

RSSi-Based Visitor Tracking in Museums via Cascaded AI Classifiers and Coloured Graph Representations

Elia Onofri¹ · Alessandro Corbetta²

¹ Department of Mathematics, Roma Tre University, Rome, Italy

E-mail: eonofri@uniroma3.it

² Department of Applied Physics, Eindhoven University of Technology, Eindhoven, The Netherlands

E-mail: a.corbetta@tue.nl

Received: 23 August 2021 / Last revision received: 9 November 2021 / Accepted: 11 November 2021
DOI: [10.17815/CD.2021.131](https://doi.org/10.17815/CD.2021.131)

Abstract Individual tracking of museum visitors based on portable radio beacons, an asset for behavioural analyses and comfort/performance improvements, is seeing increasing diffusion. Conceptually, this approach enables room-level localisation based on a network of small antennas (thus, without invasive modification of the existent structures). The antennas measure the intensity (RSSi) of self-advertising signals broadcasted by beacons individually assigned to the visitors. The signal intensity provides a proxy for the distance to the antennas and thus indicative positioning. However, RSSi signals are well-known to be noisy, even in ideal conditions (high antenna density, absence of obstacles, absence of crowd, etc.).

In this contribution, we present a method to perform accurate RSSi-based visitor tracking when the density of antennas is relatively low, e.g. due to technical constraints imposed by historic buildings. We combine an ensemble of “simple” localisers, trained based on ground-truth, with an encoding of the museum topology in terms of a total-coloured graph. This turns the localisation problem into a cascade process, from large to small scales, in space and in time. Our use case is visitors tracking in Galleria Borghese, Rome (Italy), for which our method manages > 96% localisation accuracy, significantly improving on our previous work (J. Comput. Sci. 101357, 2021).

Keywords RSSi-based tracking · total-coloured graph analysis · pedestrian dynamics in museums · IoT · machine learning

1. Introduction

The behavioural analysis of museums' visitors has a long-standing multidisciplinary tradition [1,2], and underlies the capacity of profiling exhibitions, increase visitors' comfort and safety, improve public reception, increase the number of sold tickets, and enhance artworks preservation [3]. For instance, understanding visitors' dynamics unlocks focused tuning of visiting paths, collocation of pieces, and access/ticketing strategies. The recording and analysis of individual visitors' trajectories – possibly across the entire museum venue – is a great asset towards such behavioural analyses [4], allowing even to regenerate plausible visiting patterns [5,6]. Their feasibility has significantly grown during the last decades thanks, particularly, to the diffusion of Internet-of-Things (IoT) technologies [7–9], which enabled individual tracking needlessly of invasive structural modifications (e.g. as happens with overhead optical tracking sensors [10]). Nowadays, countless portable IoT devices broadcast periodically their identities on the Bluetooth and/or Wi-fi networks. Measuring at different locations the strength of these signals, the so-called RSSi (i.e. the Received Signal Strength indication), yields a mechanism to perform localisation and tracking (cf. also reviews [11,12]). Such an approach has been employed in different fields, allowing tracking in healthcare facilities [13] and smart buildings [14].

From an operational point of view, one provides (consenting) individuals (or groups) with a portable IoT *beacon* (note that personal mobile phones, modulo privacy and randomisation issues, could be used likewise [3]). Antennas preemptively deployed measure the RSSi of the periodic advertisements of the beacon. Thus, the visitor position can be roughly approximated with that of the antenna measuring the highest RSSi value [15] yielding an antenna-scale resolution of the museum. In principle, high antennas densities could also allow precise localisation through signal tri/multi-lateration [16–18]. However, in historic buildings, a frequent location for museums, massive antennas deployments are impossible due to architectural constraints, while room-level tracking allows sufficient insights. Additionally, even in optimal conditions (e.g. line of sight to the beacon, absence of radio interference) RSSi values typically suffer from high fluctuations [19]. Literature provides useful filtering techniques like Bayesian [20] and Kalman [21] filterings, capable to improve distance readings up to 40% in accuracy. Nevertheless, museums often feature complex geometries rich of obstacles and – especially in old constructions – a wide mixture of thick and thin walls with narrow and wide doors. In combination with crowding, these yield even noisier RSSi signals with a quick decay, causing ambiguous or even void positioning readings. These constraints likewise jeopardise the success of any approach based on instantaneously “maximum RSSi” readings (argMax), even when the ambition is the sole room-level localisation. To deal with these limitations, we proposed in [4] a method based on an end-to-end Multi-Layer Perception (MLP) neural network. This MLP significantly improved room-level localisation performances over naïve argMax policies by considering short RSSi time sequences. In particular, the localisation (binary) accuracy ‘acc’ (percentage of correct room-level localisation outputs) was increased from 54.7% (argMax policy) to 85.8% (MLP), considering an experimental dataset with 900 real-life trajectories collected by the same authors in the Galleria Borghese museum in Rome, Italy. Analysing signals in short time windows enabled higher robustness against

the fluctuations of RSSi. However, extending the time window still proved insufficient to tackle environments with a coarse/irregular antennas distribution, for which the MLP localisation is close to random. In general, typical museum geometries are an obstacle to end-to-end learning, including physically disconnected areas, although geometrically close, that can be ambiguously represented by RSSi readings. It is reasonable thus, that a methodology as an MLP trained on short RSSi windows is incapable of resolving these aspects: no “expert” knowledge, such as constraints dictated by museum topology (preventing, e.g., “nonphysical jumps” amongst rooms) is injected.

In this work, we propose a methodology to perform RSSi-based visitor tracking within a sparse antennas grid that combines an ensemble of trained classifiers, henceforth referred to as *localisers*, with an expert encoding of the museum topology in terms of a room-scale total-coloured graph. Since we replace a “complex” predictor with multiple simpler ones, which operate at different space- and time- scales, and thanks to the injection of expert knowledge, we manage to significantly increase the localisation accuracy up to $\text{acc} \approx 96.8\%$ for our use case at Galleria Borghese. Additionally, the total-coloured graph representation of a museum enforces a room-scale metric $\mathcal{D}$ and, consequently, a (room-level) trajectory-scale metric $\mathcal{W}$. We employ such metrics as indicators to compare reconstructed trajectories with the corresponding ground-truth, providing a more insightful performance analysis than the simple binary accuracy.

The rest of the paper is organised as follows: in Sec. 2, we review our experimental campaign at Galleria Borghese – the use-case for our methodology – and we discuss the basics of RSSi-based tracking; in Sec. 3 we introduce a topological representation for museum-like environments, based on total-coloured graphs; these are at the basis of the cascaded localiser method that we present in Sec. 4. The method performances are addressed in Sec. 5. A final discussion in Sec. 6 closes the work.

2. RSSi-based visitor tracking at Galleria Borghese

In this section, we discuss the basics of RSSi-based tracking considering our measurement installation at the Galleria Borghese (GB) museum, Rome, Italy, which we thoroughly described in our previous work [4]. GB hosts hundreds of art pieces in a relatively small area, and consists of 21 exhibition volumes divided into two floors, which are accessible via 3 entrances. Due to logistic constraints, visits are partitioned in 2 hour-long slots, non-overlapping. Like many other cultural heritage sites, GB has heavy limitations when it comes to set up a data-gathering infrastructure, e.g. no internet connection and, even worse, a very limited number of electrical outlets to power any sensing infrastructure. Room-level tracking based RSSi-measurements, demanding at least one available electrical plug per room, is thus a viable way to bypass these constraints. Tracking at room-level entails (anonymously) associating to each visitor their room-level trajectory $\mathbf{t}$ satisfying

$$\mathbf{t} = (t_0, t_1, \dots, t_{T-1}), \quad (1)$$

where t_t states the room (within a finite virtual room set $\{r_1, r_2, \dots, r_n\}$) in which the visitor is located at discrete time t , considering a regular time sampling $t \in 0, 1, \dots, T-1$.

We employ small commercial Bluetooth-Low-Energy (BLE) beacons that are assigned to each guest at the beginning of the visit and given back at the end. Each beacon broadcasts a signal carrying its unique identifier which we captured by a fleet of $A = 14$ receiving antennas deployed as good as allowed by the available electrical outlets (ten antennas on the first floor, roughly one per room, and four on the second floor which we consider as a single virtual room). We employ single-board Raspberry Pi computers as antennas, which are programmed to forward every 5 seconds the intensity of the beacon signal strength compared to the ground noise (along with its timestamp) to a remote database. Such a reading is often called Received Signal Strength intensity (RSSi) and it is provided as a relative index in arbitrary units. No unified standards are designed to what concerns RSSi values, therefore we employed the most commonly used: values are strictly negative and the stronger are the signals, the closer they are to zero. Such a system proved to be reliable with dozens of beacons broadcasting at the same time within a small area.

In Fig. 1(b) we report the RSSi readings by the different antennas for the trajectory sketched Fig. 1(a). With respect to Fig. 1 we report different critical issues: (i) the signal recovered by adjacent antennas may overlap continuously as highlighted by the inset, (ii) no signal is detected for a long amount of time, especially on the second floor (minutes $85 \sim 95$), (iii) antennas from the first floor may capture the beacon signals, even when it is on the second floor, (iv) some random antennas may detect beacons even after the visitor left the exhibition area (e.g. at the cafeteria or the wardrobe). In the following, we consider a dataset made of 900 individual beacons RSSi readings, gathered during 13 2h-long visit slots in the period between June and August 2019. For further details we refer to [4].

The gathered RSSi data does not have uniform temporal sampling so, we resample the signal with a fixed time-stepping $\Delta t = 10$ s. The choice of Δt is a trade-off between signal granularity and the desired resolution. Our choice yields 6 samples per minute: such a time interval is small in comparison to the typical permanence and traversal time of a room while allowing redundant RPi advertisement signals (RPi transmission frequency: 5 seconds). Given the captured RSSi from a single beacon, the output of this resampling procedure is a $A \times T$ matrix $\mathcal{R}$ (where T is the duration of the visit divided by Δt). In [4, 22] we proposed two methodologies to process such a matrix to estimate individual trajectories, the most effective of which was based on an MLP neural network operating on $\mathcal{R}$ after a row-normalisation (average over time = 0, standard deviation over time = 1), resulting in $\text{acc} = 85\%$ localisation accuracy. Specifically, our MLP was 3-layers deep, its input was a two-minutes long symmetric time window (i.e. $A \times (6 + 1 + 6)$ dimensional) and the localisation output was delivered via a soft-max activation defined over the room set.

In Sec. 4 we will introduce a new methodology that improves these results by achieving a global mean accuracy over 96% as we include assumptions on the topological structure of the museum. In the next Section, we present the graph interpretation of the museum at the basis of the method.




Figure 1 (a) Galleria Borghese floor plan. We report the position of antennas (a_{xy}) and rooms (r_x). Red lines represent closed doors. We include in blue a sketch of the trajectory whose RSSI measurements are in panel (b). The trajectory begins on the first floor (F1) and follows a counter-clockwise path before reaching the second floor (F2). (b) RSSI measurements for a real visitor trajectory (sketched in panel (a)). The readings from the $A = 14$ antennas are coloured depending on the virtual room (thus, measurements taken by any of the antennas on the second floor come with the same colour, yellow triangles). Each room change is represented by a different background colour.

3. Room-scale representation of Museums as total-coloured graphs

We consider here a generic museum divided into multiple floors; usually, a visitor explores one floor at time instead of changing it continuously. Two kinds of connections arise from this example: same and different floor links. However, we can think of other different architectural constraints that may influence a visit, e.g., different buildings. Similarly, we can provide conceptual subdivisions that concern rooms: e.g. if the exhibition is organised in different historical time periods, the typical visitor will likely explore a thematic area at time instead of traversing them randomly.

In the following, we formalise this approach by first sketching the museum as a total-coloured graph and then extracting its emerging clustering. We also provide a metric definition to make useful insights available when it comes to comparing trajectory reconstruction methods.

Extraction of a total-coloured graph In this paragraph we provide, by means of graph theory [23], a natural formalisation of a museum as a graph equipped with different colouring to represent its architectural and conceptual constraints (cf. Fig. 2).

A graph is a mathematical structure made of two sets, V and E , representing, respectively, entities and connections between them. Such entities are often called vertices and, in our representation, they symbolise the museum (virtual) rooms $V = \{r_1, r_2, \dots, r_n\}$ (i.e. single rooms or aggregates thereof depending on the antenna density).

We formalise rooms connections through ordered couples of rooms $e = (r_i, r_j)$, often called edges. In regular museums, connections are often bi-directional, meaning both (r_i, r_j) and (r_j, r_i) belongs to the edge set $E = \{e_1, e_2, \dots, e_m\}$; however, also one-way connections may occur (cf. Fig. 2(c)).




Figure 2 Use cases for total-coloured graphs in three fictitious museums. (a) A museum inspired by Galleria Borghese (Rome, Italy) with three floors and two kinds of exhibitions. In Sec. 3 we build our methodology on this example. (b) A segment of a museum where all the three floors host two kinds of exhibitions. Without considering the vertex-colour information, the two trajectories T_1 and T_2 would be very similar, however, they clearly are opposite and therefore different. (c) An example of a directed museum with multiple kinds of exhibitions based on Vatican Museums (Rome, Italy). Entrance and exit artificial rooms are at different locations.

We enforce a conceptual subdivision of the rooms by means of a set of *typological* information or, in graph theory jargon, colours. Similarly, we specify architectural constraints, e.g. being a door or a staircase, equipping also edges with a colour. More formally, a vertex-colouring (resp. edge-colouring) is a function γ that maps $v \in V$ (resp. $e \in E$) in a set of natural numbers $C \subset \mathbb{N}$. A graph (V, E, γ) in which both vertices and edges are supplied with colours is called *total-coloured* graph.

We add additional artificial rooms, marked with a dedicated colour, to represent the entrances/exits of the museum. These rooms will serve in the trajectory reconstruction process as placeholders for the visitors/beacon before the visit begins, after it ends, and, possibly, whenever no signal is detected. Conceptually, these rooms shall be sources or sinks of the graph, namely nodes with no incoming or outgoing edges, respectively.




Figure 3 Colour-contractions applied to the graph representation in Fig. 2(a). (a) vertex-colouring: the two kinds of exhibitions join in one single cluster each, suggesting that visitors explore them independently. (b) edge-colouring: by contracting over door edges ($\{\leftrightarrow\}$) we obtain a floor-based representation, prompting the idea that visitors would rarely deal with stairs if it is not needed. (c) total-colouring: contracting over both vertex-colouring and edge-colouring summaries the hints from the clustering (a) and (b), detecting four areas, likely visited independently.

Room clustering via colour-contraction We now consider *colour-contraction*, a methodology to extract different kind of clustering, i.e. partitions¹ $(K_1, K_2, \dots)$ of the room set V , from a coloured graph. This methodology, originally designed as a fast method to rescale big sets of data [24,25], allows us to infer insights from the typological information stored as colours in vertices and edges.

Depending on which colouring we rely on, we can build different partitions of V by joining two rooms r_i and r_j in the same cluster as follows:

vertex-colouring: if r_i and r_j are connected and share the same colour, then they are in the same cluster, say K_l , that is:

$$\text{if } (r_i, r_j) \in E \text{ and } \gamma(r_i) = \gamma(r_j) \text{ then } r_i, r_j \in K_l. \quad (2)$$

Note that two rooms do not need to be directly connected to belong to the same cluster. This aggregates all the rooms sharing the same conceptual information which are reachable one from another (cf. Fig. 3 (a)).

edge-colouring: if r_i and r_j are connected by an edge of some specific colour $C' \subsetneq C$, then they are in the same cluster, say K_l , that is:

$$\text{if } e = (r_i, r_j) \in E \text{ and } \gamma(e) \in C' \text{ then } r_i, r_j \in K_l. \quad (3)$$

As an example, by selecting door connections, this technique contracts all the rooms within the same floor of the museum in single clusters (cf. Fig. 3 (b)).

We can also combine these two methodologies, therefore obtaining a clustering of the rooms based on both architectural and conceptual constraints (cf. Fig. 3 (c)).

Trajectory and room metric enforced by the total-coloured graph To evaluate the quality of our trajectory reconstruction process (to be discussed in Sec. 4), we shall rely on a distance between the reconstructed trajectory and the corresponding ground-truth. A

¹A partition of a set V is a set of disjoint subsets $K_1, K_2, \dots \subseteq V$ such that their union gives the entire set.

total-coloured graph induces a room-level metric $\mathcal{D}$ and, on this basis, a trajectory-level metric $\mathcal{W}$, which we will use for evaluation purposes.

A “vanilla” distance between two rooms/vertices for a colourless graph is defined by the least number of edges traversed to reach one room from the other. Determining such a distance is a well-known graph theory problem, dubbed the *shortest path problem*, which naturally extends to include edges weighing a non-unit distance (i.e. weighted edges) [26].

In the setting of coloured graphs, the colours themselves can be interpreted as weights, e.g. the weight of a colour can be proportional to the time employed to traverse that specific kind of edge. Therefore, the assignment of a weight $w(\gamma(\cdot))$ to each edge-colour enforces also a distance notion between rooms².

Whenever two rooms come in different colours, since their typological information differs (cf. Fig. 2(b)), a penalty β is applied to their mutual distance. Note, whenever $\beta > 0$, the metric $\mathcal{D}$ does not satisfy the relaxation principle [27] (which can be still applied by adding β *a posteriori*).

The distance $\mathcal{D}$ is discrete in the vertices of the graph, therefore it can be represented as a $n \times n$ matrix. The distance is commutative if and only if all the connections in the underlying graph are bi-directional, therefore no assurance can be made *a priori* on the symmetry of $\mathcal{D}$. Formally, we have:

$$\mathcal{D}_{i,j} = \mathcal{D}(r_i, r_j) = \beta \cdot \mathbb{1}_{\{\gamma(r_i) \neq \gamma(r_j)\}} + \sum_{\substack{e \in \Gamma \\ \Gamma \text{ shortest path} \\ \text{between } r_i \text{ and } r_j}} w(\gamma(e)), \quad (4)$$

where $\mathbb{1}_{\{\varepsilon\}}$ is the indicator of the event ε , i.e. $\mathbb{1}_{\{\gamma(r_i) \neq \gamma(r_j)\}} = 1$ if $\gamma(r_i) \neq \gamma(r_j)$ and 0 otherwise.

Given two room-level trajectories $\mathfrak{s}$ and $\mathfrak{t}$ spanning T discrete time instants (cf. Eq. (1)), we can extend the room-level metric $\mathcal{D}$ to a trajectory-level metric $\mathcal{W}$ as the component-wise sum of the distances between instantaneous location. In formulas, it holds

$$\mathcal{W}(\mathfrak{s}, \mathfrak{t}) = \sum_{t=0}^{T-1} \mathcal{D}(\mathfrak{s}_t, \mathfrak{t}_t), \quad (5)$$

where $\mathfrak{s}_t$ and $\mathfrak{t}_t$ are the rooms in which trajectories are at time t . It is assumed that $\mathfrak{s}$ and $\mathfrak{t}$ have the same length T . In an applied perspective, we need to compare also trajectories of different lengths [4]. To this aim, we pad the end of the shorter trajectory with a virtual exit room such as its length matches the longer trajectory.

4. Cascaded trajectory reconstruction based on colour-clustering

In this section, we show how the multiple clustering of museum rooms we obtained from the total-coloured graph representation lay the foundation for cascaded classifiers that improve room-level localisation accuracy over the methods reviewed in Sec. 2.

²As an additional feature, it can be useful to reduce the cost of the first door transition, actually decreasing the weight of short transitions that can occur if visitors stand still by the entrance door.




Figure 4 A possible design for a two-layer cascaded selector built upon the museum in Fig. 2(a), as enforced by the clustering in Fig. 3(b).

Our clustering creates a partition of V in aggregates of rooms $K_1, K_2, \dots$ (clusters) that are, by construction, explored by visitors one at a time. Thus, our localisation approach first identifies the visitor in one amongst the different room clusters, say K_l , and, only afterwards, in a specific room within K_l . From a practical point of view, this corresponds to building one “large-scale” localiser that classifies at cluster-level $K_1, K_2, \dots$, and multiple “fine-scale” localisers, each operating within a single cluster, and returning a specific room. Each localiser, at either large- or small- scale, operates on all or a subset of $A' \leq A$ RSSi signals, possibly down-sampled in time. As RSSi are highly fluctuating, analysing data from a time window of, say, $T' \leq T$ samples, instead of the single measurements is mandatory. We refer to the time window as symmetric if data from both the future and the past is used, and asymmetric otherwise (necessary for online trajectory reconstruction). In the case of the “large-scale” classifier, considering a time window including a down-sampled version of the RSSi signal yielded better performance since it limited the input dimension of the classifier and prevented noisy inputs.

Considering the example in Fig. 3(b), we would employ three localisers: a floor selector $\mathfrak{F}$ fed with the signal from all the antennas that outputs one amongst the three areas (F0, F1 or F2) of the museum; two room selectors $\mathfrak{R}_1$ and $\mathfrak{R}_2$ (to be mutually employed depending on the output of $\mathfrak{F}$) that use the signals from the receivers on their respective floor to output the most probable room within that floor (note that the area F0 has a single room associated to, so it does not need a selector $\mathfrak{R}_0$; cf. Fig. 4).

Individual localisers We now briefly outline different choices for individual localisers. Depending on the amount of data available one could opt for heuristics or data-driven methods.

Heuristic algorithms heavily depend on the museum topology. Convolutional kernel approaches (ConvKer) yield a baseline. Specifically, we consider the argMax amongst the antennas, applied to a weighted mean of the RSSi signals within the time window.

Many options are available for classification algorithms trained with ground-truth data [28].

We specifically consider the following three, widely used in the literature:

Multi-Layer Perceptron (MLP). An MLP [29] is an artificial neural network made of an ensemble of units (perceptrons) organised in at least three consecutive layers as a directed graph. Data is assigned to the first layer units (which are $A' \cdot T'$) and flows into the next layers as a non-linear activation function (sigmoid) applied to a linear combination of the values of the previous layer. The last layer provides a classification output as a probability over A' classes via a soft-max function. The training procedure iteratively back-propagates through the layers modifying the linear combination weights.

Long-Short Term Memory (LSTM). A LSTM [30] is a particular kind of recurrent neural network (RNN), i.e. artificial neural networks where data can also propagate backwards, designed specifically for time series. LSTM are fed with data from the A' sensors time bin by time bin and progressively fill up an internal cell state whose of finite size. The data flow from the input and the state to the output, being controlled by multiplicative and non-linear gates. The training procedure back-propagates through time in the cell state, modifying the gates weights.

Random Forest (RF). A RF [31] is a predictor consisting of an ensemble of randomised decision trees, i.e. weaker tree classifiers where each internal node represents a query over the data features and each leaf is a classification label. The RF classifies following the majority vote over the tree ensemble, following the concept that weaker classifiers trained on different subsets of data learn various aspects from the features. The training procedure uses different dataset bootstrapping in order to train deterministically the various trees individually.

5. Trajectory reconstruction results in Galleria Borghese

In this section, we discuss the performance of our method on the basis of experimental data gathered in our measurement campaign in Galleria Borghese (see Sec. 2). In the following, we first build a total-coloured graph corresponding to the floor plan of GB (cf. Fig. 1(a)), and then we design a cascaded localiser relying on the corresponding clustering enforced by the architectural constraints. Finally, we present the performances of our methods by means of four indicators based on the metric enforced by the graph.

According to the disposition of antennas within the different exhibition areas, we chose to reconstruct the two visiting floors with different resolutions, building trajectories at room-scale on the first floor while having a single comprehensive room on the second one. We assigned a single colour to the so-formed $R = 9$ rooms. We include an entrance/exit artificial room connected to the three entry rooms of the museum (r_4 , r_5 and r_9) representing the ticket office (where we also provided visitors with beacons), assigning to it a different colour. The point of interest for the clustering phase consist, however, in the kinds of connections here represented in black (doors) and orange (stairs) (cf. Fig. 5(a)).




Figure 5 (a) Graph representation of Galleria Borghese (cf. Fig. 1(a) with the corresponding edge-clustering enforced by the orange connections ($\leftrightarrow$). (b) Structure of a cascaded selector based on the clustering in (a). (c) The distance matrix obtained weighting ‘2’ the door connections, ‘15’ the staircases connection and ‘0’ the different type of exhibitions. A discount factor of ‘ $1/2$ ’ is applied for the first door transition (cf. Footnote 2). The corresponding numeric values are reported in App. A (cf. Tab. 3).

Building upon the emerging clustering, we develop a two-layer cascaded localiser (cf. Fig. 5(b)) built of two predictors: (i) a floor predictor $\mathcal{F}$, whose input are the data from all the A antennas; (ii) a sculpture floor (F1) room predictor $\mathcal{R}$, whose input are the data from all the $A' \subset A$ antennas displaced in the rooms from F1.

In Tab. 1, we compare the localisation accuracy we achieve employing the methods presented in Sec. 4, considering both online and offline perspectives (i.e. involving asymmetric resp. symmetric time windows). Additionally, we report the results obtained via a single classifier (as in [4]) as a benchmark.

On side of the sole accuracy, we include in the comparison other metrics induced by our total-coloured graph and based on the room-distance matrix $\mathcal{D}$ (cf. Fig. 5(c), Tab. 3). We evaluate these metrics on a test set, $\mathcal{T}$, disjoint from the training set consisting of about 20% of the labelled data. The metrics considered are

acc the “vanilla” accuracy, given by the fraction of correct predictions over the total number of samples.

$\overline{\mathcal{D}}$ mean sample displacement, given by the average distance $\mathcal{D}$ between the sample, “instantaneous”, predictions and the corresponding ground truth.

$\overline{\mathcal{D}}^*$ the mean sample displacement error, given by the average distance $\mathcal{D}$ restricted to wrongly labelled samples, and the corresponding ground truth.

$\overline{\mathcal{W}}$ the mean trajectory displacement, given by the average distance $\mathcal{W}$ between the reconstructed trajectory and the corresponding ground truth.

The proposed method proves not only to increase the accuracy but also to commit much smaller errors when a room is misclassified. In particular, when adopting $\mathcal{F} = \text{RF}$ and $\mathcal{R}$

Symmetric time window $[-6, +6]$					Asymmetric time window $[-6, 0]$				
Methodology	Accuracy	$\overline{\mathcal{D}}$	$\overline{\mathcal{D}}^*$	$\overline{\mathcal{W}}$	Methodology	Accuracy	$\overline{\mathcal{D}}$	$\overline{\mathcal{D}}^*$	$\overline{\mathcal{W}}$
Single Localiser					Single Localiser				
ArgMax	0.6562	4.3755	12.727	3150.3	ArgMax	0.6562	4.3755	12.727	3150.3
ConvKer	0.7645	3.1693	13.458	2281.9	ConvKer	0.7105	3.4537	11.930	2486.6
MLP	0.8679	1.4402	10.903	1037.0	MLP	0.8525	1.5483	10.497	1114.7
LSTM	0.9267	0.5761	7.8605	414.84	LSTM	0.8721	1.1459	8.9600	825.10
RF	0.8900	1.2150	11.046	874.84	RF	0.8662	1.3524	10.108	973.76
Floor Localiser					Floor Localiser				
ArgMax	0.8673	1.9905	15	–	ArgMax	0.8673	1.9905	15	–
ConvKer	0.9005	1.4925	15	–	ConvKer	0.8986	1.5210	15	–
MLP	0.9045	1.4325	15	–	MLP	0.9011	0.9480	15	–
LSTM	0.9650	0.5250	15	–	LSTM	0.9480	0.7800	15	–
RF	0.9386	0.9210	15	–	RF	0.9258	1.1130	15	–
Floor Localiser with downsampling ($\Delta t = 5\text{min}$)					Floor Localiser with downsampling ($\Delta t = 5\text{min}$)				
ArgMax	0.9148	1.2780	15	–	ArgMax	0.9148	1.2780	15	–
ConvKer	0.8201	2.6984	15	–	ConvKer	0.8283	2.5754	15	–
MLP	0.9949	0.0764	15	–	MLP	0.9932	0.1020	15	–
LSTM	0.9940	0.0900	15	–	LSTM	0.9966	0.0509	15	–
RF	0.9966	0.0509	15	–	RF	0.9957	0.0644	15	–
Sculpture floor Localiser					Sculpture floor Localiser				
ArgMax	0.8111	0.3314	1.7546	–	ArgMax	0.8111	0.3314	1.7546	–
ConvKer	0.8974	0.1327	1.2938	–	ConvKer	0.8258	0.2165	1.2429	–
MLP	0.9493	0.0593	1.1714	–	MLP	0.9043	0.1072	1.1212	–
LSTM	0.9507	0.0521	1.0588	–	LSTM	0.9058	0.1260	1.3385	–
RF	0.9435	0.0709	1.2564	–	RF	0.8928	0.1390	1.2973	–
Multiple Localiser (Floor + Sculpture)					Multiple Localiser (Floor + Sculpture)				
RF – LSTM	0.9677	0.0816	2.5267	58.760	LSTM – LSTM	0.9413	0.1250	2.1311	90.068
					LSTM – RF	0.9337	0.1326	2.0008	95.510

Table 1 Localisation performance achieved employing a combination of selectors (Sec. 4) on experimental data collected at Galleria Borghese. Best results are highlighted in yellow. Performances with single localisers are reported as a reference. All the localisers have been trained on a desktop computer in less than five minutes, being the RF the fastest in training ($1/30$ of MLP training) and *LSTM* the slowest (three times the MLP training). The ConvKer employs a triangular linear kernel.

= LSTM, the mean error committed by the method is $\overline{\mathcal{D}} \approx 0.08$ compared to $\overline{\mathcal{D}} \approx 0.57$ by a single LSTM localiser. This corresponds to an effective displacement restricted to misclassifications of $\overline{\mathcal{D}}^* \approx 2.5$, where 1 means the visitor is positioned in an adjacent room and 3 means the visitor is located 2 rooms away from the ground truth. We note that this major increase is due to the “large-scale” classifier generally labelling correctly (accuracy $> 99.5\%$) the floor a visitor is at. $\mathfrak{R} = \text{LSTM}$ proves to be the best sculpture floor selector, and, overall, using a symmetric time window provides higher results than the online approach. In the case of the floor selector $\mathfrak{F}$, online and offline approaches based on RF and LSTM delivered comparable results. This aspect is probably due to the fact that the floor classification problem is more regular and simpler than the room classification.

6. Conclusion

In this work, we addressed the issue of increasing the accuracy of room-level tracking of museum visitors based on RSSi analyses. RSSi-tracking leverages on the measurement of the intensity (RSSi) of self-broadcasting signals emitted by beacons individually assigned to the visitors. As such, it is a flexible non-invasive method to acquire insights on the guest dynamics with very low hardware requirements, and it is able to overcome most of the museums architectural problems. However, RSSi are by nature highly fluctuating, requiring more advanced policies than the “maximum-RSSi” in order to reconstruct trajectories.

Here, via a graph theory approach, we proposed an encoding for the room-level structure of a museum, including topological information (expert-knowledge) in terms of colours. Based on the graph encoding we obtained room- and trajectory- metrics and different room-level clustering. These enable a cascaded localisation approach built of classifiers, organised in layers, that predict the position of a beacon carried by a visitor via progressive refinements of its location. The advantages are multiple: (i) each predictor may employ a different subset of the antennas, thus lowering the noise of the readings, (ii) lower dimensional feature and output space, hence simpler training procedures, (iii) different classifiers and time-scales which may suit better the distinct classification tasks. Reaching an accuracy of $> 96.7\%$ for our use case in Galleria Borghese, Rome, the proposed approach greatly improves on our previous methodology. By means of the metric induced by the graph, we get even more insights on the trajectory reconstruction quality, which has a mean displacement error smaller than 0.08 rooms.

The cascaded structure proposed can be extended to consider multiple progressive clustering as it happens, e.g., for a museum composed of multiple buildings, in which each building is made of multiple floors. In general, we consider a clustering with k refinements (i.e. $k + 1$ clustering built hierarchically), we can build a cascaded classifier made of $k + 2$ layers of predictors. Our future work will focus on an automatic approach, based e.g. on PCA or SelectK, to identify optimal features and time-scales for each localiser layer.

Acknowledgements The authors acknowledge Emiliano Cristiani who envisioned and coordinated the experimental campaign at Galleria Borghese. Acknowledged is also the support of Roberto Natalini, Sara Suriano, Massimiliano Adamo, Pietro Centorrino, and all the staff of Galleria Borghese. EO acknowledges the research project “SMARTOUR: Intelligent Platform for Tourism” (No. SCN_00166) funded by the Ministry of University and Research with the Regional Development Fund of European Union (PON Research and Competitiveness 2007-2013). AC acknowledges the support of the Talent Scheme (Veni) research programme, through project number 16771, which is financed by the Netherlands Organization for Scientific Research (NWO).

References

- [1] Robinson, E.S.: The behavior of the museum visitor. New Series, Number 5. Publications of the American Association of Museums, Washington, D. C. (1928). ERIC: ED044919
- [2] Feng, Y., Duives, D., Daamen, W., Hoogendoorn, S.: Data collection methods for studying pedestrian behaviour: A systematic review. *Building and Environment* **187**, 107329 (2021). doi:[10.1016/j.buildenv.2020.107329](https://doi.org/10.1016/j.buildenv.2020.107329)
- [3] Yoshimura, Y., Sobolevsky, S., Ratti, C., Girardin, F., Carrascal, J.P., Blat, J., Sinatra, R.: An analysis of visitors’ behavior in the Louvre Museum: A study using Bluetooth data. *Environment and Planning B: Planning and Design* **41**, 1113–1131 (2014). doi:[10.1068/b130047p](https://doi.org/10.1068/b130047p)
- [4] Centorrino, P., Corbetta, A., Cristiani, E., Onofri, E.: Managing crowded museums: Visitors flow measurement, analysis, modeling, and optimization. *Journal of Computational Science* **53**, 101357 (2021). doi:[10.1016/j.jocs.2021.101357](https://doi.org/10.1016/j.jocs.2021.101357)
- [5] Balzotti, C., Briani, M., Corbetta, A., Cristiani, E., Minozzi, M., Natalini, R., Suriano, S., Toschi, F.: Forecasting visitors behaviour in crowded museums. In: *Proceedings of the 9th International Conference on Pedestrian and Evacuation Dynamics (PED2018)*, Lund, Sweden, August 22-24, 2018 (2018). doi:[10.17815/CD.2020.82](https://doi.org/10.17815/CD.2020.82)
- [6] Pluchino, A., Garofalo, C., Inturri, G., Rapisarda, A., Ignaccolo, M.: Agent-based simulation of pedestrian behaviour in closed spaces: A museum case study. *Journal of Artificial Societies and Social Simulation* **17**(1), 16 (2014). doi:[10.18564/jasss.2336](https://doi.org/10.18564/jasss.2336)
- [7] Falk, J.H.: *Identity and the museum visitor experience*. Routledge, Taylor & Francis Group, London and New York (2016). doi:[10.4324/9781315427058](https://doi.org/10.4324/9781315427058)
- [8] Yalowitz, S.S., Bronnenkant, K.: Timing and tracking: Unlocking visitor behavior. *Visitor Studies* **12**(1), 47–64 (2009). doi:[10.1080/10645570902769134](https://doi.org/10.1080/10645570902769134)

- [9] Chianese, A., Piccialli, F.: Designing a smart museum: When cultural heritage joins iot. In: 2014 Eighth International Conference on Next Generation Mobile Apps, Services and Technologies, pp. 300–306 (2014). [doi:10.1109/NGMAST.2014.21](https://doi.org/10.1109/NGMAST.2014.21)
- [10] Pouw, C.A., Willems Joris Van Schadewijk, F., Thureau, J., Toschi, F., Corbetta, A.: Benchmarking high-fidelity pedestrian tracking systems for research, real-time monitoring and crowd control. In: submitted (2021). [Arxiv-2108.11719](https://arxiv.org/abs/2108.11719)
- [11] Deak, G., Curran, K., Condell, J.: A survey of active and passive indoor localisation systems. *Computer Communications* **35**(16), 1939–1954 (2012). [doi:10.1016/j.comcom.2012.06.004](https://doi.org/10.1016/j.comcom.2012.06.004)
- [12] Sadowski, S., Spachos, P.: Rssi-based indoor localization with the internet of things. *IEEE Access* **6**, 30149–30161 (2018). [doi:10.1109/ACCESS.2018.2843325](https://doi.org/10.1109/ACCESS.2018.2843325)
- [13] Iadanza, E.: Designing an indoor real-time location system for healthcare facilities. In: Mediterranean Forum–Data Science Conference: First International Conference, MeFDATA 2020, Sarajevo, Bosnia and Herzegovina, October 24, 2020, Revised Selected Papers, vol. 1343, p. 110. Springer Nature (2021). [doi:10.1007/978-3-030-72805-2_8](https://doi.org/10.1007/978-3-030-72805-2_8)
- [14] Varma, P.S., Anand, V.: Random forest learning based indoor localization as an iot service for smart buildings. *Wireless Personal Communications* **117**(4), 3209–3227 (2021). [doi:10.1007/s11277-020-07977-w](https://doi.org/10.1007/s11277-020-07977-w)
- [15] Piccialli, F., Cuomo, S., Schiano di Cola, V., Casolla, G.: A machine learning approach for IoT cultural data. *Journal of Ambient Intelligence and Humanized Computing* (2019). [doi:10.1007/s12652-019-01452-6](https://doi.org/10.1007/s12652-019-01452-6)
- [16] Wang, X., Bischoff, O., Laur, R., Paul, S.: Localization in wireless ad-hoc sensor networks using multilateration with rssi for logistic applications. *Procedia Chemistry* **1**(1), 461–464 (2009). [doi:10.1109/GLOCOM.2009.5425237](https://doi.org/10.1109/GLOCOM.2009.5425237)
- [17] Xue, W., Qiu, W., Hua, X., Yu, K.: Improved wi-fi rssi measurement for indoor localization. *IEEE Sensors Journal* **17**(7), 2224–2230 (2017). [doi:10.1109/JSEN.2017.2660522](https://doi.org/10.1109/JSEN.2017.2660522)
- [18] Yang, B., Guo, L., Guo, R., Zhao, M., Zhao, T.: A novel trilateration algorithm for rssi-based indoor localization. *IEEE Sensors Journal* **20**(14), 8164–8172 (2020). [doi:10.1109/JSEN.2020.2980966](https://doi.org/10.1109/JSEN.2020.2980966)
- [19] Beder, C., Klepal, M.: Fingerprinting based localisation revisited. A rigorous approach for comparing RSSI measurements coping with missed access points and differing antenna attenuations. In: Proceedings of the International Conference on Indoor Positioning and Indoor Navigation (IPIN 2012), November 13–15, 2012 (2012). [doi:10.1109/IPIN.2012.6418940](https://doi.org/10.1109/IPIN.2012.6418940)

- [20] Mackey, A., Spachos, P., Song, L., Plataniotis, K.N.: Improving ble beacon proximity estimation accuracy through bayesian filtering. *IEEE Internet of Things Journal* **7**(4), 3160–3169 (2020). [doi:10.1109/JIOT.2020.2965583](https://doi.org/10.1109/JIOT.2020.2965583)
- [21] Zhou, C., Yuan, J., Liu, H., Qiu, J.: Bluetooth indoor positioning based on rssi and kalman filter. *Wireless Personal Communications* **96**(3), 4115–4130 (2017). [doi:10.1007/s11277-017-4371-4](https://doi.org/10.1007/s11277-017-4371-4)
- [22] Centorrino, P., Corbetta, A., Cristiani, E., Onofri, E.: Measurement and analysis of visitors' trajectories in crowded museums. In: *Proceedings of 2019 IMEKO TC-4 International Conference on Metrology for Archaeology and Cultural Heritage*, Florence, Italy, December 4-6, 2019, pp. 423–428 (2019). [ArXiv:1912.02744](https://arxiv.org/abs/1912.02744)
- [23] Diestel, R.: *Graph Theory*, 5th edn. Springer Publishing Company, Incorporated (2017). ISBN: 978-3-662-53622-3
- [24] Mansour, N., Ponnusamy, R., Choudhary, A., Fox, G.C.: Graph contraction for physical optimization methods: a quality-cost tradeoff for mapping data on parallel computers. In: *Proceedings of the 7th International Conference on Supercomputing*, pp. 1–10 (1993). [doi:10.1145/165939.165942](https://doi.org/10.1145/165939.165942)
- [25] Lombardi, F., Onofri, E.: Graph contraction on attribute-based coloring. *Procedia Computer Science* (2022). The 13th International Conference on Ambient Systems, Networks and Technologies (ANT)
- [26] Ahuja, R.K., Mehlhorn, K., Orlin, J., Tarjan, R.E.: Faster algorithms for the shortest path problem. *Journal of the ACM (JACM)* **37**(2), 213–223 (1990). [doi:10.1145/77600.77615](https://doi.org/10.1145/77600.77615)
- [27] Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: *Introduction to Algorithms*, 2nd edn. The MIT Press (2001). ISBN: 978-0-262-03293-3
- [28] Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press (2016). URL <http://www.deeplearningbook.org>
- [29] MacKay, D.J.C.: *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, USA (2002). ISBN: 978-0-521-64298-9
- [30] Lara-Benítez, P., Carranza-García, M., Riquelme, J.: An experimental review on deep learning architectures for time series forecasting. *International Journal of Neural Systems* **31**(3) (2021). [doi:10.1142/S0129065721300011](https://doi.org/10.1142/S0129065721300011)
- [31] Biau, G., Scornet, E.: A random forest guided tour. *Test* **25**(2), 197–227 (2016). [doi:10.1007/s11749-016-0481-7](https://doi.org/10.1007/s11749-016-0481-7)

A. Distance Matrices

We report here the distance matrices $\mathcal{D}$ in the case of the fictitious museum floor plan in Fig. 2(a) and for Galleria Borghese, respectively in Tab. 2 and Tab. 3.

	r_0	r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9	r_{10}	r_{11}	r_{12}
r_0	0	11.5	10.5	11.5	12.5	12.5	11.5	10.5	11.5	20.5	20.5	21.5	21.5
r_1	11.5	0	1.5	2.5	3.5	1	2.5	3.5	4.5	10	10.5	11.5	11
r_2	10.5	1.5	0	1	2	2.5	3	4	3	10.5	10	11	11.5
r_3	11.5	2.5	1	0	1	3.5	4	3	2	11.5	11	12	12.5
r_4	12.5	3.5	2	1	0	4.5	3	2	1	12.5	12	13	13.5
r_5	12.5	1	2.5	3.5	4.5	0	1.5	2.5	3.5	11	11.5	12.5	12
r_6	11.5	2.5	3	4	3	1.5	0	1	2	12.5	12	13	13.5
r_7	10.5	3.5	4	3	2	2.5	1	0	1	13.5	13	14	14.5
r_8	11.5	4.5	3	2	1	3.5	2	1	0	13.5	13	14	14.5
r_9	20.5	10	10.5	11.5	12.5	11	12.5	13.5	13.5	0	1.5	2.5	1
r_{10}	20.5	10.5	10	11	12	11.5	12	13	13	1.5	0	1	2.5
r_{11}	21.5	11.5	11	12	13	12.5	13	14	14	2.5	1	0	1.5
r_{12}	21.5	11	11.5	12.5	13.5	12	13.5	14.5	14.5	1	2.5	1.5	0

Table 2 Distance matrix $\mathcal{D}$ from the total-coloured graph in Fig. 2(a). We weight ‘1’ the door connections ($\leftrightarrow$), ‘10’ the staircase links ($\leftrightarrow$) and +0.5 the distance between two rooms not sharing the same room-colour.

	r_0	r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9
r_0	0	16	19	16	15	15	16	19	16	15
r_1	16	0	1	3	5	1	7	5	3	20
r_2	19	1	0	1	3	3	5	7	5	18
r_3	16	3	1	0	1	5	3	5	7	16
r_4	15	5	3	1	0	7	1	3	5	15
r_5	15	1	3	5	7	0	5	3	1	22
r_6	16	7	5	3	1	5	0	1	3	16
r_7	19	5	7	5	3	3	1	0	1	18
r_8	16	3	5	7	5	1	3	1	0	20
r_9	15	20	18	16	15	22	16	18	20	0

Table 3 Distance matrix $\mathcal{D}$ from GB total-coloured graph in Fig. 5(a). We weight ‘2’ the door connections ($\leftrightarrow$), ‘15’ the staircase links ($\leftrightarrow$) and we apply a discount of $1/2$ on the first door transition (cf. Fig. 5(c) for heat map corresponding to $\mathcal{D}$).