define a layer hierarchy as $k \geq 2$ ordered layers G_i ($0 \leq i \leq k$) of G that are only consecutively connected by joint edges. Similar to [17], we exclude “overlap” relations from layer hierarchies, but contrary to it, we also exclude “equal” relations to prohibit node repetition and instead favor a proper hierarchy. Instead of [17]’s “inside” and “coveredBy”, we assume “contains”, “covers”, and a corresponding top to bottom joint edge direction. Furthermore, we account for the fact that virtually any indoor environment is characterized by a basic three-layer hierarchy consisting of: a “Building” layer, a “Floor” layer, and a “Room” layer. The latter is loosely named as it may actually contain any type of room-level navigable spatial cell, such as rooms, chambers, halls, lobbies, cellars, terraces, corridors, hallways, big staircases. Therefore, G includes k layers representing static hierarchical levels of spatiosemantic granularity ($3 \leq k \leq m$). Other layers are optional and may also integrate with this core layer hierarchy.

It is thus evident that there can be layer hierarchies that comprise either topographic layers, or semantic layers, or both. Our core hierarchy is indeed a mixed one. The “Building” and “Floor” layers are spatially defined, since the architectural structure alone is mostly enough to determine which space constitutes a building and which space constitutes a floor. The “Room” layer is predicated both spatially and semantically: it should not contain cells of vastly different sizes, but it may contain cells whose boundaries are not necessarily physical (e.g. functionally independent subspaces of a big hall or of a great room).

Additionally, two optional layers are proposed for typical cases: a “Building Complex” root layer and a “Region of Interest (RoI)” leaf layer (Figure 2). We define the “Building Complex” layer to represent the indoor space of a site comprised of multiple buildings, such as a hospital spanning multiple attached wings or a university campus spanning multiple independent edifices. We define the “RoI” layer to represent navigable sub-room level spatial cells of application-specific interest, such as “you are here” map installations in a shopping mall or individual exhibit displays in a museum (Figure 4). The “Building Complex” and “RoI” layers are only relevant per case. When present, they can be properly integrated into the aforedescribed core layer hierarchy: “Building Complex” \rightarrow “Building” \rightarrow “Floor” \rightarrow “Room” \rightarrow “RoI”. Then, a “Floor” object describes a single building’s floor level (e.g. FloorA1 \neq FloorB1 in Figure 2). Ad-hoc refinements of the hierarchy are still possible in extremely particular cases (e.g. architectures with indistinguishable floor levels) as long as joint edges represent “contain” or “cover” relations and do not skip layers.

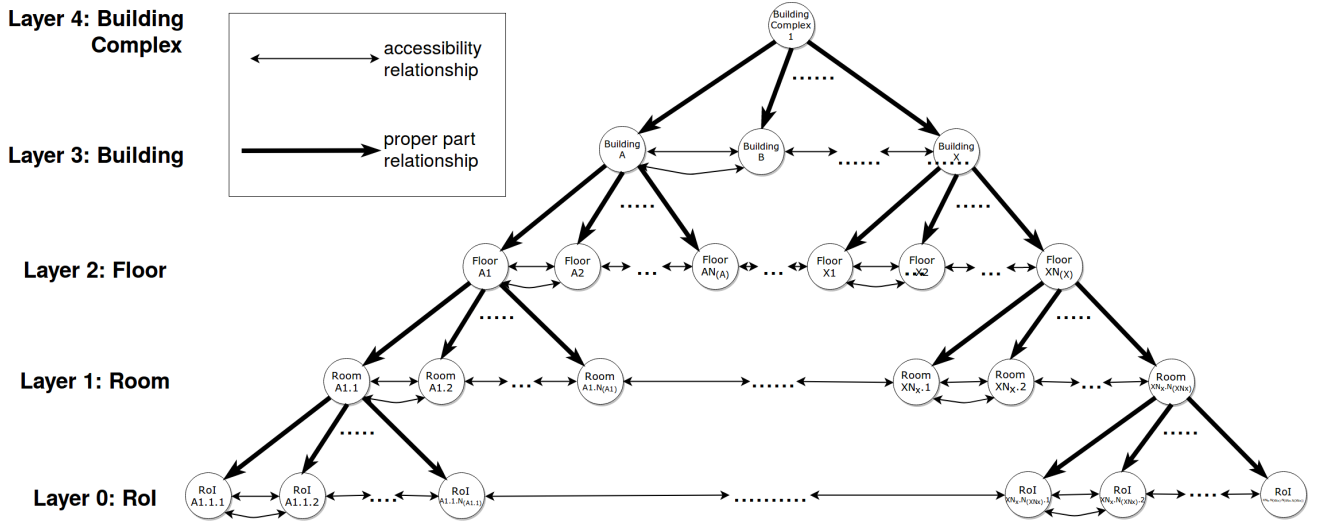


Figure 2: The required core layer hierarchy extended with a multi-building root layer and an intra-room region layer.

A static predefined layer hierarchy (e.g. Figure 2), as opposed to local ad-hoc node subdivision (e.g. Figure 1), allows a structured reasoning about the trajectories at multiple levels of granularity. By only allowing “proper part” types of relationships, we allow inference of a MO’s location at all levels of granularity above the detection data level. This in turn allows developing reasoning mechanisms to cope with missing or uncertain location information. It also enables the identification of certain types of movement patterns at the “room” level for instance, and at the same time of other types of patterns at the “floor” level, from the same trajectory dataset. Finally, hierarchies simplify the conceptual indoor space data model thanks to the transitivity of parthood (isomorphic to set inclusion) in classical mereology: a layer hierarchy only needs to connect to other layers or layer hierarchies at the lowest possible level, since a relation (e.g. “overlap”) between two nodes will also hold between their predecessors.

3.3 Semantic Indoor Trajectory Modeling

Automatically collected raw movement data typically consist of spatiotemporal records, out of which individual trajectories can be extracted. Depending on the application and on the type of MO, only the evolution of its representative location may be relevant (e.g. museum visit analysis) or perhaps also its shape and parts’ movements (e.g. sports performance analysis). In the former case, a trajectory is typically represented as a sequence of timestamped spatial points, as explained in Subsection 3.2. Due to a building’s clearly separated spaces, we consider regions (instead of points) as our primary primitive spatial entities, in the spirit of Qualitative Spatial Representation [10] and IndoorGML’s cellular space [19], both described in Section 2.

Definition 3.1 (semantic trajectory). A semantic trajectory is defined as the couple of its spatiotemporal $trace$ and the set A_{traj} of semantic annotations describing it in its entirety, as given by the following equation:

$$T_{ID_{mo}, t_{start}, t_{end}} = (trace_{ID_{mo}, t_{start}, t_{end}}, A_{traj})$$

where ID_{mo} is the identifier of the MO concerned, t_{start} and t_{end} are the trajectory’s starting and ending timestamps. Moreover, $trace_{ID_{mo}, t_{start}, t_{end}}$ represents the spatiotemporal aspect

of the trajectory defined as a sequence of timestamped semantically annotated presence periods/intervals at states of the indoor space graph G .

The second element of the couple in Def. 3.1 is a non-empty set of semantic annotations characterizing the trajectory in its entirety. A trajectory semantic annotation $a_{traj} \in A_{traj}$ is not confined within specific types of information, but would typically be chosen to represent an activity, a behavior, or a goal showcased by the complete trajectory. These terms are often ambiguously used in trajectory literature. Here, we consider an “activity” to concern more targeted/conscious actions than a “behavior”, which concerns less intentional actions or reactions. Both describe the actuality of movement. A “goal” might instead concern the potentiality of movement (e.g. a disrupted activity).

Definition 3.2 (semantic trajectory trace). A semantic trajectory trace is defined as following:

$$trace_{ID_{mo}, t_{start}, t_{end}} = (e_i, v_i, t_i^{start}, t_i^{end}, A_i)_{i \in [1, n]}$$

where $e_i = (v_{i-1}, v_i) \in \bigcup_{i=0}^m E_i^{acc}$ is the transition, i.e. boundary crossed, that led the MO from state v_{i-1} to state v_i at time t_i^{start} , where it stayed until time t_i^{end} . Moreover, A_i is a potentially empty set of semantic annotations describing that specific stay. Given that each layer’s NRG is a multigraph, it is generally useful to know the specific transition e_i (e.g. which door, staircase, or elevator was used), albeit optional².

For example, the spatiotemporal trace of a museum visitor’s 3-hour visit (on a given day) might resemble the following:

$$trace_{ID_{vis}, 11:30:00, 14:28:00} = \{ (_, room001, 11:30:00, 11:32:35, \emptyset), (door012, hall003, 11:32:31, 11:40:00, \emptyset), \dots (door005, room006, 14:12:00, 14:28:00, \emptyset) \}$$

We define a semantic subtrajectory as being for all practical purposes a semantic trajectory (similar to how a mathematical subsequence is itself a sequence) but necessarily referable to some other main semantic trajectory:

²For applications where individual transitions bear a dynamic semantic load (e.g. setting off an alarm with some probability), we can extend the TM with semantic transition annotations, effectively substituting e_i with $e_i^{sem} = (e_i, A_i^{trans})$.

Definition 3.3 (semantic subtrajectory). Given a semantic trajectory

$$T_{ID_{mo}, t_{start}, t_{end}} = (trace_{ID_{mo}, t_{start}, t_{end}}, A_{traj})$$

a semantic subtrajectory of it is defined as:

$$T'_{ID_{mo}, t'_{start}, t'_{end}} = (trace'_{ID_{mo}, t'_{start}, t'_{end}}, A'_{traj})$$

iff $trace'$ is a proper subsequence of $trace$:

$$t_{start} \leq t'_{start} < t'_{end} < t_{end} \text{ or } t_{start} < t'_{start} < t'_{end} \leq t_{end}.$$

A subtrajectory's set of semantic annotations A'_{traj} may or may not be the same as that of its main trajectory A_{traj} , contrary to [8], where a subtrajectory is enriched with different types of semantic information than its main trajectory.

Moreover, in the following, we define an episode of a semantic trajectory as any particularly meaningful part of it.

Definition 3.4 (episode). Given a semantic trajectory

$$T_{ID_{mo}, t_{start}, t_{end}} = (trace_{ID_{mo}, t_{start}, t_{end}}, A_{traj})$$

an episode of it is defined as

$$T'_{ID_{mo}, t'_{start}, t'_{end}} = (trace'_{ID_{mo}, t'_{start}, t'_{end}}, A'_{traj})$$

iff (1) $T'_{ID_{mo}, t'_{start}, t'_{end}}$ is a semantic subtrajectory of $T_{ID_{mo}, t_{start}, t_{end}}$,

(2) $A'_{traj} \neq A_{traj}$, and (3) it satisfies a given spatiotemporal and/or semantic predicate:

$$P_{ep} : T'_{ID_{mo}, t'_{start}, t'_{end}} \rightarrow \{true, false\}$$

where P_{ep} is domain-dependent and user-defined.

Moreover, an episodic segmentation of a semantic trajectory is simply any subset of its episodes that covers it time-wise. Contrary to typical literature practice, we allow an episodic segmentation to contain episodes that overlap in time, since the exact same movement part may have multiple meanings depending on the broader context. An example illustrative of the museum domain is given in the next Section.

Finally, the SITM is event-based in the sense that, only a change of the spatial cell that the MO is located in, or a change of the semantic information regarding the MO's presence in that cell, needs to be accompanied by a new tuple and a corresponding timestamp. Hence, in the previous museum visit example the last presence interval could be split if the visitor changes his goal while in *room006* (which hosts both exhibits and the gift shop): (*door005, room006, 14:12:00, 14:21:45, goals: ["visit"]*) and (*_room006, 14:21:46, 14:28:00, goals: ["visit", "buy"]*). This modeling approach allows us to integrate different data sources in order to semantically enrich the trajectory.

4 THE LOUVRE CASE STUDY

In this section, we present a compelling trajectory dataset from the world's most frequented museum, the Louvre Museum.

4.1 Visitor Movement Dataset

In July 2016, the Louvre launched its official "My Visit to the Louvre" smartphone application, which takes advantage of a large Bluetooth Low Energy (BLE) beacon infrastructure³ and the smartphone's accelerometer and compass, in order to estimate the visitor's (lat, long) coordinate position within the museum. This is accomplished via BLE Received Signal Strength Indicator (RSSI)-based trilateration, extended Kalman and particle filtering techniques. The app visualizes the position over a locally stored

³Around 1800 beacons installed across all five floors.

version of the museum map for navigation purposes. The Louvre has already been the object of visitor mobility research in the past leading to interesting conclusions [27], but the current beacon infrastructure offers improved tracking coverage and continuity.

In the obtained dataset, raw geometric positions have already been spatially aggregated into 52 non-overlapping zones. Each zone corresponds to a large polygonal area of the museum (Figures 3 and 5) specified by the museum administration in such a way so as to reflect a single exhibition theme (e.g. Italian paintings) but also only extend within a single floor. In more detail, our dataset consists of 4,945 visits (continuously collected from 19-01-2017 to 29-05-2017, where each visit consists of a sequence of timestamped "zone detections", i.e. detections of the visitor's smartphone inside a certain zone. The duration of a visit ranges from 0 sec (potential error) to 7 hours, 41 min and 37 sec, whereas the duration of a zone detection ranges from 0 sec (potential error) to 5 hours, 39 min and 20 sec. The visits were performed by 3228 different visitors using both the iPhone and Android app versions. Out of them, 1227 were "returning" visitors who made 1717 second/third visits, although not necessarily on different days. The dataset includes 20,245 zone detections and 15,300 (intra-visit) zone transitions in total.

Unfortunately, the trajectories obtained from the dataset are sparse, since a visitor may delay launching the app or stop using it early in the visit for a variety of reasons, ranging from battery depletion to lack of engagement or sporadic navigation-only usage. Moreover, around 10% of the zone detections have a duration of zero value, forcing us to filter them out as detection errors.

4.2 Model Instantiation

In order to instantiate the STIM presented in Section 3 for the Louvre case, we need first to represent Louvre's indoor spaces according to the proposed graph-based model. This is done in Figure 2. Although the Louvre's multi-layered graph is prohibitively large to be included in this paper, we cite hereafter its correspondences with respect to Figures 3 and 5: Layer 4 is instantiated as the whole "Louvre Museum", Layer 3 as its three wings ("Richelieu", "Denon", and "Sully") as well as the "Napoleon" area (under the Pyramide), Layer 2 as a wing's five different floors (-2, -1, 0, +1, +2), Layer 1 as a floor's rooms and halls (hundreds in total), and Layer 0 as a room's exhibits (several hundreds of the most important ones). In addition, we add a semantic layer that happens to fall right between Layer 2 and Layer 1, representing the thematic zones of our dataset as described in Subsection 4.1 (Figure 3). Layer 4 actually represents a level above any specific building, denoting whether a visitor is at the Louvre in general. Layer 3 treats each wing of the museum as a separate building because its spaces and usage are practically equivalent to that of a typical building. In Layer 0, we opted to define a RoI as the predefined spatial area of engagement with the corresponding exhibit, outside of which a visitor is certainly not paying attention to it. For simplicity, a RoI includes the area physically taken up by the exhibit itself and its display installation (i.e. no holes).

Finally, an interesting space modeling decision concerns whether or not to assume that the spatial region represented by a node in layer $i + 1$ is fully covered by the union of the spatial regions represented by its child nodes in layer i . For example, is a floor fully covered by the rooms it contains (Figure 2)? Although not explicitly stated, the IndoorGML standard and related works (e.g. [17]) seem to adhere to a full-coverage hypothesis. This has the advantage that accessibility relations need only be captured at




Figure 3: Choropleth map of visitor detections in the Louvre’s 11 ground floor polygonal zones.

the lowest possible level of the hierarchy, from where they can be inferred for the higher levels. However, it is often an unrealistic assumption [20]. In Figure 4 for instance, the RoIs of the displayed exhibits do not completely cover their room’s surface.

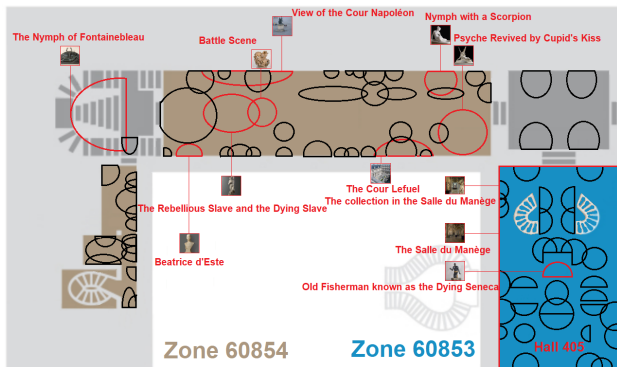


Figure 4: Indicative representation of the RoIs contained within zones 60854 and 60853 of the Louvre.

Having instantiated the Louvre’s indoor space representation, the SITM is used to extract (from the zone detection data) the Louvre visit trajectories as sequences of presence intervals in the museum’s thematic zones. Figure 6 depicts the accessibility topology of the 30 zones present in the dataset, which was extracted by hand on site and can therefore also assist in filtering out data errors. The figure’s lower part corresponds to the -2 floor of the museum, and a short sub-visit of a random visitor in February 2017 is drawn over it: at time t_1 the visitor was detected in Zone60887 (i.e. E in Figure 5) for a duration of δt_1 , and at time t_2 he was detected in Zone60890 (i.e. S in Figure 5) for a duration of δt_2 . From the zone layer NRG (Figure 6) we can infer that although never detected there, the visitor must have passed from Zone60888 (i.e. P in Figure 5). In our SITM, this would be captured with the addition of an extra tuple in the sequence, e.g.: (*checkpoint002, zone60888, 17:30:21, 17:31:42, {goals:[“cloakroomPickup”,“souvenirBuy”,“museumExit”]}*)

The semantics of places also offer us valuable insight about the visitor’s trajectory. For instance, we know that the visitor disappearing after Zone60890 is normal because it is one of the Louvre’s exit zones (through the Carrousel Hall). Also, Zone60887




Figure 5: A Louvre visit trajectory may contain two overlapping “exit museum” and “buy souvenir” episodes.

hosts the temporary exhibition of the Louvre which requires a separate ticket to enter. Thus, we would expect that $\delta t_1 \gg \delta t_2$. There are many such interesting examples, where cell semantics could help, not only explain the results of, but potentially even redesign, existing sequential pattern mining methods.

It is now more apparent why our SITM allows for overlapping episodes instead of requiring mutually exclusive episode predicates (as in [26] for example). For instance, if a given visitor (Figure 5) has visited the temporary exhibition (hosted in E) and wishes to leave the museum, he may take the path $E \rightarrow P \rightarrow S \rightarrow C$ before his trace disappears, as he is leaving the museum through the Carousel exit (C). However, he may also want to first buy something from the souvenir shops (hosted in S). Hence, when considering a goal-related episodic segmentation of his trajectory, we may tag the whole $E \rightarrow P \rightarrow S \rightarrow C$ part with the “exit museum” goal and its $E \rightarrow P \rightarrow S$ subsequence with the “buy souvenir” tag. More generally, any part of the MO’s trajectory may correspond to multiple episodes (goal-related or otherwise).

5 CONCLUSIONS AND FUTURE WORK

In this work, we presented an indoor space representation based on the IndoorGML standard [19] and using a hierarchical graph structure similar to [17]. The main difference is that we require a static hierarchy of three basic layers (building, floor, room) and propose two more typical layers (building complex, intra-room region of interest), thus avoiding ad-hoc subdivisions of space. Motivated by our case study involving a museum visitor mobility dataset, containing spatially aggregated timestamped detections, we instantiated the space representation, also adding a case-specific semantic layer of “thematic zones” that matches the granularity of our data. We also explained how a sequence




Figure 6: Based on the chain topology of zones, a visitor's presence in Zone 60888 (blue zone) can be inferred.

of presence intervals in symbolic indoor areas, coupled with semantic annotations, and flexible concept definitions, can produce a Semantic Indoor Trajectory Model (SITM) that adopts good practices from state-of-the-art semantic outdoor TMs.

We will next focus on developing new data mining methods that exploit the expressiveness of the SITM, and on proposing semantic similarity metrics for trajectories (e.g. for visitor profiling). In the future, it would be interesting to integrate the indoor space representation with formal ontologies of cultural heritage information (e.g. CIDOC Conceptual Reference Model [12]). Also, modeling conceptual instead of physical trajectories could be compelling in the museum domain, where an interpretation of visitor movement based on “focus of attention” is sometimes even more important than one based on physical presence. With regards to the Louvre case, it would be of interest to account for the problem of data sparsity by restructuring longer indicative visits from the actual fragmented zone sequences. However, the data can already provide some interesting insight albeit at a coarse level of granularity (e.g. floor-switching patterns).

ACKNOWLEDGMENTS

The authors would like to thank Anne Krebs for her cooperativity and her multifaceted help and Artus Gosselin for his contribution in developing the code for Figure 6.

This work is supported by the TRAJECTOIRES project funded by the French Heritage Science Foundation (EUR-17-EURE-0021).

Karine Zeitouni's work in this paper has been supported by the MASTER project that has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie-Słodowska Curie grant agreement N. 777695.

REFERENCES

- [1] Imad Afyouni, Cyril Ray, and Claramunt Christophe. 2012. Spatial models for context-aware indoor navigation systems: A survey. *Journal of Spatial Information Science* 1, 4 (May 2012), 85–123.
- [2] Imad Afyouni, Cyril Ray, and Christophe Claramunt. 2017. *Representation: Indoor Spaces*. American Cancer Society, 1–12.
- [3] Luis Otavio Alvares, Vania Bogorny, Bart Kuijpers, Jose Antonio Fernandes de Macedo, Bart Moelans, and Alejandro Vaisman. 2007. A Model for Enriching Trajectories with Semantic Geographical Information. In *Proceedings of*

- the 15th Annual ACM International Symposium on Advances in Geographic Information Systems (GIS '07)*. ACM, New York, NY, USA, 22:1–22:8.
- [4] Gennady Andrienko, Natalia Andrienko, Peter Bak, Daniel Keim, Slava Kisilevich, and Stefan Wrobel. 2011. A conceptual framework and taxonomy of techniques for analyzing movement. *Journal of Visual Languages & Computing* 22, 3 (2011), 213 – 232.
- [5] Gennady Andrienko, Natalia Andrienko, and Marco Heurich. 2011. An Event-based Conceptual Model for Context-aware Movement Analysis. *International Journal of Geographical Information Science* 25 (2011), 1347–1370.
- [6] Thomas Becker, Claus Nagel, and Thomas H. Kolbe. 2009. Supporting Contexts for Indoor Navigation Using a Multilayered Space Model. In *2009 Tenth International Conference on Mobile Data Management: Systems, Services and Middleware*. 680–685.
- [7] Vania Bogorny, Carlos Alberto Heuser, and Luis Otavio Alvares. 2010. A Conceptual Data Model for Trajectory Data Mining. In *Proceedings of the 6th International Conference on Geographic Information Science (GIScience'10) (Lecture Notes in Computer Science)*. Springer-Verlag, Berlin, Heidelberg, 1–15.
- [8] Vania Bogorny, Chiara Renso, Artur Ribeiro de Aquino, Fernando de Lucca Siqueira, and Luis Otavio Alvares. 2014. CONSTAnT – A Conceptual Data Model for Semantic Trajectories of Moving Objects. *Transactions in GIS* 18, 1 (2014), 66–88.
- [9] Juan Chen, Anthony G. Cohn, Dayou Liu, Shengsheng Wang, Jihong Ouyang, and Qiangyuan Yu. 2015. A survey of qualitative spatial representations. *The Knowledge Engineering Review* 30, 1 (2015), 106–136.
- [10] Anthony Cohn, Brandon Bennett, John Gooday, and Mark Gotts. 1997. Qualitative Spatial Representation and Reasoning with the Region Connection Calculus. *Geoinformatica* 1 (1997), 275–316.
- [11] Abdoulaye Diakité, Sisi Zlatanova, and Ki-Joune Li. 2017. About the subdivision of indoor spaces in IndoorGML. In *12th 3D Geoinfo Conference*. 41–48.
- [12] Martin Doerr, Christian-Emil Ore, and Stephen Stead. 2007. The CIDOC Conceptual Reference Model: A New Standard for Knowledge Sharing. In *Tutorials, Posters, Panels and Industrial Contributions at the 26th International Conference on Conceptual Modeling (ER '07)*. Australian Computer Society, Inc., 51–56.
- [13] Max Egenhofer and John Herring. 1992. Categorizing binary topological relations between regions, lines and points in geographic databases. *The 9-Intersection, Formalism and Its Use For Natural-Language Spatial Predicates* 94 (1992), 1–28.
- [14] Max J. Egenhofer and Franzosa Robert D. 1991. Point-set topological spatial relations. *International Journal of Geographical Information Systems* 5, 2 (1991), 161–174.
- [15] Stephen C. Hirtle and John Jonides. 1985. Evidence of hierarchies in cognitive maps. *Memory & Cognition* 13, 3 (1985), 208–217.
- [16] Christian Søndergaard Jensen, Hua Lu, and Bin Yang. 2010. Indoor - A New Data Management Frontier. *IEEE Computer Society Data Engineering Bulletin* 33, 2 (2010), 12–17.
- [17] Hae-Kyong Kang and Ki-Joune Li. 2017. A Standard Indoor Spatial Data Model - OGC IndoorGML and Implementation Approaches. *ISPRS International Journal of Geo-Information* 6, 4 (2017), 1–25.
- [18] Mikko Kivelä, Alex Arenas, Marc Barthelemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. 2014. Multilayer networks. *Journal of Complex Networks* 2, 3 (2014), 203–271.
- [19] Jiyeong Lee, Ki-Joune Li, Sisi Zlatanova, Thomas H. Kolbe, Claus Nagel, and Thomas Becker. 2018. *OGC IndoorGML*. Technical Report. Open Geospatial Consortium.
- [20] Hua Lu, Chenjuan Guo, Bin Yang, and Christian S. Jensen. 2016. Finding Frequently Visited Indoor POIs Using Symbolic Indoor Tracking Data. In *Proceedings of the 19th International Conference on Extending Database Technology (EDBT2016)*. 449–460.
- [21] Christine Parent, Stefano Spaccapietra, Chiara Renso, Gennady Andrienko, Natalia Andrienko, Vania Bogorny, Maria Luisa Damiani, Aris Gkoulalas-Divanis, Jose Macedo, Nikos Pelekis, Yannis Theodoridis, and Zhixian Yan. 2013. Semantic Trajectories Modeling and Analysis. *ACM Comput. Surv.* 45, 4 (2013), 42:1–42:32.
- [22] Donna J. Pequet. 2002. *Representations of Space and Time*. Guilford Press.
- [23] Jochen Renz. 2002. *Qualitative Spatial Reasoning with Topological Information*. Springer-Verlag, Berlin, Heidelberg.
- [24] Stefano Spaccapietra, Christine Parent, Maria Luisa Damiani, Jose Antonio de Macedo, Fabio Porto, and Christelle Vangenot. 2008. A Conceptual View on Trajectories. *Data and Knowledge Engineering* 65, 1 (2008), 126–146.
- [25] Zhixian Yan, Dipanjan Chakraborty, Christine Parent, Stefano Spaccapietra, and Karl Aberer. 2011. SeMiTri: A Framework for Semantic Annotation of Heterogeneous Trajectories. In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT/ICDT '11)*. ACM, New York, USA, 259–270.
- [26] Zhixian Yan, Christine Parent, Stefano Spaccapietra, and Dipanjan Chakraborty. 2010. A Hybrid Model and Computing Platform for Spatio-semantic Trajectories. In *The Semantic Web: Research and Applications (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 60–75.
- [27] Yuji Yoshimura, Anne Krebs, and Carlo Ratti. 2017. Noninvasive Bluetooth Monitoring of Visitors' Length of Stay at the Louvre. *IEEE Pervasive Computing* 16, 2 (2017), 26–34.