Data Science and Engineering
https://doi.org/10.1007/s41019-026-00358-6
Article in Press

Unleashing the Power of UWB for Indoor Mobility Analytics: A Museum Case Study

Received: 19 March 2026
Accepted: 10 April 2026
Published online: 27 June 2026

Cite this article as: Vecchia D., Hachem F., Molteni D. *et al.* Unleashing the Power of UWB for Indoor Mobility Analytics: A Museum Case Study. *Data Sci. Eng.* (2026). https://doi.org/10.1007/s41019-026-00358-6

Davide Vecchia, Fatima Hachem, Davide Molteni, Maria Luisa Damiani & Gian Pietro Picco

Unleashing the Power of UWB for Indoor Mobility Analytics: A Museum Case Study

Davide Vecchia[1], Fatima Hachem[2], Davide Molteni[1], Maria Luisa Damiani[2], Gian Pietro Picco[1*]

[1*]Department of Information Engineering and Computer Science (DISI), University of Trento, via Sommarive, 9, Povo, Trento, I-38123, Italy.
[2]Department of Computer Science, University of Milano, via Celoria, 18, Milano, I-20133, Italy.

*Corresponding author(s). E-mail(s): gianpietro.picco@unitn.it; Contributing authors: davide.vecchia@unitn.it; fatme.hachem@unimi.it; davide.molteni@unitn.it; maria.damiani@unimi.it;

Declarations

Ethics approval and consent to participate

The manuscript is based on data from tracking 1500+ visitors at the Museum of Science (MUSE) in Trento, Italy. As explicitly stated on p. 4 of the manuscript: *"Informed consent and other standard procedures in place at the museum for similar studies were properly followed; no personal information was collected."*

Availability of data and materials

As stated on p. 7 and elsewhere in the manuscript, we intend to publicly release all datasets in it upon acceptance, e.g., on Zenodo.

Funding

This work is partially supported by the Italian government via the NG-UWB project (MUR PRIN 2017), and by the project SERICS (PE00000014) and the ICSC National Research Centre for High Performance Computing, Big Data and Quantum Computing (CN00000013), both under the NRRP MUR program funded by the NextGenerationEU.

Authors' contributions

Davide Vecchia and Davide Molteni performed the in-field acquisition of UWB trajectories, later analyzed and interpreted with Gian Pietro Picco. Fatima Hachem and Maria Luisa Damiani designed and performed the processing to extract the stop and visit information from trajectories along with the corresponding analysis, later revised and expanded with Davide Vecchia and Gian Pietro Picco. Davide Vecchia, Fatima Hachem, Maria Luisa Damiani, and Gian Pietro Picco contributed to the writing of the manuscript.

Acknowledgments

We are grateful to M. Lanzingher, Director of MUSE, and V. Cozzio, Head of IT services, for making this study possible, and to D. Dal Piaz and D. Tombolato for their support.

Unleashing the Power of UWB for Indoor Mobility Analytics: A Museum Case Study

Abstract

Ultra-wideband (UWB) radios enable decimeter-level positioning. This is posited to unlock unprecedented sophistication and detail in indoor mobility analytics. However, to date industry and academia have resorted to primitive analyses based only on raw trajectories, unable to ascertain *where* individuals stop and *for how long*, crucial in many domains. Among these is the museum providing the real-world context and dataset for this paper. We deploy a UWB localization system in a 40×15 m^2 area containing 42 exhibits and track 1500+ visitors across 3 months. We confirm that commonplace analyses relying on UWB trajectories alone offer practical insights, yet cannot capture the two key dimensions above. Instead, we exploit state-of-the-art techniques to extract higher-level semantic trajectories directly capturing visitor behavior, and distill a multitude of detailed, multi-layered, actionable insights. Our experience concretely highlights the untapped, disruptive potential of UWB-based mobility analytics—and a way to seize it. To foster further research, in museums and beyond, we publicly release our large UWB dataset (9+ million positions) along with the one resulting from our higher-level analyses.

Keywords: Indoor mobility analytics, stop-move detection, ultra-wideband (UWB), positioning and tracking, museums and cultural heritage

1 Introduction

Motivation. The large-scale outdoor positioning offered by Global Navigation Satellite Systems (GNSS) is today an integral part of our lives. Nevertheless, positioning and tracking are also increasingly available *indoor*, thanks to recent advances in wireless communication. Several competing technologies exist [1–3]. These include WiFi and BLE, whose main advantage is their pervasive availability; however, these radios incur errors of several meters when used to estimate the distance between devices. In contrast, ultra-wideband (UWB) radios rely on narrow pulses ($\leq$2 ns) to improve time-of-arrival estimation and separation from multipath components, yielding decimeter-level accuracy in distance estimation (*ranging*), empowering applications with significantly finer-grained positioning capabilities. As a consequence, although UWB is less pervasive than the technologies above, its adoption is soaring, as witnessed also by their increasing penetration in the market for consumer electronics [4].

Among the many use cases envisioned for UWB, *indoor mobility analytics* is arguably one of the most prominent. For instance, as reported by the FiRa consortium [5], a smart retail system can track the paths followed by customers inside a shop and their proximity to the point of interest (POIs) represented by items on sale; this can provide information about the shopping behavior, e.g., whether and how the customer behavior is affected by the shop layout and advertisements. Similar considerations apply in other domains where it is important to determine how people make use of a target space and items in it, e.g., cultural and trade shows [6, 7], museums [8], or heritage sites [9]. In principle, the same information could be extracted from video footage. Nevertheless, unlike UWB, the use of cameras raises privacy concerns and suffers from occlusions especially in crowded environments [10].

Still, the opportunities offered by UWB have been largely neglected by the scientific community, despite the renewed interest in mobility data science [11]. We are not aware of techniques or studies geared towards UWB and exploiting its potential. Crucially, *no* *large, real-world UWB trajectory dataset for human mobility is publicly available* (§2), in turn stifling the development of techniques for UWB-based mobility analytics. As for commercial endeavors, the advertised product capabilities are typically limited to the creation of heatmaps based on the density of trajectory points vs. the target area—useful, yet rather primitive in terms of the insights enabled. Indeed, in the applications above it is important to find *where each person stops*, e.g., in front of an item on sale; heatmaps are unable to discriminate between people passing by and people stopping, let apart for how long. Individual trajectories contain this information but only after extra processing sifts stops from trajectories, calling for dedicated techniques.

Focus: Enabling museum analytics. These considerations are front and center in the use case we focus on, enabled by a collaboration with a museum, that *i)* provides a concrete motivation grounded in the vast body of museum literature, and *ii)* offers a real-world setup for collecting the large dataset we use to showcase analyses and findings enabled by our UWB-based methodology.

The museum curators are interested in gaining quantitative information on the behavior of visitors, e.g., to understand the relation between the layout of exhibits and their fruition by visitors, akin to the smart retail scenario above. In this context, the crucial question is: *"At what exhibits do they stop and for how long?"*

Answering quantitatively this question offers a primary indicator of heightened attention and visitor interest [12], whose centrality in museums has been recognized since 1935, when the pioneering work of Melton [13] introduced the practice of external observers manually recording the movements of visitors. Despite the obvious human effort, this so-called "timing and tracking" practice is still in use today [8, 14]. Nevertheless, individual stop durations are rarely recorded by human trackers, due to the high complexity and inaccuracies they introduce.

UWB offers unique advantages by yielding a fine-grained resolution for *both* space and time and in an *automated* fashion, thus removing the need for on-site human trackers and supporting an area size and number of simultaneous visitors well beyond the scale practically addressable by them.

Real-world setup and long-term campaign. Unfortunately, to date these potential gains are only hypothetical, due to the above lack of reported experiences. Therefore, we conducted our own long-term study at the museum premises, informed by the needs and experience of the curators.

We deployed a permanent UWB localization infrastructure in a 40×15 m^2 area (Figure 1a) on the 2^{nd} of the 5 museum floors, thus explored by the vast majority of visitors. The area hosts 42 exhibits (Figure 1b), the POIs in our context, whose fruition happens via different actions: *i) reading* text panels whose content is often split across multiple sections, likely causing slight movements; *ii) watching* looping videos; *iii) inspecting* items, e.g., soil tray, corals, tiger head, large insect model; *iv) interacting* with a touch screen. The number and nature of the exhibits yields a high diversity of visitor paths and increased value in the quantitative insights we derive. The MAIN and SPONSOR sub-areas are defined by curators based on different goals, providing another dimension of analysis. Finally, most exhibits, often delimited by colored panels (Figure 1a), are very close to each other: the maximum distance between adjacent POIs in our setup (Figure 1b) is only 47 cm, with a median of 8 cm. Therefore, reliably discerning them would be impossible with the poor resolution offered by Bluetooth (§2), reasserting the need for the high accuracy offered by UWB.

In our 3-month campaign, during the high-influx summer period, each visitor entering the museum was given the option to wear a UWB tag. Informed consent and other standard procedures in place at the museum for similar studies were properly followed; no personal information was collected. Overall, *1573 visitors volunteered to participate in the study.*

Goals and contributions. The resulting huge trove of real-world data is key to our multi-faceted contributions, whose overall goal is to investigate the extent to which

(a) Bird's eye view.

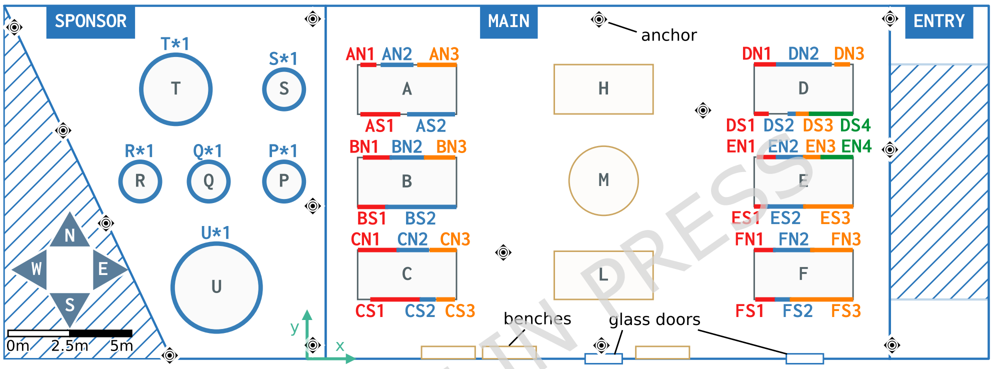

(b) Main elements in the area, layout of POIs, anchor placement.

Fig. 1: Museum target area. Visitors enter from one of the pathways in the entry area (right of Figure 1b). The POIs are denoted by the table they belong, followed by the cardinal point of the side they are placed on (or '∗' for those in SPONSOR) and a number. The central element `M` is the globe shown in Fig 1a; the rectangular elements `H` and `L` and the benches on both sides of it. The origin and Cartesian axes of the 2D reference system used for positioning is shown at the bottom of the SPONSOR area.

UWB-based mobility analytics can be performed *in practice*, paving the way to similar endeavors in related domains [5].

We first analyze the pure UWB trajectories, overall containing 9+ million positions, without any higher-level processing. This serves as a baseline mirroring the state of the art, whose value is amplified by the magnitude of our dataset. We show that this basic analysis already offers surprising insights to the curators, yet insufficient to answer the core question about where and for how long visitors stop.

The data mining community addressed similar questions for GPS outdoor trajectories, devising techniques [15, 16] that discriminate the *stops* near points of interest (POIs) from the *moves* between them. However, applications are, e.g., the study of migrating animals [17], stopping across large geographical areas and for months at a time, or of people [18] during their daily lives (e.g., home, workplace, stores) lasting hours at urban scale. Our museum case study demands instead a much finer spatio-temporal resolution where stops may be separated by <1 m (Figure 1) and last only seconds [12].

Recent studies closed the gap between these two realms. The work in [19] showed, for the first time, that existing techniques designed for coarse GPS trajectories are applicable, with proper configuration, to the much finer UWB ones. These results, obtained in highly controlled conditions with accurate ground truth, were later confirmed and extended in [20] based on actual trajectories from volunteers, providing a validated set of techniques to extract the position and durations of stops from UWB trajectories, associate them to POIs, and capture the key notion of *visit* to a POI.

The study presented here builds upon these earlier techniques and results by applying them to the large-scale dataset gathered during our real-world experimental campaign. In doing so, we retain very high confidence in the quantitative results and findings we obtain because our campaign has been carried out *in the same museum and with the same hardware, localization system, and environmental conditions* as in [19, 20]. This ensures that the corresponding techniques and configurations, whose accuracy and reliability has already been quantitatively ascertained and validated against ground truth in the controlled and semi-controlled setups of [19, 20], can be safely *applied "in the wild"* to our large-scale campaign without requiring another effort-demanding preliminary validation.

However, the work described here goes significantly beyond [19, 20], and in dimensions other than the mere scale of the dataset collected. Specifically, we provide a conceptual and technical framework for the joint use of UWB and stop-move detection at scale where UWB trajectories are abstracted into *semantic trajectories* [21–23] that *directly* capture the mobility patterns of visitors via the association of their stops with the POIs visited. In turn, this enables us to distill a multitude of multi-layered, actionable insights whose detail and depth is currently unmatched in the literature and cannot be achieved by raw UWB trajectories alone. We expect that, by showing *how* this can be achieved in a real-world setting and *what practical benefits* it can bring, we raise awareness about the untapped, disruptive potential of UWB for mobility analyses and pave the way for seizing it, beyond museums, fueling a new wave of research jointly tackling localization and its exploitation for mobility data science. To this end, we publicly release all our datasets [24], filling a fundamental gap across these two communities.

2 Background and Related Work

Observing visitor behavior. Museums are central to the preservation of cultural heritage and fulfill a fundamental educational role through exhibitions. Nonetheless, ensuring visitor engagement remains a challenge. In this respect the collection of information about the movement of visitors through exhibition spaces plays a key role, as it enables data-driven analyses to characterize the behavior of visitors and their level of engagement, therefore informing the decisions of the museum curators concerning the content and layout of exhibits. *Timing and tracking* is the term coined by the museum research community to denote the systematic process of observing, following, and recording visitor behavior within areas containing sets of exhibits [8].

This approach was popularized by the study by Serrell [12, 25] in 110 exhibitions of 62 museums, where human trackers recorded the total time spent in the exhibition area and number of visited POIs for 8000+ visitors.

Technology-based approaches. Timing and tracking remains widely used in museums [14, 26], as well as in retail [27] and other visitor-centric contexts [8]. Still, the high human effort involved and consequent difficult application to large areas and long observations motivated the exploration of alternative approaches, in line with the development of modern *smart museums* aiming to enhance visitor engagement and education through the use of digital technologies [28, 29].

Several approaches have been proposed, typically based on wireless radio technologies, cameras [10], or their combination [30]. Among these, the former are the vast majority because, unlike cameras, do not pose high privacy concerns and do not suffer from occlusions or light conditions. As for radio-based technologies, although experiences with other radios exist, e.g., RFID [31], Bluetooth Low Energy (BLE) is the technology of choice, thanks to its energy efficiency and pervasiveness in consumer devices, notably including smartphones.

In a typical scenario, small BLE beacons are placed in rooms or other designated places, while visitors carry their own smartphone or a dedicated device; based on their communication exchanges, *proximity* between the beacon and the user device can be determined. For instance, in [32] beacons are placed in the rooms of the Louvre museum in Paris, France, to analyze the frequency of visits and overall stay (hours). Similarly, the work in [33] analyzes the flow of visitors through the rooms of the Galleria Borghese museum in Rome, Italy, with the goal of devising models supporting analysis and simulation. The experience [34] in the CoBrA museum of modem art in Amstelveen, The Netherlands, similarly aims at predicting the visitor movement, this time based on the identification of "hotspots" inside the exhibit rooms. The work in [35], in the Monumental Cemetery museum of Pisa, Italy, complements the information obtained from beacons with BLE information "crowdsourced" from nearby devices, improving the ranging accuracy.

Technological limitations. This concise review of exemplary works evidences how these experiences generally yield a coarse spatio-temporal resolution, targeting rooms or areas inside them. In the few cases when individual POIs are the subject of study, these are placed far from each other, e.g., 1–5 m in [35]; in comparison, the maximum distance between adjacent POIs in our setup (Figure 1b) is only 47 cm, with a median of 8 cm.

The reason for the coarse resolution of existing approaches lies in the inherent limitations of detecting proximity via BLE. In its simplest form, proximity is determined simply by the reception of a packet; in [36], a fixed beacon determines the time a nearby user device has spent in a room as the time difference between the first and last packets received from it. A similar scheme but with flipped device roles is used in [34], using beacons at transmitters and user devices as receivers.

Other works improve over these very coarse estimates by similarly exploiting beacons as transmitters but determining proximity at the user device via the Radio Signal Strength Indicator (RSSI) available for each received packet. RSSI is a measure of signal strength that can be used to estimate the distance from the sender, i.e., the beacon, as the radio signal is known to attenuate as this distance increases (*path loss*) according to well-known models. Several museum experiences rely on this basic technique, e.g., [30, 37]. Nevertheless, the BLE literature shows that deriving accurate distance estimates is not straightforward: the decrease is not linear, path loss models require environment-dependent parameters, RSSI is very sensitive to obstacles and multipath effects, and BLE receivers determine its value with poor accuracy (±6 dBm) [38]. As a result, an abundant literature exists on improving BLE-based distance estimation, typically via filtering, signal processing techniques, and machine learning, which also found its way into museum experiences [33, 35, 39, 40]. Still, regardless of the technique used, BLE-based methods yield errors of *meters* [41], therefore relegating their applicability to museum scenarios with sparse POIs.

Another reason for the coarse resolution is the fact that these approaches require a beacon to be placed on *each* POI to be automatically identified. This obviously does not scale well to common situations where many POIs are present, like the 42 exhibits in the relatively small museum area we considered (Figure 1b). An alternative is to employ *localization* schemes, which exploit information from only a few beacons to determine the position of the target in a reference coordinate system. Unfortunately, the resulting positions are affected by the same large errors of the ranging estimates they are based upon [41], as witnessed by experiences in museums [40] and other contexts [42].

These limitations are well-known and have been addressed by subsequent versions of Bluetooth that increased accuracy in different ways. Bluetooth 5.1 introduced *direction finding* [43], i.e., the ability to determine the angle of arrival (AoA) or of departure (AoD) of the signal to/from an antenna array; based on the angles from several beacons, localization can be directly computed via triangulation. However, antenna array are bulky and expensive; therefore, they are typically used as fixed beacons covering the area of interest, while user devices compute their own position via AoD. Direction finding is known to provide sub-meter accuracy in controlled environments [44]; however, its accuracy in real-world scenarios is significantly lower [45], especially in cluttered environments, and can easily reach errors of several meters.

More recently, Bluetooth 6.0 introduced *channel sounding* [46] as an additional option that, by combining phased-based ranging (PBR) and round-trip time (RTT) techniques, provides more accurate distance estimation. The target accuracy of channel sounding is $\pm 10\,\text{cm}$. However, it has been recently shown [47] that the actual accuracy is worse and very sensitive to the orientation, height, power level, channel, and ranging algorithm used; depending on their combination, errors can be up to several meters. This technology is comparatively less mature and therefore has margins of improvement, potentially becoming a valid alternative for accurate ranging. However, another significant drawback is that, unlike RSSI and direction finding, channel sounding requires the two devices taking part in the ranging to setup and maintain a connection. This constraint is compatible with its motivating use cases, e.g., keyless vehicle entry, but is impractical in dynamic scenarios with many potential targets, like those typical of museums.

To the best of our knowledge, no museum experience exploiting direction finding or channel sounding have been reported in the literature.

High spatio-temporal accuracy: UWB. Nevertheless, Bluetooth is only one of the available technologies for ranging and localization, among which UWB radios are commonly considered the most accurate option [1–3]. UWB radios operate based on very short pulses ($\leq$2 ns) that, in comparison with narrowband radios like Bluetooth, better separate the signal from multipath components and offer very fine-grained time resolution, e.g., 15 ps in the UWB hardware we use. This enables UWB devices to accurately estimate the signal time of flight τ and, once multiplied by the speed of light c, estimate their distance $d = c \times \tau$ with decimeter-level accuracy and similar precision, i.e., without the long-tailed error distributions up to several meters that plague Bluetooth. These capabilities are at the core of the growing popularity of UWB, witnessed by the many Real-Time Location Systems (RTLS) commercially available and its increasing penetration in consumer electronics [4], including smartphones.

Interestingly, UWB is also being increasingly deployed in museums, which makes the contribution of this paper practically and immediately relevant. However, products focus on augmenting the user experience via automated guides [48, 49]; in the few cases where analytics are considered, these are limited to heatmaps [50] directly derived from UWB timestamped positions, whose limitations we already hinted at (§1) and further discuss later (§5.2).

Nevertheless, the potential of UWB for determining and analyzing the trajectories of museum visitors is still unexplored, with the exception of two related works. In [19], a controlled setup with accurate ground truth from cameras enabled a rigorous assessment of stop-move techniques and their configurations across 70,000+ positions containing 200+ pre-determined stops in front of museum exhibits. The work in [20] later confirmed and expanded the previous one by analyzing ∼220,000 positions from the trajectories of 10 volunteers moving and stopping freely across exhibits. In these works, and in the very same environment and experimental conditions as ours (§1), the authors reported the accuracy of UWB positioning as characterized by a median error of 46 cm and a 75^{th} percentile of 57 cm, with 96.2% of the samples with sub-meter error. These results confirms the superior accuracy of UWB over BLE, despite the non-line-of-sight (NLOS) conditions caused by the cluttered environment and the human body directly behind the UWB tag, inducing signal attenuation [51, 52].

Finally, although several publicly available UWB datasets exist, to the best of our knowledge *no large, real-world UWB trajectory dataset for human mobility is publicly available*. In fact, many datasets are generated with targets other than humans, e.g., drones and mobile robots (see e.g., [53] and its Table 1), sometimes even collected outdoors [54]. When human trajectories are present in the dataset, they are recorded only as a convenient way to explore the accuracy of UWB with a mobile target in addition to fixed ones; moreover, the number of humans involved is typically very small, often even a single person [55]. In other cases, e.g., [56], the dataset focuses on humans but does not contain position trajectories, rather raw signals (UWB channel impulse response, CIR) meant to support *device-free* passive localization or human activity recognition (HAR), in line with the vast body of literature based on radars, including UWB ones [57]. The only publicly available dataset for UWB trajectories of human subjects is the aforementioned one in [20]. In comparison, our dataset (§5) is *order of magnitudes larger*, as it contains $\sim$157× more subjects (1573 vs. 10) and $\sim$51× more UWB positions (11,335,936 vs. 220,000).

Extracting mobility patterns. The use of UWB ensures high accuracy and fine-grained spatio-temporal resolution for the analyses extracting higher-level mobility patterns from human trajectories. Stop-move detection techniques are the most established [15, 16], originally applied to outdoor GPS trajectories [17, 18, 58–61]. However, these studies are focus on much coarser spatio-temporal resolutions, like those determined by animals migrating across large geographical areas and staying in a given area for months, or by the urban scale relevant when analyzing the behavior of people during the stops, typically lasting hours, in the places relevant to their daily lives (e.g., home, workplace, stores). In contrast, the analysis of mobility patterns in *indoor* environments like ours remains a challenging problem [62–64], despite significant advances across multiple domains, e.g., including the development of shop recommender systems [65], the analysis of the flow of people in large spaces [66], or the determination of customer interest in a structured space [62, 63, 67]. However, like in the case of museums, these typically rely on BLE or other narrowband technologies, thus inheriting the above limitation of coarse spatio-temporal stop resolution. Unfortunately, studies of indoor mobility exploiting UWB are very scarce. In [68], a new stop detection technique is presented and evaluated on a small UWB dataset of 6 users and 70k points. The work in [69], is based on a larger dataset of 800k+ points, is instead focused on the specific application of spatial pedagogy, extracting metrics specifically geared towards this domain and not directly applicable to ours.

In this work, we build upon [19, 20] whose results showed, for the first time, that existing techniques designed for coarse GPS trajectories are applicable to the much finer UWB ones and identified the configuration yielding the best performance. Among the techniques considered, SeqScan [70] was found the most effective, yielding

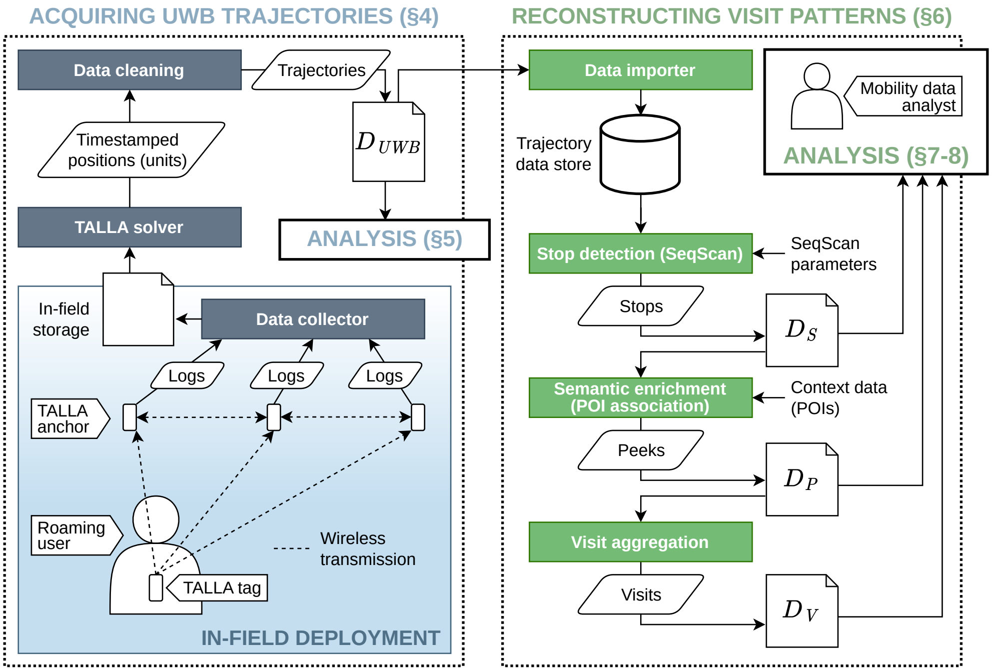

Fig. 2: Organization of the study, processing toolchain, and paper structure.

a stop detection F-score of ∼90% despite the above reported UWB positioning error. As discussed later (§6.2), we employ the same technique and its best configuration identified in these works, as they were validated in the very same museum area and hardware infrastructure we use.

Finally, there is a large body of literature concerned on extracting higher-level *semantic* trajectories [21, 22], describing movement as a sequence of *visits* to POIs. These include applications in museums [23], which are nonetheless based on BLE and therefore restricted to a coarse-grained spatio-temporal resolution. In contrast, we exploit UWB-based stop detection to support the definition of semantic trajectories capable to discriminate among POIs very close to each other, enabling analysis with unprecedented spatio-temporal detail.

3 Overview

Figure 2 illustrates the overall organization of our study and associated processing toolchain, both mirrored in the structure of the paper. The first step is the acquisition of UWB trajectories (§4) via the TALLA [71] state-of-the-art *real-time localization system (RTLS)* installed at the museum premises. The in-field fixed localization anchors and mobile user tags are equipped with the TALLA firmware coordinating the packet exchanges enabling localization and time synchronization. The resulting packet logs are sent by anchors to the TALLA solver, which exploits time-difference-of-arrival (TDoA) techniques to compute the tag position $\mathbf{p}_i = (x_i, y_i)$ at timestamp t_i based on the different propagation times of a *single* packet transmitted in broadcast by the tag and received by all N anchors at once. This technique is significantly more scalable, hence preferable in our museum context, compared to the computation of position based on individual distances between the tag and the N anchors, e.g., via two-way ranging [72], which would require a minimum of $2N$ packets per position instead of a single one.

The *trajectory* of the tag can therefore be formalized as a sequence $T = \langle(\mathbf{p}_1, t_1)\dots(\mathbf{p}_n, t_n)\rangle$ of individual spatio-temporal *units* $(\mathbf{p}_i, t_i)$. Figure 3 shows an example of trajectory from our dataset, in which each dot represents the position of the corresponding trajectory unit. Nevertheless, visitors may decide to skip the target area or their tag may malfunction; further, positions are subject to occasional outliers induced by communication noise. Therefore, a crucial and non-trivial *data cleaning* step is necessary to generate the final dataset of valid UWB trajectories.

We analyze this basic dataset to derive interesting insights (§5) germane to our context. Moreover, key to this paper, we use it as input to the second step *reconstructing the visitor behavior* (§6) with a level of detail currently unmatched in the literature. We first extract the *stops* performed by visitors via the SeqScan [70] technique, whose effectiveness on UWB trajectories has been recently proven [19]. SeqScan determines a stop as trajectory *segment* $S \subseteq T$: the stop position is the centroid of the

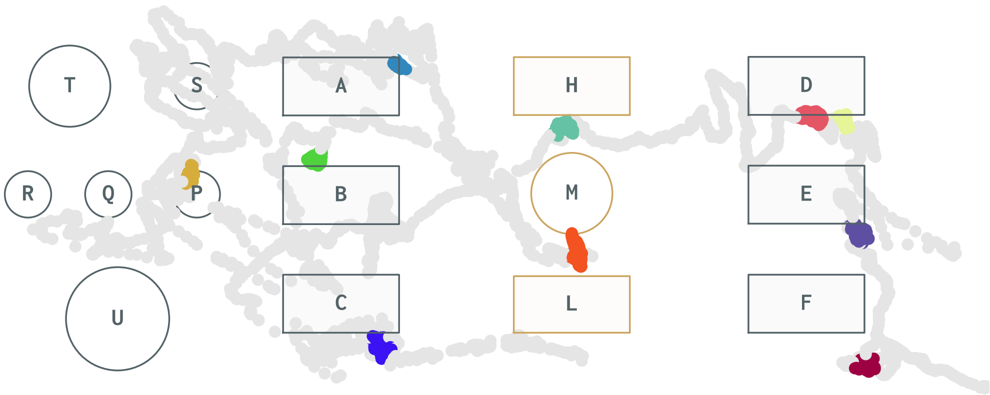

Fig. 3: A (partial) UWB trajectory whose units are segmented into moves (gray) and stops near a POI (i.e., *peeks*, colored).

units positions in S and its duration is the time interval between the first and last unit timestamp in S. We then extract the subset of stops occurring near a POI, hereafter called *peeks*, relevant to museum curators, whose units are shown in color in Figure 3.

Finally, we aggregate consecutive peeks to the same POI into a *visit* concisely capturing the visitor engagement with it. The resulting dataset, building upon and complementing UWB trajectories, is the basis for multi-layered insights focusing separately on each of the three concepts (§7) and showing how visits in particular enable powerful, at-a-glance, actionable characterizations of the interplay between visitors and POIs (§8).

4 Acquiring UWB Trajectories

We describe the first step in our processing toolchain.

4.1 In-field deployment

Tags and their management. The tags are MDEK1001 evaluation kits [73], equipped with the popular Qorvo DW1000 UWB radio [74]. Weighing only 41 g, battery included, they are easily worn on the chest of visitors, attached to a lanyard.

Tags were kept at the concierge right at the entry, where the museum staff explained the campaign goals to each visitor. Participation was completely optional; those who agreed to participate received a tag identified by a unique ID. No personal information was asked. We developed a graphical Web application used by the staff to easily mark the tag ID as handed out (or returned upon visitor exit) and *automatically* associate a timestamp. When not in use, tags were connected to pre-existing USB hubs for tablet-based guides. This aspect, and the need to avoid adding overhead on the museum staff, limited the number of available tags to 20.

Anchors. The fixed RTLS infrastructure comprises 14 anchor nodes (Figure 1b), each hosting the same UWB device used for tags, albeit equipped with a different firmware, and connected to a Raspberry Pi 3 (RPi). This powers the UWB device via USB, collects its data and sends it to the TALLA solver (§4.2), and enables remote configuration. The RPi is hidden in the false ceiling and connected to power and Ethernet via the existing PoE infrastructure. UWB devices are instead attached externally to ceiling, to avoid its signal obstruction, via existing metallic supports hosting other equipment, making them visually non-invasive.

Firmware and radio configuration. Anchors and tags host the corresponding TALLA firmware, handling the packet exchanges enabling TDoA localization. Communication occurs in slots, whose TDMA schedule repeats periodically with an epoch of 250 ms. Each device knows its fixed allocated slot(s) inside the epoch. Anchors are assigned one transmission slot in the first part of the epoch, used for synchronizing among themselves and with tags. In the rest of the epoch, anchors always listen; each tag is instead assigned 3 transmission slots, yielding a localization rate per tag of 12 Hz.

The UWB radio operates on channel 5, centered at 6489.6 MHz, with a pulse repetition frequency of 64 MHz and a preamble duration of $\sim$128 µs.

4.2 Position computation and storage

Anchors continuously generates logs of their own packet TX and of packet RX from tags and anchors, each associated with with epoch number, slot number, and a high-resolution timestamp (15.65 ps). These logs are forwarded both to a server on-site, acting as a short-term backup, and to a remote one where a daily automated procedure executes the TALLA solver with the logs as input. To ensure enough anchors with good synchronization are used for localization, we require that a packet TX from a tag is received by at least 7 anchors in the same epoch and slot. The solver computes all anchor sets matching this constraint along with geometrical ones and uses the set most recently synchronized to solve the nonlinear least-squares problem to find the tag position that minimizes the TDoA residuals [71].

Improving position accuracy. Due to the noisy nature of the wireless medium, some position estimates may have a large error. To limit the impact of these outliers on the rest of the toolchain, TALLA employs Kalman filters (KF) for both anchor clock drift estimation and tag position estimation. For the former we use a linear KF based on a constant velocity model. For localization, TALLA tracks the (x, y) coordinates using two Unscented Kalman filters (UKF) catering for different visitor mobility contexts: intervals with no acceleration (stationary) and those with acceleration (moving). These are then reconciled under Interacting Multiple Models (IMM) [75], which weights the output of the filters depending on the likelihood of the tag being in each state.

Storage. The TALLA solver outputs a trajectory as a sequence of entries $\langle \texttt{ID,x,y,t} \rangle$, where `ID` is the tag identifier and the remaining values are the position $\mathbf{p}_i = (x_i, y_i)$ and corresponding timestamp t_i associated to a trajectory unit $(\mathbf{p}_i, t_i)$. These entries are downloaded in CSV format from the server running the TALLA solver, they undergo data cleaning, and are finally uploaded into a database (PostgreSQL 13.2 64-bit, with extension PostGIS) to simplify processing via appropriate queries.

4.3 Data cleaning

Separating different uses of the same tag. The UWB trajectories output by TALLA are already associated with the corresponding tag ID. However, the same tag is typically reused by different visitors at different times, within and across days. Nevertheless, the tag pickup and return times are automatically set by our Web application (§4.1); we use them to compare against the UWB timestamps and, by inferring the use by different visitors, separate the trajectories accordingly.

Separating different tours of the same visitor. A visitor may tour the target area at different moments, e.g., returning to it after having explored other floors, to delve deeper into interesting exhibits. To capture this, we separate each tour into distinct trajectories when the temporal gap between two consecutive units is $\geq T_{gap}$. Nevertheless, the gap can also be due to localization noise leading to missing positions—at the extreme, even a single one ($T_{gap} = 1/12$ s). We set $T_{gap} = 60$ s, which provides a good balance between capturing intentional departures of the visitor from the target area and avoiding to break a trajectory prematurely due to transient noise.

Annotating positions in the target area. Visitors receive a tag when entering the museum (§4.1), not the target area. Therefore, we need a means to separate the portion of the UWB trajectory in the target area from the entry path to it (Figure 1b). Moreover, our target area is divided in two sub-areas, MAIN and SPONSOR, that bear different museum concerns and we analyze separately (§5.2). We cater for these issues by annotating each trajectory unit as $(\mathbf{p}_i, t_i, a_i)$ with an area label $a_i \in \{entry, \text{MAIN}, \text{SPONSOR}, none\}$ whose value is assigned based on where the unit position $\mathbf{p}_i$ falls in the floor plan.

Dealing with transitions and outliers. However, a visitor may move back and forth exactly at the transition between areas; the basic labeling above would create a rather fragmented annotation. Moreover, Kalman filters cannot eliminate outliers completely, e.g., non-line-of-sight (NLOS) conditions may induce transient tag misplacements. We tackle both issues as follows. First, we observe that a trajectory $T = \langle(\mathbf{p}_1, t_1, a_1)\ldots(\mathbf{p}_n, t_n, a_n)\rangle$ can be regarded as sequence $T = \langle\theta_1 \ldots \theta_m\rangle$ where each portion $\theta_i = \langle(\mathbf{p}_j, t_j, a)\ldots(\mathbf{p}_k, t_k, a)\rangle$ contains only units with the *same* value a for the label area. Then, to account for transient visitor movement or noise that unnecessarily fragments trajectories, we *merge* consecutive portions under given conditions. Trivially, if θ_i contains only 1 unit (e.g., an outlier), it is merged with θ_{i-1} and takes its area label a. The same concept extends to the case where the trajectory transiently "spills" from given area to another, i.e., from a portion θ_{i-1} into the next one θ_i, and back. However, in this case the decision about whether or not to merge θ_i into θ_{i-1} depends on whether the spatio-temporal entity of this detour is relevant. Specifically, the merge occurs only when: *i)* the duration of the detour θ_i is shorter than a threshold T_{area}, and *ii)* its units are within a spatial distance M from the border of the area of the previous area in θ_{i-1}. Hereafter, we set $T_{area} = 5$ s and $M = 1$ m.

Final clean dataset. To account for the processing described above, each unit entry is augmented as $\langle$`ID,x,y,t,tour,area`$\rangle$ by adding a `tour` identifier, to distinguishing the trajectories of returning users, and the `area` label containing the value $a \in \{\text{MAIN}, \text{SPONSOR}\}$. The final dataset contains the trajectories whose units are in the target area for a total of at least $T_{min} = 60$ s.

5 Analyzing UWB trajectories

We provide a concise characterization of the UWB dataset, then delve into its analysis to illustrate the insights it unlocks.

5.1 Dataset

The dataset (Table 1) was accrued over 58 actual days, due to days with closures and special events incompatible with the campaign; 1573 visitors volunteered to carry

Table 1: UWB dataset statistics.

Full dataset	
# tag delivered	1573
delivery errors	31
faulty tag/battery	42
# visitors	1500
UWB units	11 335 936
After data cleaning (§4.3)	
# visitors in target area	1396 (93.1%)
UWB units	9 112 682
# trajectories	1856
mean trajectory duration	7′46″
total duration	10d 0h 14′42″

Table 2: Trajectories per visitor.

#	# visitors	%
1	1039	74.43
2	269	19.27
3	74	5.30
4	13	0.93
5	1	0.07

a tag. However, in 31 cases (1.9%) this was incorrectly handled at the concierge, and in other 42 cases (2.6%) the battery malfunctioned. Overall, 1500 visitors were correctly tracked, yielding over 11 million units. Applying the data cleaning procedure (§4.3) yields *1856 trajectories containing over 9 million units* for 1396 visitors; only 104 visitors (6.9%) never stopped in the target area. Those who did spent an average of 7′24″, overall yielding the equivalent of >10 days of tracked visitor behavior. Moreover, more than a quarter of these visitors has multiple trajectories associated (Table 2), although only 6.3% have >2. Finally, the localization rate across all trajectories is 10.67 Hz, i.e., very close to the ideal 12 Hz: missing samples occur but their impact is limited.

5.2 Insights

Figure 4 illustrates the results of our analysis. The first metric, the total time visitors spend in the museum (Figure 4a), is actually derived solely from the logs of our application for tag distribution; on average, visitors remain in the museum between 2 and 3 hrs, matching the curators' expectation. We now discuss other insights enabled by the UWB dataset, exemplifying questions relevant to the curators.

When do they arrive in the target area? The museum has exhibits on floors both below and above our target area, with no strictly prescribed visit order. Based on the

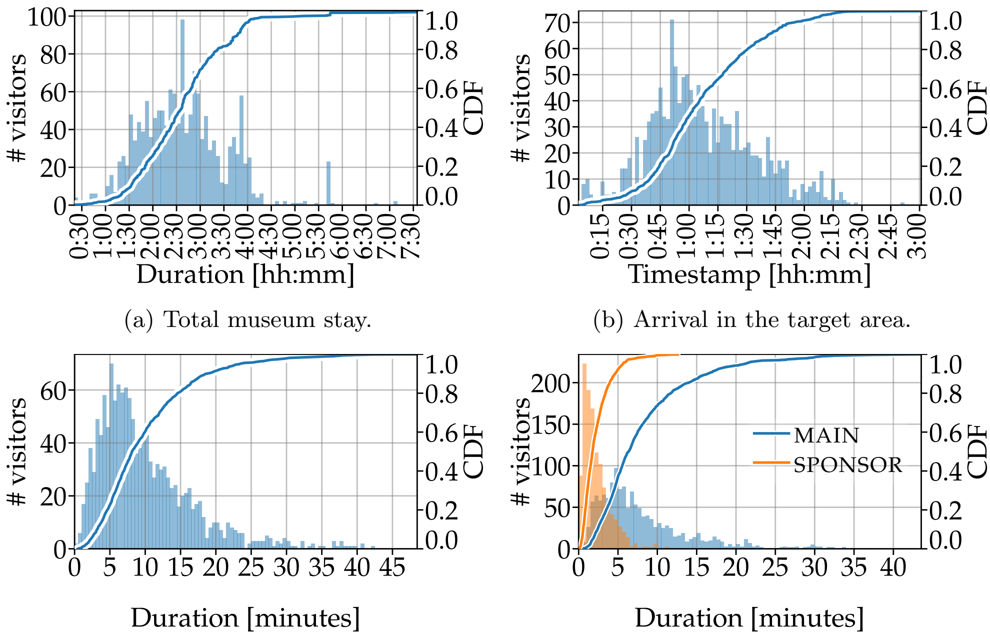

(a) Total museum stay. (b) Arrival in the target area.

(c) Stay in the target area. (d) Stay in sub-areas.

Fig. 4: Coarse-grained data analysis: UWB dataset.

difference between the start of the museum stay (logged at the concierge) and the first timestamp in the visitor trajectory, we can determine that visitors, on average, arrive after ∼1 hour, (Figure 4b) and roughly at 44% of their overall stay in the museum compatible with the placement on the 2^{nd} of 5 floors.

From which side do they arrive? The target area covers only a portion of the 2^{nd} floor. In the curators' intentions, visitors arriving from the stairs should first visit the other portion to the west, then enter the MAIN area from the south access, proceed to the SPONSOR area, and exit via the north access after going through MAIN again (Figure 1b). However, our data tells a different story. Based on the units at the start and end of the first trajectory of each visitor, this intended path is followed only in 49% of the cases; 37% of the visitors enter from north and split evenly in choosing the exit, while 14% enter and exit from south. This points to the need to signal more prominently the intended path, as the "wrong" north access is to the right of the stairs and people are naturally inclined to turn right at a fork [76].

How long do they stay? Each unit in the clean trajectories includes the area label, enabling us to derive the visitor stay in the target area, whose median duration is $8'8''$ (Figure 4c). Only 3.6% of the visitors remain for $<2'$, while $\sim$19% remain for $>15'$, up to a maximum of $48'30''$. The curators were specifically interested in assessing separately the visitor stay in each of the sub-areas, due to their different characteristics (Figure 4d). Overall, visitors spend significantly more time in MAIN: a median of $6'20''$ vs. $1'46''$ for those entering SPONSOR. This is not surprising, as MAIN is larger and must be traversed fully to access SPONSOR. A less obvious insight is that the stay distribution of SPONSOR is a lot more skewed towards very short durations, pointing to limited visitor interest in the area. This is confirmed by the fact that 211 visitors ($\sim$15%) never set foot in it.

Do they come back? UWB trajectories can also be used to understand how visitors organize their museum tour, i.e., whether they visit the target area all at once or instead take a first look and return later. To this end, we consider only the visitors with multiple trajectories (Table 2) and analyze the temporal gap separating them, amounting to $<5'$ in $\sim$70% of the cases, indicating brief detours in which arguably the interest remains centered on the target area. At the other extreme, $\sim$11% of the visitors return only after $>30'$.

Where do they go? Localization systems often provide heatmaps based on unit density to provide an aggregate view of the most frequented zones. Figure 5 shows the one we obtain from D_{UWB} via Kernel Density Estimation (KDE). Some high-density zones are clearly visible near the tables hosting the exhibits and over benches, as expected. Interestingly, other high-density areas appear at the bottom, far from the exhibit; these are close to a few benches along the wall but also to several laboratories closed to the public but whose glass door allows visitors to peek in, effectively providing

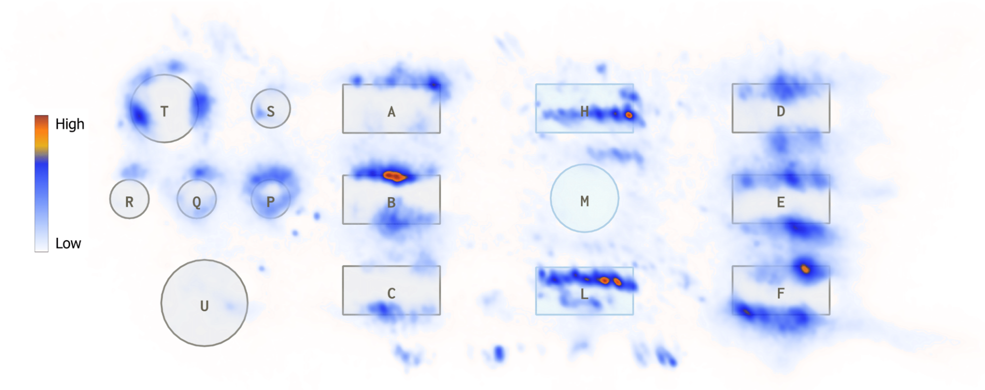

Fig. 5: KDE heatmap of the UWB dataset D_{UWB}.

another exhibit of sorts (Figure 1a). As for the real exhibits, this at-a-glance view shows also which zones are less visited, e.g., table `A` and `C` in MAIN and exhibit `R` and `U` in SPONSOR, along with the most frequented exhibits, `BN2` and `FN3`.

Key limitations. Nevertheless, analyses based on the pure UWB trajectories cannot fully answer the fundamental question *"'At what exhibits do they stop and for how long?"* (§1).

Consider the bottom of table `F` (Figure 5), hosting three exhibits. The heatmap shows a high density of units; we infer that this zone is highly "frequented". However, we are unable to discriminate whether visitors *stop* in front of the exhibits or instead simply pass by, nor we can determine which of these exhibits attracts more visitors and/or engages them for the longer time. Comparison is also impossible, for these and other exhibits in the target area. Interestingly, similar issues hold even for exhibits apparently providing a clear indication, like `BN2`: is its high unit density caused by a few engaged visitors spending a lot of time at it, or by many visitors attracted by the exhibit albeit for only a short time?

Discriminating among all these situations is key to empower curators with an understanding of how visitor interact with exhibits [25]. Still, the main appeal of heatmaps—concision—is also their limitation. As only *raw, aggregate movement* is captured, the *individual behavior* of visitors w.r.t. stops is essentially lost. A different, finer-grained type of analysis is necessary that, by fully exploiting the high spatio-temporal resolution offered by UWB trajectories, accurately extracts visitor stops ("where") and their duration ("for how long").

6 Reconstructing a visitor's tour

Understanding where visitors stop is key for museum professionals. In her seminal work, Serrell [12] identifies a visitor's stop in front of an exhibit as a primary indicator of heightened attention and interest. This is typically acquired by visually following and recording the behavior of a small sample of visitors in a small area. This very common practice, known as "timing and tracking" [8], has several drawbacks: it is effort-demanding, stop recognition is subjective, and does not scale to large spaces or long observation periods. In contrast, we present a *scalable* and *generic* framework to *automatically* identify stops near points of interest (POI) based on a UWB trajectory dataset.

6.1 Key concepts

The framework revolves around three key concepts.

Stops are the most basic component of our analyses, capturing the immobility of a visitor; two stops in a trajectory are always separated by movement. Stops can occur anywhere; in front of an exhibit, but also far from them, e.g., when resting at a bench or picking up a phone call.

Peeks capture instead the *semantics* of stops; a peek is a stop occurring near a POI. The assumption embodied in this association is that the stopped visitor is devoting attention to the POI. Still, the visitor could also be looking elsewhere or conversing with someone. Unfortunately, we cannot check what the visitor is actually looking at, like a human observer would do; we therefore use proximity to a POI for

Fig. 6: Movements of a visitor around a POI.

a minimum duration as a proxy for visitor interest. This decrease in quality w.r.t. timing and tracking, is the price to pay for scalability and automation and is also partly compensated by the ability to accurately measure duration.

Finally, visitors do not always observe a POI in a single stop. Multiple stops may occur due to external conditions, e.g., other people arriving and forcing the visitor to move or the desire to observe the POI from different vantage points (Figure 6). Therefore, we generalize the notion of peek into a *visit*, a sequence of one or more consecutive peeks at the same POI, capturing a visitor's interaction with it. Visits are the high-level concept that concisely represents the information needed to characterize a visitor's tour of the museum *and*, dually, the ability of a POI to attract and engage visitors.

6.2 Processing techniques

Determining the position and duration of a stop, at the core of all the concepts above, is actually tricky. Unintended unit misplacements induced by unavoidable localization errors may be incorrectly interpreted as movement of an immobile visitor. Dually, a visitor not walking may still move the torso, hence the tag, like the one in Figure 6 bending to get take a better view; small movements should be correctly recorded yet not break the stop. Human observers can easily compensate for this "behavioral noise", which instead easily hinders algorithms assuming full immobility during a stop.

These considerations also prevent the use of straightforward techniques simply mapping units into a spatial grid, as choosing the proper cell size would be difficult due to the decimeter-level UWB resolution and the intrinsic presence of noise. A grid cell too small would run the risk of "splitting" stops across cells; instead, a grid cell too large would *i)* reduce the fine-grained resolution offered by UWB to the coarser one defined by the cell, and *ii)* run the risk of merging multiple stops into one. In essence, although stops intrinsically occur within a region, the placement and size of this region are not known a priori; therefore, they cannot be reconciled effectively with a statically defined grid without losing either spatial resolution, accuracy in stop detection, or both.

Instead, the techniques we employ allow transient movement within a stop to tolerate both localization errors and behavioral noise. We describe how we use them to extract each of the concepts above from the UWB dataset, constituting the second phase of our processing toolchain (Figure 2).

Stop detection. In the literature concerning mobility analytics [16, 22], *stop-move* segmentation techniques are a common choice. Broadly, they divide a trajectory, consisting of a sequence of units, into temporally disjoint sub-sequences or *segments*, by labeling units as either *stop* or *move*.

These techniques were developed for GNSS, whose spatio-temporal resolution is significantly coarser than UWB. However, the works in [19, 20] recently confirmed their applicability by comparing and validating against ground truth the accuracy of multiple segmentation techniques applied to real-world UWB trajectories acquired in the same museum area targeted here. SeqScan [70] was reported to be the one showing highest stop detection accuracy and robustness to noise; we therefore employ it here.

Inspired by generic analysis techniques like DBSCAN [77], SeqScan identifies stops as density-based clusters of arbitrary shape. Unlike DBSCAN, however, SeqScan explicitly accounts for the temporal dimension, capturing the duration of stops and their temporal separation. Moreover, it caters for the aforementioned transient departures from stops, called *local noise* [70], recognized as *move* units that do not occur close to the stop position yet fall within the stop interval. These transient departures can be of arbitrary length and duration; abstractly, they indicate a *recursive* pattern (go-and-return) [78] that does not include any other stop.

However, this calls for defining the *duration* of the stop, including transient movements, and the actual *presence* around its position. This distinction, also introduced in [70], is key to reconcile the mechanics of stop detection with the spirit of timing and tracking where it is important to capture *"a visitor stopping with both feet planted on the floor"* [79].

Nicely, SeqScan is configured by only three parameters. N and ϵ specify the spatial density constraint inherited from DBSCAN: N is the minimum unit cardinality in the cluster neighborhood of radius ϵ. Hereafter, we use use $N = 24$ and $\epsilon = 15\,\text{cm}$ as these have been identified as the best choice in the study in [20], performed with the exact same environment and hardware (§1) as ours. The third parameter, ρ, captures directly the minimum presence relevant to the application; stops with presence $< \rho$ are discarded. Based on the indications of the curators, we set $\rho = 10$ s.

Figure 3, already commented upon earlier (§3), shows graphically the effect of a segmentation output by SeqScan. In the dataset D_S, stops are output as tuples $(\mathbf{c}, I, P)$, where $\mathbf{c}$ is the stop position obtained as the centroid of the positions of the units in the cluster labeled as stop, $I = [t_{start}, t_{end}]$ is the time interval delimited by the first and last unit, determining the stop *duration* $D = |I|$, and $P \leq D$ is the *presence* at the stop, obtained by excluding from I the time intervals corresponding to *move* units.

POI association. We determine peeks by selecting the stops in D_S that are near to a POIs in the target area. Proximity is based on the distance between the stop position $\mathbf{c}$ and the geometric shape representing the POI. The POIs in MAIN are placed next

Table 3: Stop, peek, and visit datasets.

	Stops (D_S)	**Peeks** (D_P)	**Visits** (D_V)
#	9735	8048	6793
# visitors	1323	1248	1248
Total duration	5d 19h 51′5″	3d 20h 36m 27s	3d 21h 11m 19s
Total presence	5d 02h 15m 38s	3d 07h 18m 16s	3d 07h 18m 16s

to each other on the side of tables and are often delimited by colored tiles (Figure 1a); we model them as line segments. Instead, the circular tables in SPONSOR host content that cannot be as easily separated; therefore, we consider the whole table as a POI (Figure 1b).

To preserve the semantics of peeks (§6.1) capturing visitor attention, we associate a stop and a POI only when their distance is below a threshold R; if multiple POIs match the constraint, only the nearest is recorded. A reasonable value of R depends on several factors, e.g., the layout, nature, and expected interactions with POIs. Those in our target area are text panels, screens, and interactive elements; after consulting with the curators, we set $R = 1$ m.

The peek tuples $(\mathbf{c}, I, P, o)$ output in the dataset D_P are obtained by augmenting the original stop tuple in D_S with the identifier o of the associated POI.

Visit aggregation. Finally, consecutive peeks at the same POI are straightforwardly aggregated into visits represented by tuples (I_V, P_V, o), where I_V is the interval defined by the beginning (t_{start}) of the first peek in the visit and the end (t_{end}) of the last one, P_V is the sum of the presence of peeks, and o is the POI identifier shared by all peeks.

7 Analyzing Stops, peeks, and visits

Applying the processing in §6.2 to the 1856 trajectories from 1396 visitors in D_{UWB} yields the three datasets summarized in Table 3. We discuss their characteristics, then exploit them to distill insights on visitors and POIs.

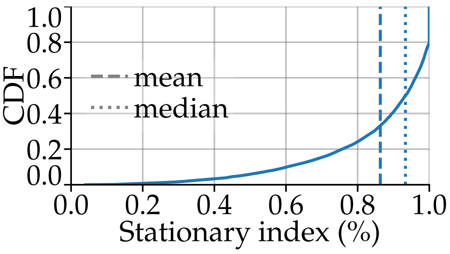

Fig. 7: Stationary index: presence over duration.

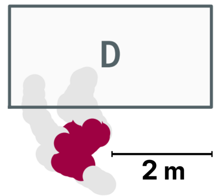

Fig. 8: A visitor moves while stopped at a POI.

7.1 Datasets

Presence vs. duration. We begin by inspecting the characteristics of the 9735 stops in D_S, as these are the building blocks for extracting the higher-level peeks and visits. Our techniques allow units labeled as *move* to appear among *stop* ones, to tolerate effectively transient movements induced by behavioral or localization noise. This yields the distinction between presence P and duration D (§6.2), whose values should be very close in the majority of cases, matching the intuition that a stop captures stationary behavior; the opposite would point to an incorrect configuration of SeqScan.

We investigate this aspect via the *stationary index* [70] $P/D \in [0, 1]$. In D_S (Figure 7) the index has mean 0.86 and median 0.93, and 20.75% of the stops are fully stationary ($P/D = 1$); only 5.89% of the stops has an index < 0.5 confirming the soundness of our configuration.

Still, these cases do occur occasionally. Figure 8 shows the real trajectory of a visitor that stops at a POI, moves away at ≤ 1.5 m from the stop position multiple times (for 2–6 s) then returns, yielding a stationary index of 0.45. Although these cases are a minority, they are accurately accounted for by the distinction between presence and duration.

Stops vs. UWB units. Another validation of the soundness of our stop detection, albeit highly qualitative, is offered by visually comparing the heatmap of UWB units in D_{UWB} (Figure 5) against a similar heatmap for stops in D_S (Figure 9). It is easy to see that the two largely match in terms of the areas with highest density, although the density value is different. The reason is that Figure 5 considers the density of raw UWB units, while Figure 9 considers the density of the stop *positions* derived from UWB units by jointly considering the spatial and temporal dimensions. In any case, we do not exploit this heatmap further because its explanatory power is limited. First, it provides only a spatial dimension; capturing the temporal one would require a separate heatmap, unable to capture the interplay between the two dimensions. Second, this heatmap still suffers from the intrinsic limitation of offering only an *aggregate* view of all stops, rather than the *individual* one necessary to determine, e.g., the ability of a POI to attract and engage visitors. Ascertaining these and other relevant metrics requires an analysis both stop position *and* duration, further captured by the higher-level notions of peeks and visits.

Stops vs. peeks. Among the stops in D_S, 8048 (82.6%) are "promoted" as peeks in D_P; as desirable, visitors tend to stop at POIs rather than away from them. Moreover, the stationary index of peeks is comparable to stops, with mean 0.87 and median 0.93. However, somewhat surprisingly, the mean duration $D = 41.4$ s and presence $P = 35.5$ s of peeks is $\sim$20% lower w.r.t. generic stops in D_S ($D = 51.7$ s, $P = 45.2$ s), meaning stops away from POIs are longer. We take a brief detour to analyze this aspect, which is also the opportunity to showcase the flexibility of our approach.

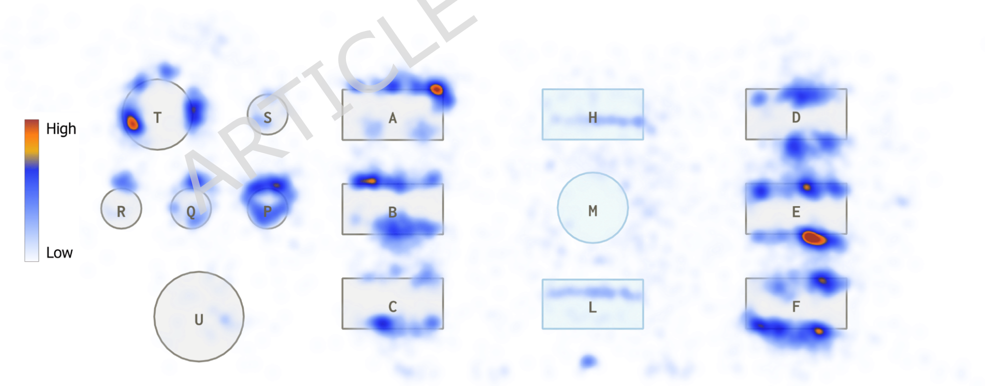

Fig. 9: KDE heatmap of the stops dataset D_S.

Where do stops away from POIs occur? These stops are easily selected by removing from D_S all stops in D_P. We can then process the resulting dataset based on proximity within $R = 1$ m as for peeks, this time considering as POIs elements in the target area *other than* exhibits, namely the benches in H, L, M. Results show that, in 47% of the cases, stops not occurring at POIs occur in one of these elements, accounting for 8% of all stops in D_S. Also, their mean duration is $2'30''$, much longer than peeks. These considerations explain the differences between stops and peeks highlighted above.

At the same time, these results could directly inform the curators on how visitors use benches, e.g., whether they are insufficient or over-provisioned; similarly, results from different POI definitions could inform the safety manager about whether visitors stop in front of emergency exits or fire extinguishers. Although in this paper we focus on the fruition of exhibits by museum visitors, *simply changing the set and definition of POIs caters for different applications.*

Visits and their dynamics. Aggregating the peeks in D_P yields the 6793 visits in D_V, whose mean duration and presence are 49.9 s and 42.0 s, respectively. These values are higher than for peeks, as expected. Still, only 14% of the visits contain multiple peeks and, of these, 77% contains exactly two. This is likely due to the small size of POIs, and explains the modest impact of aggregation on the visit stationary index P_V/D_V. Indeed, for a large POI like `T*1` we find 42% of visits have multiple peeks, on average 1.67 peeks per visit, nicely capturing the peculiar visit pattern w.r.t. other POIs. Interestingly, repeated visits to a POI are rare (3%) and their temporal separation large ($17'16''$ on average).

7.2 Insights: Visitors

We now derive practical insights based on our datasets, starting from a characterization of visitors.

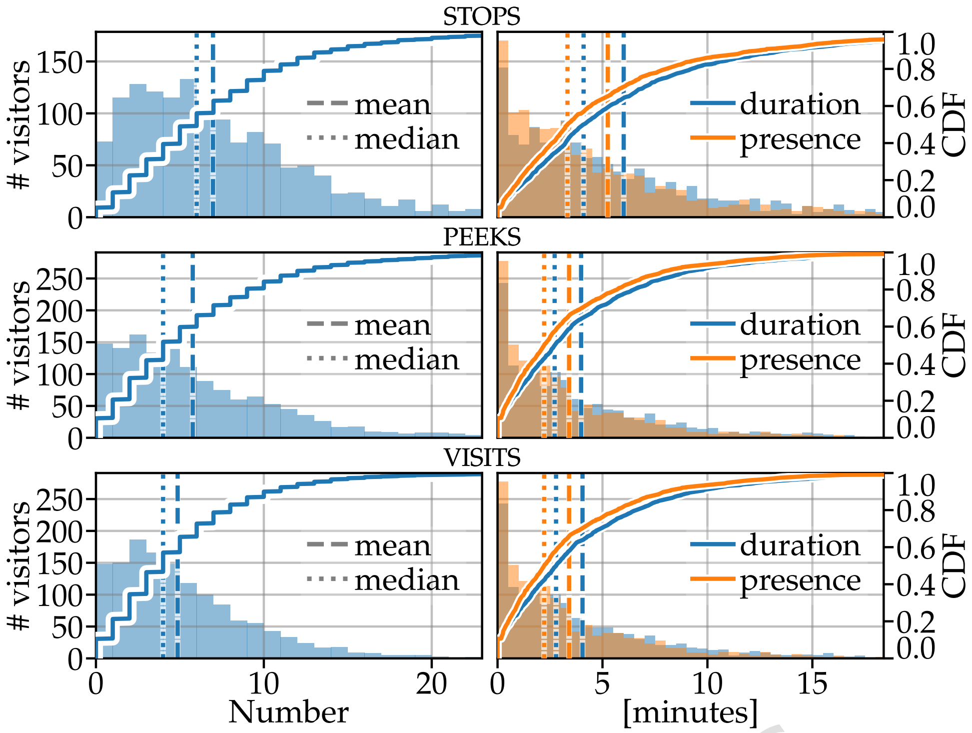

Fig. 10: Number of visitors vs. number of stops, peeks, and visits and their total duration and presence.

How many visitors engage with exhibits? Among the 1323 visitors stopping at least once in the target area, 1248 do so near a POI, thus have at least one peek and visit associated (Table 3); they are 94% of those stopping, 89% of those with a clean trajectory, and 83% of the 1500 visitors carrying a tag.

The above is a concrete example of the magnifying lens visits offer w.r.t. UWB data. Based on the latter, we could have concluded that 93% of visitors engage with at least one exhibit in the target area, i.e., those with a clean trajectory (Table 1): the more accurate estimate based on visits is 10% lower. Similarly, UWB data tells us that 100% of the visitors with a clean trajectory spent time in MAIN (by construction, §4.3) and 85% in SPONSOR; however, the data in D_V tells us that only 86% of the visitors actually visit at least one exhibit in MAIN, and only 55% in SPONSOR.

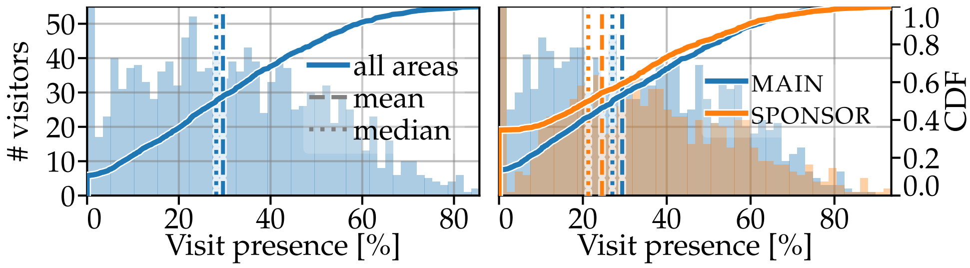

Fig. 11: Visit presence over entire stay in target area.

How long are they engaged? Visitors display wildly different behaviors, immediately visible from the number of stops, peeks, and visits they make and their total duration and presence (Figure 10). To comment our results more concisely hereafter we focus only on total presence, as it is more conservative in capturing visitor engagement.

Interestingly, all distributions are right-skewed, the common case based on knowledge from museum science [12]. The average visitor has 6.97 stops, 5.77 peeks and 4.87 visits, and is engaged with POIs for a total presence of $3'25''$. However, this is skewed by few visitors making many visits and/or remaining at POIs for long, as the median visit presence is just $2'13''$; the 95^{th} percentile of total visit presence is $\text{p95} = 11'10''$, with a maximum of $36'7''$. Moreover, presence varies significantly between MAIN and SPONSOR, coherently with the differences in overall stay observed earlier (§5.2). On average, visitors are present at POIs in MAIN for a total of $2'44''$ ($\text{p95} = 8'55''$) vs. only $48''$ ($\text{p95} = 3'1''$) in SPONSOR.

Again, visit presence serves as a more accurate proxy of visitor engagement w.r.t. the overall stay of visitors in the target areas as computed from UWB data §5.2) or commonly done via technologies offering coarser-grained resolution (§2). On the other hand, the ratio between the two offers a rich characterization of each visitor behavior. On average, visitors spend ∼30% of their time engaged

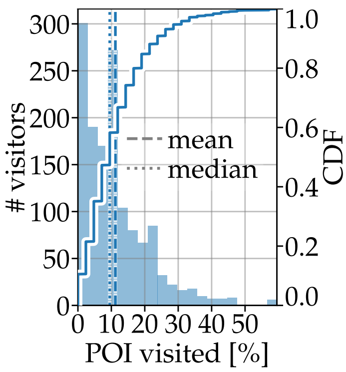

Fig. 12: Fraction of POIs visited.

with POIs, but with a high degree of variability (Figure 11, left), further reasserting that estimating engagement with the overall stay duration would generally yield large errors. Similar considerations hold when analyzing sub-areas (Figure 11, right).

How many POIs do they visit? The fine-grained data in the datasets output by the second step of the toolchain (Figure 2) enables an accurate answer. Interestingly, *none* of the visitors experiences the exhibition in full; visitors engage with 11.2% of POIs on average (median 9.5%), with a maximum of 59.5% (Figure 12).

7.3 Insights: POIs

The question is then which POIs are more "interesting"; we answer it by exploring the data from another, dual perspective, and investigate the *attractiveness* and *engagement* power of POIs based on the number and presence of their visitors.

Which POIs attract the most (or least) visitors? Attractiveness changes significantly across POIs (Figure 13a). `T*1` counts 474 visitors, almost 1 in 3 (31.6%). At the other extreme, `DS1` counts only 32 (2.1%). `T*1` is the largest table in SPONSOR, a rich POI with maps, wooden structures, vertical panels, and screens, while `DS1` is the small white tile visible on the corner of the table in Figure 1. Other highly frequented POIs are `ES3`, hosting a tiger head and ivory sculptures, and `P*1`, showcasing pieces produced by a renowned engineering company. At the lowest ranks we also find `FN1` and `FS1`, both very similar to `DS1`. Overall, the relative attractiveness of POIs matches intuition; yet, visits allow us to *quantify* their differences, and sometimes hold surprises. For example, `ES2`, a large screen cycling through various topics and full of animations, does not attract many visitors; this might be due to the nearby `ES3` capturing their attention. The fine-grained, quantitative lens we offer via visits allows curators to flag similar situation where expectations and data do not align, possibly informing a redesign of the exhibit layout.

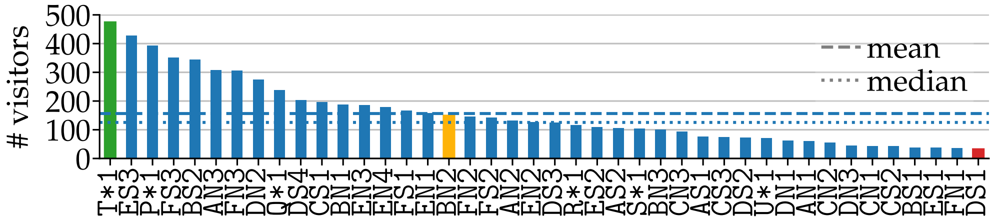

(a) Attractiveness.

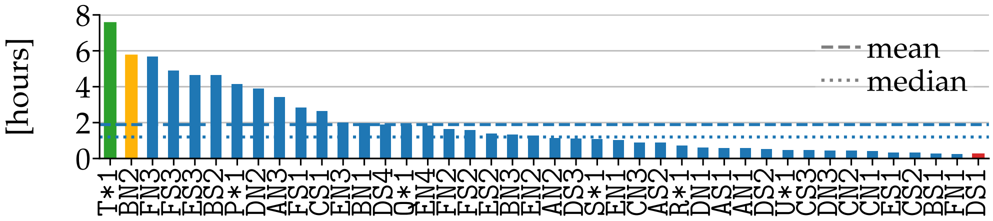

(b) Engagement (by total visit presence).

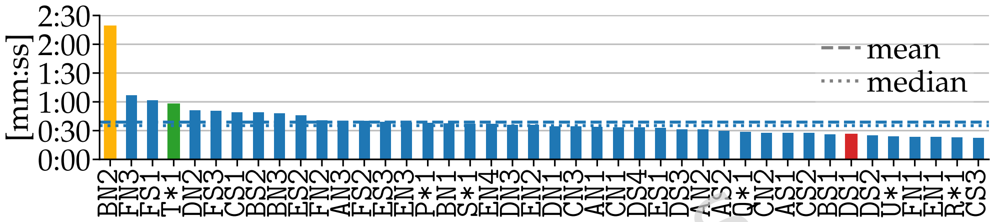

(c) Engagement (by mean visit presence per visitor).

Fig. 13: Ranking of POIs.

What about *groups* of POIs? The "magnification" of this lens can be adapted by curators, who may focus on the attractiveness of *groups* of POIs. A group can be defined by a physical containment relation (e.g., same table), a logical one (e.g., same theme or content type), or a combination thereof, possibly yielding *arbitrary POI hierarchies*. Once the containment relation is properly encoded, it can be straightforwardly exploited by dedicated processing.

As an example, we report an analysis of attractiveness by using a hierarchy defined by physical containment based on *i)* tables, as already encoded in the POI descriptor (Figure 1b), including the single-POI tables `P`–`U` in SPONSOR, and *ii)* and their further clustering in three zones: MAIN-EAST (tables `D`–`F`), MAIN-WEST (tables `A`–`C`), and

SPONSOR. For each of these groups at different layers, we analyze attractiveness as the number of visitors *stopping at one or more POIs* in the group.

Table `E` is the most attractive, visited by 712 visitors, while `C` is the least visited, with only 338. Interestingly, the tables in MAIN-EAST accrue more visitors than those in MAIN-WEST, 1055 vs. 868. This is likely due to the proximity of the former to the *only* two entry accesses, which may also explain the even lower number of visitors in SPONSOR (773), placed even further than MAIN-WEST. Overall, 1206 visitors stop at least once at a POI in *either* MAIN-EAST or MAIN-WEST.

Which POIs engage visitors the longest? Attractiveness provides only a partial view of the ability of a POI to capture the attention of visitors. Indeed, a POI may attract many visitors that leave shortly after, as the content is not very engaging or, dually, attract few visitors that nonetheless remain engaged for a long time. Therefore, understanding engagement requires an analysis of *both* these complementary aspects, shown in Figure 13b and Figure 13c, enabled by our notions of visit (§6.1) and presence (§6.2).

Figure 13b shows all POIs ranked by total presence, i.e., the aggregate of the time spent by *all* visitor at each POI. As with attractiveness, `T*1` is again ranked first; however, it is followed by `BN2` and `FN3`, which ranked only 17^{th} and 7^{th} in attractiveness, respectively. The reason becomes evident when analyzing the *mean* presence per visitor in Figure 13c: `BN2` is $\sim$6 times more engaging than all other POIs, followed by `FN3` whose presence is nonetheless nearly half. A look at the nature of these POIs provides a possible explanation. `BN2` hosts a touch screen for a game targeted to children, explaining both the relatively small number of visitors and the long time spent there. `FN3`, instead, showcases a tray of soil with plants, an animal skeleton, and an interactive screen. Interestingly, both POIs appear as high-density in the UWB heatmap (Figure 5), oblivious of the important differences we just outlined about their fruition in terms of numbers of visitors and time spent at the exhibit.

Finally, we can also analyze engagement using the same POI groups defined earlier; in this case, we compute the total and mean presence over the entire group. Results show a trend similar to attractiveness: the tables in MAIN-EAST accrue 48% of the total presence (68% of MAIN), MAIN-WEST 32%, and SPONSOR only 20%. The same relative ranking holds also for mean presence, with $2'9''$, $1'46''$, and $1'16''$, respectively.

8 Building a Global Picture

The automation and scalability enabled by our approach inevitably yield large datasets, whose exploitation by domain experts demands a concise, actionable view of their content.

We offer two examples: the classification of POIs and visitors into paradigmatic classes (*archetypes*) (§8.1) and an at-a-glance view of the *order* in which visitors stop at POIs. In both cases, our goal is to highlight the endless possibilities unlocked by our first-class concepts (§6.1) directly capturing the relation between visitors and POIs, whose expressive power and quantitative detail cannot be achieved with pure UWB trajectories or coarser-grained technologies.

8.1 Defining and exploiting archetypes

Thus far, we considered each insight and corresponding metric in isolation. However, our discussion of Figure 13 and of specific exhibits like `BN2` shows that the *combination* of metrics, attractiveness and engagement, enables other important insights. We generalize this concept at scale, showing how it be used to elicit archetypes of POIs and visitors.

Defining archetypes. Figure 14a places POIs in a bi-dimensional space defined by attractiveness vs. engagement. We capture archetypes by dividing the chart into *quadrants* based on, respectively, the averages of number of visitors (156) and of mean visit presence (38.9 s). We then assign a descriptive label to each quadrant, resulting

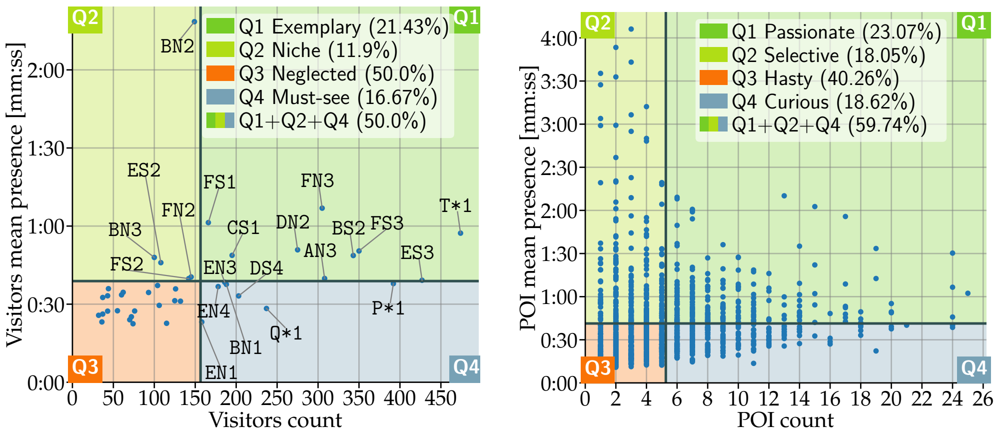

(a) POIs.

(b) Visitors.

Fig. 14: Mapping POIs and visitors into archetypes.

in a simple yet effective characterization of POIs. Quadrant Q1 contains POIs that, from the museum point of view, are *exemplary*, i.e., highly attractive *and* engaging. At the opposite, the POIs in Q3 are neither attractive nor engaging and therefore largely *neglected*. Q2 includes *niche* POIs, attracting fewer-than-average visitors yet engaged for a long time; BN2 falls into this category, albeit barely. Finally, POIs in Q4 appear to be a *must-see*, not highly engaging yet visited by many. Visualizing visitor fruition in this way immediately conveys to the museum curators the relative strength of each POI, facilitating a match against expectations as well as guiding changes to the layout.

The same approach can be applied to visitors. We characterize their attitude when visiting the museum based on the number of POIs visited vs. the mean visit presence across them and define quadrants based on their average values, 4.7 s and 37.0 s, respectively (Figure 14b). Visitors in Q1 are the ideal ones, *passionate* enough to visit a high number of POIs and spend considerable time at each of them. At the other extreme, visitors in Q3 are *hasty* in their visit, covering only a small fraction of the exhibits and for a short time. Visitors in Q2 are instead *selective*, visiting fewer-than-average POIs for a longer-than-average time. Finally, visitors in Q4 are *curious* enough to explore a large fraction of the POIs, without spending too much time at them.

Table 4: Visitor quadrants (%) per day of the week. For readability, the cells in each column follow a color scale from light to dark based on the range of values.

Day	# visitors	Q1	Q2	Q3	Q4	Q1+Q2+Q4
Mon	71	28.2	15.5	36.6	19.7	63.4
Tue	238	26.5	13.9	40.3	19.3	59.7
Wed	263	27.4	20.9	31.2	20.5	68.8
Thu	197	19.8	15.7	40.6	23.9	59.4
Fri	280	21.4	16.8	46.1	15.7	53.9
Sat	154	22.1	22.7	41.6	13.6	58.4
Sun	193	17.6	20.7	44.0	17.6	56.0

Exploiting archetypes. In our examples, we defined quadrants based on mean values. However, alternative boundary values could be set, e.g., based on common knowledge from museum science [12] or the curators' own expectations and goals. In general, based on the metrics above, the goal of the museum is to reduce the visitors and POIs in Q3, and increase them in other quadrants. Understanding which factors impact behavior and preferences is essential; quadrants are a concise and informative way to capture not only the current situation, but also to *compare at different times*, e.g., before and after changes to the POI content or layout.

As an example, Table 4 shows how the distribution of visitors across archetypes shifts w.r.t. the day of the week. Visitors appear significantly less passionate (Q1) and curious (Q4) during the weekend, yet more selective (Q2); this could match the attitude of casual visitors, who added a visit to the museum among several other weekend activities, or by the fact that the museum is typically more crowded during weekends, hampering fruition. On the other hand, the highest fraction of passionate visitors (Q1) is found on Mondays, when the museum is closed to the public but open for private tours, e.g., to schools or other groups. This analysis should be corroborated by longer campaigns; we have only 8 weekends in our dataset. However, it is illustrative of the insights enabled by our approach, which could be amplified if by knowledge about, e.g., special museum activities on given days or visitor provenance and demographics.

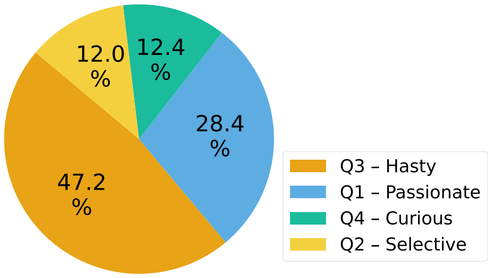

Fig. 15: Distribution of visitor archetypes among bench users.

Another example is offered by the observation made earlier that a non-negligible fraction of stops occurs at benches: we can now determine how these stops distribute across visitor archetypes. To this end, we consider all visitors who visited at least one POI and stopped at least once at a bench. Figure 15 shows a marked predominance of hasty visitors (Q3), accounting for nearly half of bench users. This indicates that, although these visitors spend only a short time on a few POIs, they exploit the wide benches as a rest area or as a place to wait for their companions. Passionate visitors (Q1) represent another substantial portion (28.4%): highly engaged visitors, spending an extended period of time in the target area, likely need occasional resting breaks before resuming their activity. The remaining users are nearly evenly split between curious (Q4) and selective (Q2) visitors.

8.2 Exploring the *order* of visits

Visits abstract UWB positions into a higher-level concept capturing the semantics of the relation between a visitor and a POI. Thus far, we discussed the individual or aggregated characteristics of visits. Nevertheless, in the same way a sequence of UWB units yields a trajectory, a sequence of visits yields a *semantic trajectory* [21, 22] that *directly* captures, at a higher-level, the behavior of a visitor touring the target area, enabling an entirely different set of analyses.

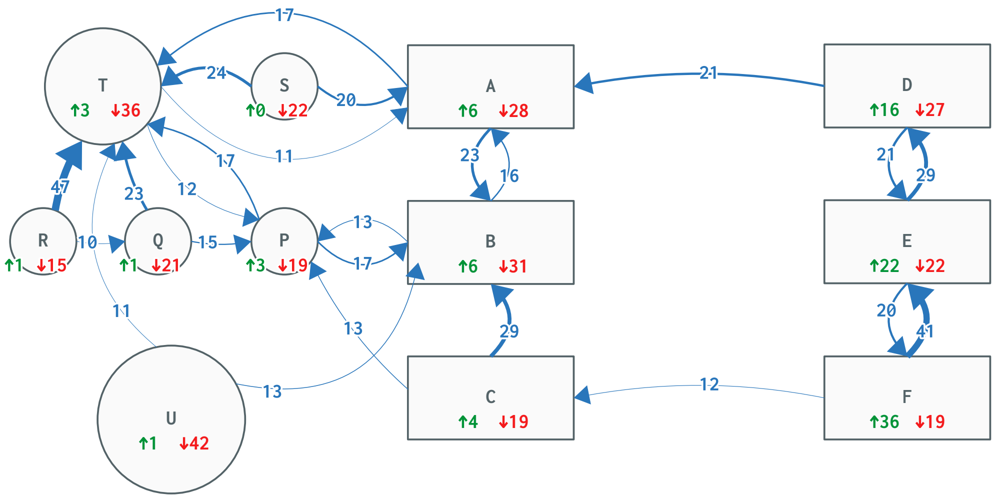

Fig. 16: Visualizing the order of visits. Values are in percentage; green/red ones are entry/exit.

Figure 16 offers an example by showing the probability of a visitor at table t_i visiting table t_{i+1} next, along with the probability that the visitor's tour starts or ends at t_i. Due to readability concerns, we show only the two highest probabilities and limit the analysis to movement across tables, in line with the POI hierarchy defined earlier (§7.3); the extension to individual POIs is straightforward.

The majority of visitors (74%) starts from D (16%), E (22%), or F (36%), coherently with the placement of entry paths and the insights from D_{UWB} (§5.2). Further, F→E (41%) and E→D (29%) are among the highest values; the opposite transitions, starting from D, are also popular. In general, the transitions among these three tables clearly show the tendency of visitors to move among them in an orderly way, before moving the other zones. In contrast, the mobility pattern in MAIN-WEST and SPONSOR is more chaotic, as shown by the low transition probabilities—considering they are the top ones—except for R→T (47%). This is confirmed by computing the entropy of the outgoing transitions of D–F towards other tables, whose value is lower than all others except R. Finally, exit transitions are *always* among the top two most likely, due to the few POIs visited on average; visitors at `U`, `T`, `B` are particularly likely to end their tour there (42%, 35%, 31%).

This analysis, directly useful to curators, could also be exploited to derive or adapt models (e.g., [80]) able to extract sequential patterns and *predict* the probability of a visitor stopping at a POI, in turn enabling *museum digital twins* [33] offering unmatched accuracy and spatio-temporal resolution.

9 Discussion and Outlook

The first structured approach to observing museum visitors dates back to the work of Melton [13] nearly one century ago, and survives today in the common practice of "timing and tracking" popularized by Serrell [12, 25] (§2). Interestingly, in these works the individual duration of *each* POI visit was deemed important, yet it was not recorded since *"Timing visitors at every stop makes tracking more complicated and harder, and it accumulates much larger quantities of data"* [12]. In contrast, state-of-the-art BLE-based systems can automatically track this temporal information. However, due to the inherent inaccuracy of this technology (§2), this is typically achieved with a coarse spatial resolution that is often inferior to human observers, who can easily discern among stops at POIs even when these are densely packed, as in our case.

Our work shows quantitatively, in a real-world setup, that seizing the best of both worlds is possible, enabled by UWB. We offer a fine-grained spatio-temporal stop resolution at the level of individual POIs, yet in an automated and scalable way. Crucially, this achievement is only partly enabled by the superior spatio-temporal quality of UWB trajectories *per se*. The leap from this raw data to semantic trajectories capturing mobility patterns as first-class concepts is necessary to truly unleash the power of UWB for indoor mobility analyses.

In doing so, we followed the footprints of the large body of literature concerned with the extraction of mobility patterns (§2). Nevertheless, the objective of this paper is *not* to propose novel techniques but to show, concretely and at scale, the currently untapped potential offered by the fine-grained analysis enabled by UWB trajectories. In this respect, the very fact that we are using well-known tools and relatively simple techniques and yet obtain a trove of detailed and actionable findings may inspire other researchers to build upon our first-of-a-kind findings towards more sophisticated and customized analyses. For instance, elaborating along the lines of Figure 16, an interesting option is to mine common visit patterns or association rules, e.g., determining that visitors stopping at a given exhibit (or sequence thereof) are then likely to visit another POI in another area. The possibilities are many, offering deeper insights into visitor behavior and informing exhibit placement, especially when it is possible to include demographic information in the analyses.

Finally, we focused on a museum scenario due to the prohibitive administrative, logistic, and deployment effort to replicate it elsewhere. Nevertheless, our findings can in principle be extended to any domain where measuring the presence of an individual near POIs is relevant. An obvious example is retail [5]: the quantitative insights offered by fine-grained detection of stops near products on display, akin to museum exhibits, once combined and informed by consumer research [27] could provide store owners with *actionable* information about their customers. In this respect, the added value offered by the *fine granularity* of the spatio-temporal information gathered and analyzed comes at the cost of installing the UWB infrastructure, which obviously grows with the size of the area covered. On the other hand, this infrastructure can also be deployed *temporarily* only in areas of particular interest and later relocated in others, which is feasible in TALLA and other RTLS systems. In our case, the museum curators, enticed by the possibilities opened by our study, manifested precisely this desire, e.g., to "illuminate" only the areas dedicated to themed, temporary exhibitions; these areas are of particular relevance, yet are placed each time in different zones of the large museum, and therefore could be better catered for by a temporary, relocatable system. The same relocation strategy in principle can be replicated in virtually any environment, significantly reducing costs.

10 Conclusions

UWB and its high spatio-temporal accuracy is posited to revolutionize indoor mobile analytics. Nevertheless, the state of the art has only scratched the surface of what is potentially achievable, focusing on the low-hanging fruits offered by raw UWB trajectories and their visualization. This paper shows that much deeper insights are unleashed once UWB is combined with state-of-the-art techniques from the research communities working on spatial data. To provide concrete evidence for this statement, validate our methodology, and contribute public datasets [24] to the research community, we focus on a museum use case. However, the implications of our study reach far beyond, to the many scenarios where determining *where* and *for how long* people stop is key, whose proliferation is bound to increase as UWB becomes increasingly pervasive on smartphones and personal devices.

References

[1] Zafari, F., Gkelias, A., Leung, K.K.: A Survey of Indoor Localization Systems and Technologies. IEEE Communications Surveys & Tutorials **21**(3), 2568–2599 (2019)

[2] Shams, P., Farahzadi, A., Rezazadeh, J., Bagheri, A.: A Survey on Indoor Positioning Systems for IoT-Based Applications. IEEE Internet of Things Journal **9** (2022)

[3] Leitch, S.G., Ahmed, Q.Z., Abbas, W.B., Hafeez, M., Laziridis, P.I., Sureephong, P., Alade, T.: On Indoor Localization Using WiFi, BLE, UWB, and IMU Technologies. Sensors **23**(20) (2023)

[4] Wikipedia: List of UWB-enabled mobile devices. [Online] https://en.wikipedia.org/wiki/List_of_UWB-enabled_mobile_devices – Last access: 2025-04-15 (2025)

[5] Consortium, F.: Smart Retail. https://www.firaconsortium.org/discover/use-cases/smart-retail. Last access: June 12, 2025 (2025)

[6] Vitali, V., Bazzani, C., Gimigliano, A., Cristani, M., Begalli, D., Menegaz, G.: Trade show visitors and key technological trends: from a literature review to a conceptual framework. Journal of Business & Industrial Marketing **37**(13), 142–166 (2022)

[7] Dane, G., Borgers, A., Kaya, D.I., Feng, T.: Visitor Flows at a Large-Scale Cultural Event: GPS Tracking at Dutch Design Week. ISPRS International Journal of Geo-Information **9**(11) (2020)

[8] Yalowitz, S., Bronnenkant, K.: Timing and tracking: Unlocking visitor behavior. Visitor Studies **12**(1), 47–64 (2009)

[9] Chrysanthi, A., Earl, G.P., Pagi, H.: Visitor movement and tracking techniques. A visitor-sourced methodology for the interpretation of archaeological sites. International Journal of Heritage in the Digital Era **1**(1), 33–37 (2012)

[10] Ferrato, A., Limongelli, C., Mezzini, M., Sansonetti, G.: Using Deep Learning for Collecting Data about Museum Visitor Behavior. Applied Sciences **12**(2) (2022)

[11] Mokbel, M., Sakr, M., Xiong, L., Züfle, A., et al.: Mobility data science: Perspectives and challenges. ACM Transactions on Spatial Algorithms and Systems **10**(2) (2024)

[12] Serrell, B.: Paying Attention: Visitors and Museum Exhibitions. American Association of Museums, Business & Economics, Arlington, VA, USA (1998)

[13] Melton, A.: Problems of Installation in Museums of Art. American Association of Museums (1935)

[14] Kosmopoulos, D., Tzortzi, K.: Visitor Behavior Analysis for an Ancient Greek Technology Exhibition. In: Proc. of Artificial Intelligence Applications and Innovations. AIAI 2021 IFIP WG 12.5 International Workshops (2021)

[15] Spaccapietra, S., Parent, C., Damiani, M.L., de Macedo, J.A., Porto, F., Vangenot, C.: A conceptual view on trajectories. Data & Knowledge Eng. (2008)

[16] Zheng, Y.: Trajectory data mining: An overview. ACM Transactions on Intelligent Systems and Technology **6**(3) (2015)

[17] Damiani, M.L., Issa, H., Fotino, G., Heurich, M., *et al.*: Introducing 'presence' and 'stationarity index' to study partial migration patterns: an application of a spatio-temporal clustering technique. International Journal of Geographical Information Science **30**(5), 907–928 (2016)

[18] Zheng, Y., Zhang, L., Ma, Z., X. Xie, X., *et al.*: Recommending Friends and Locations Based on Individual Location History. ACM Transactions on the Web **5**(1), 44 (2011)

[19] Hachem, F., Vecchia, D., Damiani, M.L., Picco, G.P.: Fine-grained Stop-Move Detection in UWB-based Trajectories. In: Proc. of IEEE International Conference on Pervasive Computing and Communications (PerCom) (2022)

[20] Hachem, F., Vecchia, D., Damiani, M., Picco, G.P.: Fine-Grained Stop-Move Detection with UWB: Quality Metrics and Real-World Evaluation. ACM Trans. on Sensor Networks (TOSN) (2025)

[21] Wang, S., Bao, Z., Culpepper, J.S., Cong, G.: A survey on trajectory data

management, analytics, and learning. ACM Computing Surveys **52**(2), 36 (2021)

[22] Parent, C., Spaccapietra, S., Renso, C., Andrienko, G., *et al.*: Semantic trajectories modeling and analysis. ACM Computing Surveys **45**(4), 32 (2013)

[23] Kontarinis, A., Zeitouni, K., Marinica, C., Vodislav, D., *et al.*: Towards a semantic indoor trajectory model: Application to museum visits. Geoinformatica **25**(2), 311–352 (2021)

[24] Vecchia, D., Hachem, F., Molteni, D., Damiani, M.L., Picco, G.P.: Dataset for "Unleashing the Power of UWB for Indoor Mobility Analytics: A Museum Case Study". Zenodo (2026). https://doi.org/10.5281/zenodo.19681004

[25] Serrell, B.: The aggregation of tracking-and-timing visitor-use data of museum exhibitions for benchmarks of thorough use. Visitor Studies **23**(1), 1–17 (2020)

[26] Cuellar, K., Jambor, E., Sorornejad, P., Talpey, S.: Improving visitor evaluation at the british museum. Technical report, Worcester Polytechnic Institute (2020)

[27] Underhill, P.: Why We Buy: the Science of Shopping. Simon & Schuster, New York (1999)

[28] Li, J., Zheng, X., Watanabe, I., Ochiai, Y.: A systematic review of digital transformation technologies in museum exhibition. Computers in Human Behavior **161**, 108407 (2024)

[29] Ivanov, R., Velkova, V.: Analyzing visitor behavior to enhance personalized experiences in smart museums: A systematic literature review. Computers **14**(5) (2025)

[30] Alletto, S., Cucchiara, R., Del Fiore, G., Mainetti, L., Mighali, V., Patrono, L., Serra, G.: An Indoor Location-Aware System for an IoT-Based Smart Museum.

IEEE Internet of Things Journal **3**(2), 244–253 (2016)

[31] Vena, A., Illanes, I., Alidieres, L., Sorli, B., Perea, F.: Rfid based indoor localization system to analyze visitor behavior in a museum. In: 2021 IEEE International Conference on RFID Technology and Applications (RFID-TA), pp. 183–186 (2021)

[32] Yoshimura, A., Krebs, A., Ratti, C.: Noninvasive Bluetooth monitoring of visitors' length of stay at the Louvre. IEEE Pervasive Computing **16**(2), 26–34 (2017)

[33] Centorrino, P., Corbetta, A., Cristiani, E., Onofri, E.: Managing Crowded Museums: Visitors Flow Masurement, Analysis, Modeling, and Optimization. Journal of Computational Science **53**, 101357 (2021)

[34] Martella, C., Miraglia, A., Frost, J., Cattani, M., *et al.*: Visualizing, Clustering, and Predicting the Behavior of Museum Visitors. Pervasive and Mobile Computing **38**, 430–443 (2017)

[35] Girolami, M., La Rosa, D., Barsocchi, P.: A crowdsensing-based approach for proximity detection in indoor museums with bluetooth tags. Ad Hoc Networks **154**, 103367 (2024)

[36] Yoshimura, Y., Sinatra, R., Krebs, A., Ratti, C.: Analysis of Visitors' Mobility Patterns through Random Walk in the Louvre Museum. Journal of Ambient Intelligence and Humanized Computing, 1–16 (2019)

[37] Pušnik, M., Galun, M., Šumak, B.: Improved Bluetooth Low Energy Sensor Detection for Indoor Localization Services. Sensors **20**(8) (2020)

[38] Nikodem, M., Trajnowicz, G., Blasio, G., Alexis Quesada-Arencibia, F.: Experimental Evaluation of Multicarrier Phase Difference Localization in Bluetooth

Low Energy. IEEE Sensors Journal **25**(1), 1548–1560 (2025)

[39] Spachos, P., Plataniotis, K.N.: Ble beacons for indoor positioning at an interactive iot-based smart museum. IEEE Systems Journal **14**(3), 3483–3493 (2020)

[40] Philippopoulos, P.I., Koutrakis, K.N., Tsafaras, E.D., Papadopoulou, E.G., Sigalas, D., Tselikas, N.D., Ougiaroglou, S., Vassilakis, C.: Cost-efficient rssi-based indoor proximity positioning, for large/complex museum exhibition spaces. Sensors **25**(9) (2025)

[41] Zhuang, Y., Zhang, C., Huai, J., Li, Y., Chen, L., Chen, R.: Bluetooth Localization Technology: Principles, Applications, and Future Trends. IEEE Internet of Things Journal **9**(23), 23506–23524 (2022)

[42] Hu, Y., Qian, F., Yin, Z., Li, Z., *et al.*: Experience: Practical indoor localization for malls. In: Proc. of ACM International Conference on Mobile Computing and Networking (MobiCom) (2022)

[43] Wooley, M.: Bluetooth Direction Finding: A Technical Overview (2019). https://www.bluetooth.com/bluetooth-resources/bluetooth-direction-finding/

[44] Sambu, P., Won, M.: An Experimental Study on Direction Finding of Bluetooth 5.1: Indoor vs Outdoor. In: 2022 IEEE Wireless Communications and Networking Conference (WCNC), pp. 1934–1939 (2022)

[45] Cominelli, M., Patras, P., Gringoli, F.: Dead on Arrival: An Empirical Study of The Bluetooth 5.1 Positioning System. In: Proceedings of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization. WiNTECH '19, pp. 13–20 (2019)

[46] Bluetooth SIG: Bluetooth Channel Sounding. https://www.bluetooth.com/

learn-about-bluetooth/feature-enhancements/channel-sounding. Last access: 2026-02-16

[47] Wieme, J., Nietvelt, R., Leemput, D., Weyn, M., Crombez, P., Berkvens, R., Poorter, E., Hoebeke, J.: Performance Evaluation of Bluetooth Channel Sounding on Commercial Hardware. IEEE Access **13**, 180862–180876 (2025)

[48] nomtek: UWB Art Guide. https://www.nomtek.com/labs-projects/uwb-art-guide/. Last access: June 12, 2025 (2025)

[49] Locatify: Solutions for Museums. https://locatify.com/solutions-for-museums/. Last access: June 12, 2025 (2025)

[50] Qorvo: Bringing the Bible to Life with Ultra-Wideband. https://www.qorvo.com/design-hub/blog/bringing-the-bible-to-life-with-ultra-wideband. Last access: June 12, 2025 (2025)

[51] Chen, Y., Wang, J., Yang, J.: Exploiting anchor links for NLOS combating in UWB localization. ACM Transactions on Sensor Networks (TOSN) **20**(3), 1–22 (2024)

[52] Le, V.A.M., Trobinger, M., Vecchia, D., Picco, G.P.: Human Occlusion in Ultra-wideband Ranging: What Can the Radio Do for You? In: Proc. of International Conference on Mobility, Sensing and Networking (MSN) (2022)

[53] Zhao, W., Goudar, A., Qiao, X., Schoellig, A.P.: UTIL: An ultra-wideband time-difference-of-arrival indoor localization dataset. The International Journal of Robotics Research **43**(10), 1443–1456 (2024)

[54] Lee, B., Choi, J., Jeong, S., Choi, J., Oh, K.-Y.: Comprehensive outdoor UWB dataset: Static and dynamic measurements in LOS/NLOS environments.

Scientific Data **12**(1), 1532 (2025)

[55] Bregar, K.: Indoor UWB Positioning and Position Tracking Data Set. Scientific Data **10**(1) (2023)

[56] Bocus, M.J., Piechocki, R.: A comprehensive ultra-wideband dataset for non-cooperative contextual sensing. Scientific Data **9**(1) (2022)

[57] Cheraghinia, M., Shahid, A., Luchie, S., Gordebeke, G.-J., Caytan, O., Fontaine, J., Herbruggen, B.V., Lemey, S., Poorter, E.D.: A Comprehensive Overview on UWB Radar: Applications, Standards, Signal Processing Techniques, Datasets, Radio Chips, Trends and Future Research Directions. IEEE Communications Surveys & Tutorials **27**(4), 2283–2324 (2025)

[58] Palma, A., Bogorny, V., Kuijpers, B., Alvares, L.O.: A clustering-based approach for discovering interesting places in trajectories. In: Proc. of ACM Symposium on Applied Computing (2008)

[59] Buchin, M., Kruckenberg, H., Kölzsch, A.: Segmenting trajectories by movement states. In: Advances in Spatial Data Handling: Geospatial Dynamics, Geosimulation and Exploratory Visualization, pp. 15–25. Springer, Berlin (2013)

[60] Hwang, S., VanDeMark, C., Dhatt, N., Yalla, S.V., *et al.*: Segmenting human trajectory data by movement states while addressing signal loss and signal noise. International Journal of Geographical Information Science **32**(7), 1391–1412 (2018)

[61] Bermingham, L., Lee, I.: A Probabilistic Stop and Move Classifier for Noisy GPS Trajectories. Data Mining and Knowledge Discovery (2018)

[62] Teng, S., Ku, W., Chuang, K.: Toward mining stop-by behaviors in indoor space.

ACM Transactions on Spatial Algorithms and Systems **3**(2) (2017)

[63] Li, H., Lu, H., Chen, G., Chen, K., *et al.*: Toward translating raw indoor positioning data into mobility semantics. ACM/IMS Transactions on Data Science **1**(4), 1–32 (2020)

[64] Zhou, Y., Chen, Y., Pi, D.: Discovery of stay area in indoor trajectories of moving objects. Expert Systems with Applications **170**, 114501 (2021)

[65] Dedhia, K., Dagli, S., Konkar, M., Katre, N.: A proposed approach to indoor navigation and store recommendation. In: Proc. of Smart Technologies, Communication and Robotics (STCR) (2021)

[66] Marini, G.: Towards indoor localisation analytics for modelling flows of movements. In: Adjunct Proc. of ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proc. of ACM International Symposium on Wearable Computers (UbiComp/ISWC) (2019)

[67] Yaeli, A., Bak, P., Feigenblat, G., Nadler, S., *et al.*: Understanding customer behavior using indoor location analysis and visualization. IBM Journal of Research and Development **58**(5/6), 1–12 (2014)

[68] Hayama, T., Nariai, T., Nagatomo, K.: Extracting stay regions from uwb indoor trajectory and its evaluation. In: International Journal of Service and Knowledge Management, vol. 4, pp. 27–40 (2020)

[69] Martinez-Maldonado, R., Echeverria, V., Schulte, J., Shibani, A., *et al.*: Moodoo: Indoor positioning analytics for characterising classroom teaching. In: Artificial Intelligence in Education. Lecture Notes in Computer Science, vol. 12163 (2020)

[70] Damiani, M.L., Hachem, F., Issa, H., Ranc, N., *et al.*: Cluster-based Trajectory

Segmentation with Local Noise. Data Mining and Knowledge Discovery **32**, 1017–1055 (2018)

[71] Vecchia, D., Corbalán, P., Istomin, T., Picco, G.P.: TALLA: Large-scale TDoA Localization with Ultra-wideband Radios. In: Proc. of International Conference on Indoor Positioning and Indoor Navigation (IPIN) (2019)

[72] IEEE 802.15.4-2015, Standard for Low-Rate Wireless Networks

[73] Qorvo: MDEK1001. https://www.qorvo.com/products/p/MDEK1001. Last access: June 12, 2025. (2025)

[74] DecaWave Ltd.: DW1000 Data Sheet, version 2.19 (2017)

[75] Blom, H.A., Bar-Shalom, Y.: The Interacting Multiple Model Algorithm for Systems with Markovian Switching Coefficients. IEEE Transactions on Automatic Control **33**(8), 780–783 (1988)

[76] Bitgood, S.: An analysis of visitor circulation: Movement patterns and the general value principle. Curator: the Museum Journal **49**(4), 463–475 (2006)

[77] Ester, M., Kriegel, H.P., Sander, J., X.Xu: A Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In: Proc. of ACM International Conference on Knowledge Discovery & Data Mining (KDD) (1996)

[78] Berger-Tal, O., Bar-David, S.: Recursive Movement Patterns: Review and Synthesis Across Species. ESA Ecosphere **6**(9), 1–12 (2015)

[79] Serrell, B.: Paying attention: The duration and allocation of visitors' time in museum exhibitions. Curator: The Museum Journal **40**(2), 108–125 (1997)

[80] Zhang, J.-D., Chow, C.-Y., Li, Y.: LORE: exploiting sequential influence for location recommendations. In: Proc. of ACM International Conference on Advances

in Geographic Information Systems (SIGSPATIAL) (2014)