

User Behaviour in the Wild: Analysing Naturalistic Log Data of a Pedestrian Navigation App

Robert Jackermeier*, Bernd Ludwig

Information Science, University Regensburg, Germany, [firstname.lastname]@ur.de

* Corresponding Author

Abstract:

Knowledge about real-life user behaviour is an important factor for the design of navigation systems. Prompted by the observation that users tend to use our campus navigation app in unexpected ways, we present a naturalistic study of navigation logs. The data set consists of sensor and interaction data from over 4600 sessions, collected over a span of several months from hundreds of users. In our analysis, we demonstrate how the core concepts from navigation literature, i. e. wayfinding and locomotion, can be observed, but also point out differences to previous studies and assumptions. One of our main findings is that the application is mostly used to plan routes in advance, and not to navigate along them. Furthermore, detailed case-studies of actual navigation sessions provide a unique insight into user behaviour and show that persons are often not focused on their navigation task but engaged otherwise. Based on these results, we formulate design implications that do not only apply to future iterations of our application, but can be seen as best practices for pedestrian navigation apps in general.

Keywords: pedestrian navigation, naturalistic study, user behaviour

1. Introduction

Information about a user's context is an essential requirement for any system that aims to provide targeted assistance. In personal navigation systems, determining the exact location is usually of primary interest, but other factors such as user behaviour should be considered just as important. Navigation success depends on whether the system can use this information to decide when and how to present new instructions. As a step towards this goal, in the present study we explore what insights can be gained based on naturalistic log data of our campus navigation app.

For the past 4 years, we have been providing a web-based navigation system for our university campus, with about 50-100 daily visits leading to 15-30 navigation sessions during weekdays, and up to 500 visits during large conferences or when a new semester begins. This underlines the necessity for such a system and indicates that many people are — at least partly — unfamiliar with the environment. In order to provide additional assistance during the whole navigation process, we released the system as a smartphone application, with the main advantage of access to device sensors and functionality that can be used for indoor positioning. Being able to collect more sensor data also allows for additional insights about the behaviour during navigation, which is especially of interest for the present work. Looking at the log files however, we noticed that usage patterns differed from what we had expected based on previous studies, e. g. causing worse positioning performance and aborted navigation sessions.

This leads us to the questions at hand: How is a (pedestrian) navigation system actually used in real life? And how does user behaviour in the wild differ from that in controlled experiments?

In the remainder of the paper, we first look at related work, then briefly describe the app and data collection method. After a gaining an overview of the data set, we highlight several aspects in more detail by means of illustrative examples, followed by a discussion of the results and possibilities for future work.

2. Related Work

Many studies about smartphone use in general employ naturalistic methods, e. g. by investigating the influence of context and interactions with other people (Brown et al., 2014) or by analysing app usage patterns (Böhmer et al., 2011, Jones et al., 2015). To our knowledge however, no studies focusing on smartphone-assisted pedestrian navigation behaviour based on naturalistic log data have been published so far, most likely due to the lack of systems in academic use.

In their study of pedestrian outdoor navigation, Webber et al. observed three different strategy groups, ranging from persons who need constant support to ones that are mostly inattentive (Webber et al., 2012). Test persons were sent on predefined routes, with an experimenter following behind, intervening only if needed. A true naturalistic study was done by Ruggiero et al., who collected a large data set of walking trajectories by means of a camera setup in order to derive mathematical models, but the experiment was confined to a small indoor area (Ruggiero et al., 2018). Finally, Brajdic and Harle investigated the influence of different carrying modes of mobile devices on step detection. In their semi-controlled experiment, they roughly instructed test persons to walk along a path in varying speeds of their choosing, carrying devices in different positions and orientations (Brajdic and Harle, 2013).

In an earlier study of ours (Jackermeier and Ludwig, 2018),

12 test persons were sent on 3 routes (ranging from 151 to 190 metres) each, with the navigation app as guidance and an experimenter following behind. The same setup was used in our latest controlled study, where 100 test persons had to navigate along a route of more than 300 metres across multiple buildings on the campus.

We observe that both in our own studies and in the literature, the experiment setup usually consists of predefined test routes that have to be completed from start to finish, with some help by the experimenter if necessary, e. g. when the test person gets lost (Webber et al., 2012). As the experiments are carried out on special test devices and not on the test persons' personal devices, there is no possibility of interruptions due to received messages or phone calls, and therefore the system under evaluation always stays in the foreground.

Furthermore, we made the experience that test persons, being in an experiment situation, usually pay attention to the navigation instructions and are carrying the device in front of their body, pointing the screen towards them, even without being told to do so. Consequently, they are ready to interact with the system at any time.

In the following sections, we will explore how well these observations and assumptions hold up in a naturalistic scenario.

3. Navigation App and Data Set

The navigation app used in this study was designed for Android smartphones and tablets, providing the same basic functionality as our existing web application. This entails the calculation and presentation of arbitrary routes across the campus, text-based instructions that rely on nearby landmarks, and the possibility to swipe through these instructions and to zoom and pan the map. Since we expected the app to be used as a companion during walks across the campus, an indoor/outdoor localisation system is coupled with the navigation logic. Nearby landmarks can be displayed on demand, with the option to take pictures and add new ones. In case of erroneous localisation, users can correct their position, automatically triggering a re-routing process, if needed. The app's user interface during navigation, with schematic map, route segments, navigation instructions and estimated user position, is shown in Fig. 1. The app is available for free in the Google Play Store.

Users can opt in to anonymously share their usage data for research purposes. Over a period of almost ten months, 4655 sessions from 643 different devices were registered. At the beginning of the semester, the number of daily sessions peaked at over 300 (suggesting that many of our users are students), the regular amount being about 10–20 on weekdays, with very little activity on weekends (see Fig. 2).

In every session, several different kinds of data are collected:

- start, destination, and the route that was calculated by the routing server
- interactions with the user interface (zooming, panning, switching instructions, ...)
- sensor data (detected steps, device orientation and acceleration, Wi-Fi, GPS, ...)

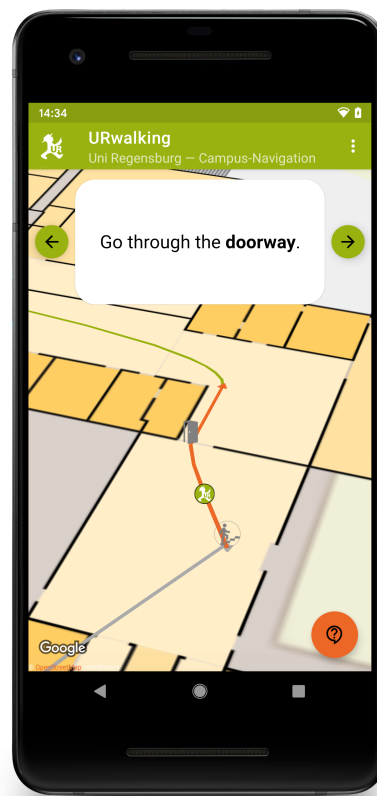


Figure 1. User interface of the app during navigation.

- position during navigation (as estimated by the app's positioning system)
- device information (randomly generated device ID, Android version, available sensors)

A session begins as soon as the navigation view opens, i. e. after the route has been calculated. Since the user's progress keeps being tracked while the app is in the background, sessions continue until it is explicitly closed.

As is common when analysing naturalistic log data, some cleanup was inevitable. All sessions that originated from our own devices, may it be due to development, testing, or controlled experiments, were excluded. Since not all Android devices contain the same types of sensors, they produce different logs, sometimes even with implausible timestamps, which led to difficulties in determining the true session length when the app was closed unexpectedly. For these cases, a heuristic based on several different sensor logs was used.

4. Finding Actual Navigation Sessions

During a quick glance over the log files, it became apparent from different factors such as session length or starting position that many sessions do not contain any navigation on campus, contrarily to the originally intended use of the app. Since the upcoming analysis focuses on sessions where actual navigation happens, those that aren't of interest need to be filtered out first.

4.1 Session Length

The length of a session can be measured in multiple ways. A look at the duration reveals a long-tail distribution (Fig.

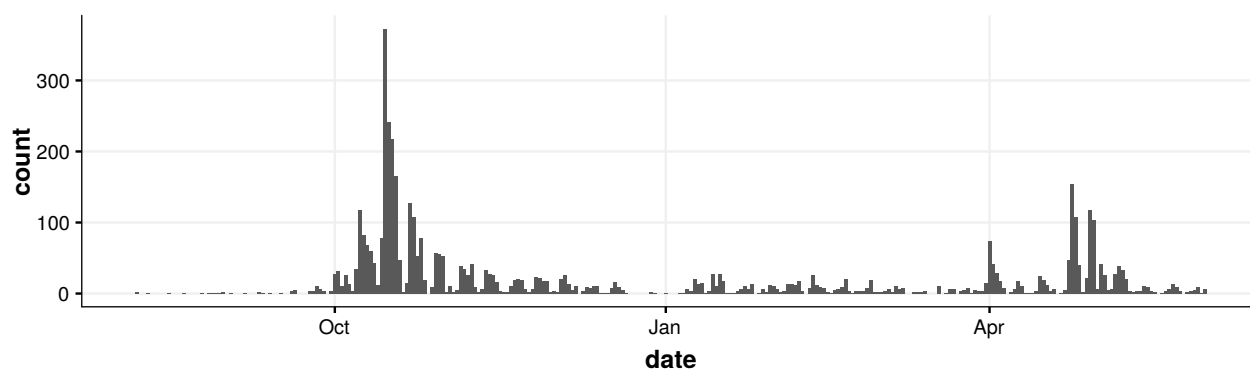


Figure 2. App usage over time. Notice the periodic weekly patterns and the peaks when a new semester starts (in October and — to a lesser extent — in April).

3), where the median session length amounts to just 80 seconds. This result is in line with the observations for general app usage (Böhmer et al., 2011). Since — as mentioned above — a navigation session continues while the app is in the background, defining session length by duration is somewhat problematic. In fact, some of the longer sessions can only be explained by periods of other activities happening inbetween.

A better way to measure session length is by counting the amount of steps a user takes, especially in combination with the chosen route. In a large portion of all sessions (57%), we only observe a small amount of steps (0–10), and therefore cannot consider them to be actual navigation sessions. Between 10 and 100 steps are registered in another 16% of the sessions, while the rest (26%) consists of more than 100 steps. Taking the route into account, we find that — assuming an average step length of 0.7 metres — in 13% of the cases the walked distance equals or exceeds the route length. This, of course, does not necessarily mean that the navigation was finished successfully. Users can optionally tell us in the app that they reached their destination, which happened in 3.8% of the cases. In 22% of all sessions the route (length) is completed at least halfway; in 29% the walked distance amounts to a quarter or more.

We conclude from these observations that many users do not use the app as a companion, and even those who initially intended to follow the route often abort the navigation process after a while.

4.2 Start Positions

The initial user position is determined by Android’s builtin location services, which in turn are leveraging GPS, Wi-Fi and cellular information. While positioning accuracy inside buildings is usually not high enough to pinpoint exact locations, it is definitely possible to distinguish between locations on and off the university campus.

About 46% of all sessions are started on the campus, while 24% originate from somewhere else. In the remaining 30%, no user position was detected, either due to the session being too short, technical difficulties, or system settings preventing access. Routes can only be calculated within the campus borders, therefore well-known locations such as the bus station often serve as manually entered starting points when the actual user location is off-campus. A Wilcoxon rank-sum test indicates that sessions started on

the campus last significantly longer (Mdn = 2.36 min) than those started elsewhere (Mdn = 0.97 min, $W = 1461408$, $p < 0.01$).

In this context, we can also investigate the distances between the initial user position and the starting point of the route. While it is not evident from the log files whether a user opted to use the current position as starting point, the respective coordinates are available. As a heuristic, we can assume that distances below 25 metres are likely due to the “start from current position” feature — the device’s positioning accuracy is usually lower than that inside buildings. This happens in only 17% of all sessions, while in another 8% the distance lies between 25 and 50 metres. In all other cases it is either even higher (45%) or, as described above, could not be determined (31%).

To conclude the overview over the data set, we observed that many users evidently just use the application to look up routes in advance, without planning to engage in actual, guided navigation. Sessions are oftentimes very short, do not originate from where the route begins, or are aborted after a little while.

5. Behaviour during Navigation

For the analysis of user behaviour during actual navigation, the sessions are filtered using the following criteria, based on what we’ve learned until now and considering the size of the campus: A session has to start within the university campus, with a route of 100 metres or more that is completed at least half-way. Furthermore, the session needs to last at least 2 and at most 30 minutes. After applying these restrictions, we are left with 395 sessions from 211 different users/devices, to which all of the following figures and observations relate. In order to illustrate the different aspects of the analysis, we introduce each concept by means of a representative example before applying the method to the data set in general.

5.1 Activity Patterns and Pauses

The first aspect of interest is whether different navigation phases as described in the literature can be observed from our data. Navigation is generally thought of as a combination of wayfinding and locomotion (Montello, 2005). Therefore, one would expect to see short orientation periods, followed by locomotion, maybe with a planning phase at the beginning.

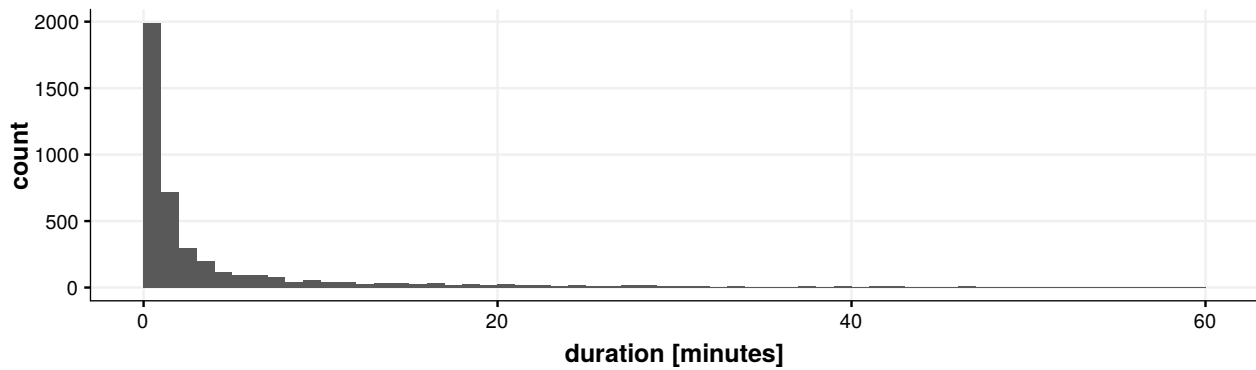


Figure 3. Distribution of session durations. (308 sessions > 60 minutes omitted for clarity.)

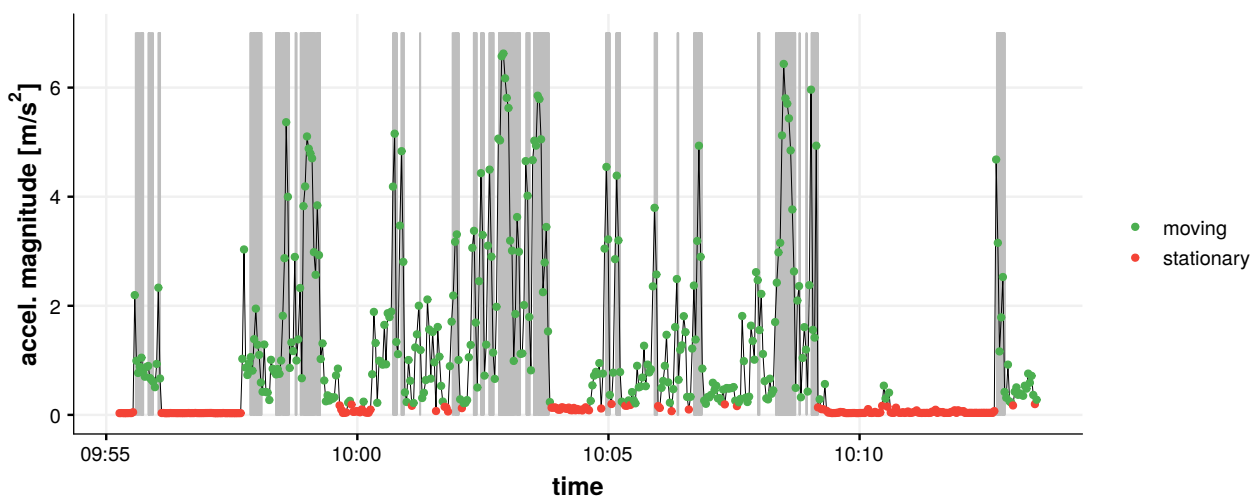


Figure 4. Example for a navigation session with intermittent stationary or inactive periods. Grey areas indicate when steps were detected.

One of the best ways to infer this information is by analysing accelerometer readings in order to find out when the device is almost stationary, meaning that the user either is standing still or put it down somewhere. Fig. 4 shows an example for a session where this behaviour can be observed. After filtering out noise using a moving average with a window size of 200 (corresponding to an interval of 2 seconds at 100 Hz sampling frequency), a threshold of 0.2 m/s^2 is applied to distinguish between a stationary and a moving device, indicated by red and green dots respectively. The flat red line near the beginning suggests that the device is resting completely, while the intermittent periods of low activity later on likely correspond to phases where the user stops and needs to re-orient himself.

Looking at the data set in general, we find that these (almost) stationary periods amount to 15.9% of navigation time across all sessions.

Moments of pausing or hesitation during the navigation where the device is not necessarily at rest can also be identified by looking at the data reported by Android's step detector sensor. Again in Fig. 4, the periods where steps were detected at most five seconds apart are marked in grey, revealing a lot of inactivity (i. e. lack of locomotion) throughout the session, coinciding with but also going beyond the stationary phases mentioned before.

Indeed, when applied to the whole data set, we find that

users spend only 59.3% of their time walking actively, pausing five times per session on average. With a stricter definition where all pauses longer than one second are counted — still significantly longer than a typical step, even for slow gait (Oberg et al., 1993) —, this number further decreases to 49.3% (Mdn = 16 pauses / session).

Until now, we've assumed that a user's attention is exclusively on the navigation system. As mentioned in Section 3 however, navigation sessions and data collection do continue in the background (since e. g. the position has to be tracked during these periods as well). For about half of the sessions, detailed logs about when the app itself is paused, i. e. other apps were used inbetween or the screen was turned off, are available.

Fig. 5 shows an interesting case where the periods during which the app is in the background (marked in grey) correspond to a higher activity (in terms of measured acceleration magnitude). Our log data does not provide information whether another app was used during this time or whether the screen was simply turned off. The latter seems more plausible for this example, with the higher acceleration caused e. g. by a swinging arm or periodic step motion (see next section).

In the data set in general, these background periods occur at least once in 145 out of 208 sessions, with a median

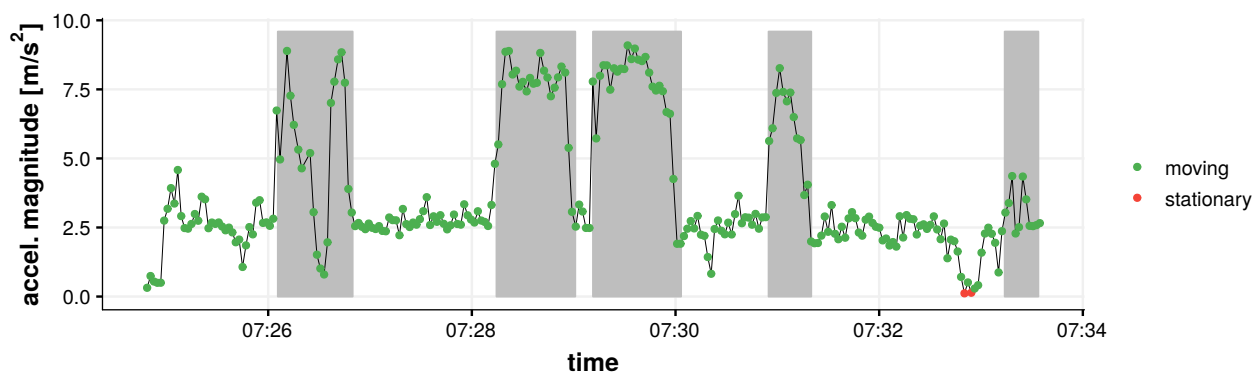


Figure 5. Example for a navigation session where the app is running in the background on several occasions (marked in grey).

amount of 2 per session and a long-tail distributed duration (Mdn = 27 seconds). This means that persons are often either interrupted, e. g. by phone calls or messages, or turn off the device deliberately; the reason for the latter could simply be that they don't need navigation assistance at the moment.

5.2 Device Orientation and Carrying Mode

The fact that users sometimes put the app aside — physically or metaphorically — needs to be explicitly considered in a navigation concept. Instead, our initial assistance strategy relied on the user being able to receive instructions at any time and place.

As mentioned above, based on what we had observed in previous studies, we expected users to hold the device in front of their body during navigation, pointing roughly in walking direction. This assumption can be tested by looking at the device orientation, measured by Android's Game Rotation Vector sensor, which fuses data from accelerometer and gyroscope.

While without ground truth it is not possible to make any statements about whether the compass direction coincides with the walking direction, we can infer the current carrying mode by looking at the angle in Z direction, i. e. whether the device is held horizontally or pointed upwards or downwards. In contrast to e. g. Brajdic and Harle, who investigate six different modes or positions (Brajdic and Harle, 2013), we only distinguish between two cases that can be recognised relatively easily from unlabelled data:

In what we'll call active usage mode — in the literature referred to as “texting mode” (Klein et al., 2018) — the device needs to be about horizontally or tilted upwards, or specifically somewhere in the range of -20 and $+60$ degrees that accounts for most of the values in controlled experiments where the phone was always held that way. Contrarily, we call it passive mode when the device is pointing downwards at an angle below -20 degrees, which usually means it is swinging besides the body, or pointing almost straight up or down, indicating that it was put away into a pocket. For our analysis, the exact carrying position (front/back pocket, bag, ...) does not matter. Instead, the two modes reflect whether the user can immediately react to new navigation instructions (and therefore use the app in its intended way) or not.

Fig. 6 illustrates the observations on a real example. As above, a moving average is applied that helps smoothing

the signal in case of large variance, e. g. when the device is swinging in an outstretched arm, which otherwise might lead to many mode changes. The device is initially held almost horizontally, then after a while pointed straight down (red dots, passive mode), most likely put into a pants pocket, before being pulled out again. Over the course of the navigation process, this procedure is repeated several times.

When the same session is visualised in all three dimensions, additional patterns emerge: The graph shown in Fig. 7 can be thought of as the trace left behind by an imaginary laser pointer attached to the device along the local Y axis, i. e. parallel to the long side of the screen, while the device is located at the origin of the coordinate system. Clearly visible are the clusters near the middle (with z approximately between 0 and 0.5) and at the bottom, corresponding to active and passive mode. Whenever a mode change occurs, a line can be seen connecting the two clusters. Note that we use the unfiltered data set here, allowing us to identify repeating patterns near the bottom, created by periodic walking motion with the device in the pocket.

Applying this analysis to all navigation sessions, we observe that active mode accounts for 71.8% of all time. However, in just about a quarter (26.0%) of the sessions, the device is held in active mode over the course of the whole navigation process. In all remaining sessions at least one mode change happens, with a median amount of 4, occurring from very early on in many sessions, peaking roughly at the two-minute mark.

Fig. 8 shows the device orientation distribution across all navigation sessions. Notice the sharp peaks at -1 and 0 , corresponding to the device pointing straight down (presumably in a pocket) and laying flat, respectively. Besides that, users mostly opt to hold the device angled slightly upwards.

In comparison, in our latest controlled study (mentioned in Section 2), the overall time in active mode was over 99%, and in 77 out of the 100 runs, the device was always held that way. The median amount of mode changes across all runs therefore is zero. Test persons were given the same app that is publicly available, the only difference being that they had to explicitly confirm the completion of each navigation step.

In our earlier study (Jackermeier and Ludwig, 2018), test persons were not required to interact at all; in fact, the ex-

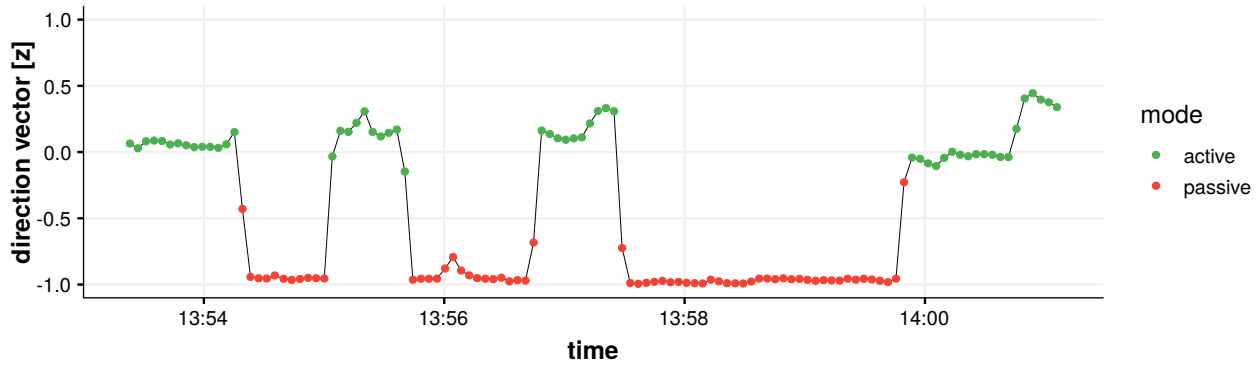


Figure 6. Example for a navigation session with multiple mode changes, suggesting that the user is either familiar with or able to partially memorise parts of the route.

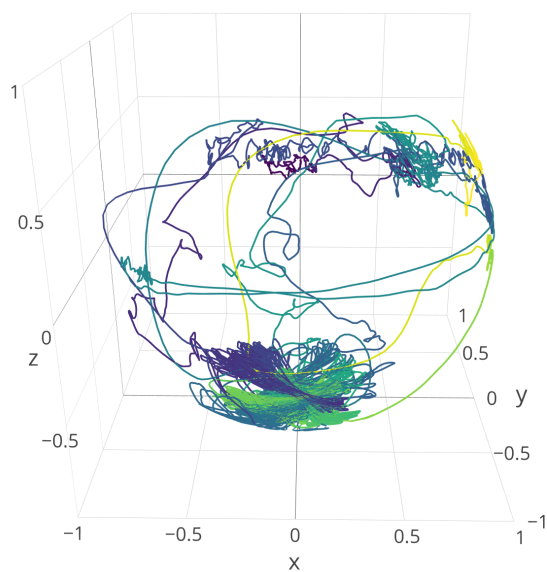


Figure 7. The same session as in Fig. 6, visualised in 3D. The coloured trace on the unit sphere indicates device orientation over time.

perimeter was remotely controlling the app, thereby guaranteeing the test person to be always provided with actionable and valid navigation instructions. But still, all of them opted to keep the device in active mode the whole time, consequently no mode changes were detected at all.

In contrast to these experiments, the naturalistic data presented so far proves that the assumption of a user who can interact with the navigation app at any time does not hold up in reality.

5.3 Explicit Interactions

Until now, the analysis focused on the sensor capabilities a native application provides. Additional insight into user behaviour can be gained by the explicit interactions with the app's interface, i. e. when the user pans or zooms the map, switches between instructions manually etc.

By default, the map follows the user position, keeping the current route section centred. Instructions change automatically when a new route section is entered, but can also be switched manually by either pressing the buttons next to them or by swiping them left or right (see Fig. 1).

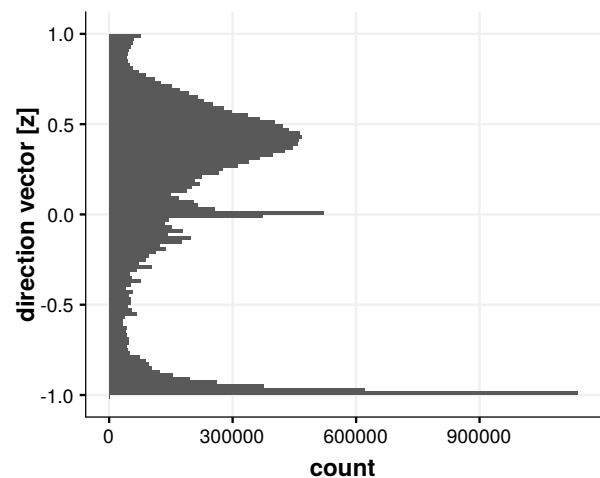


Figure 8. Device orientation across all navigation sessions.

Fig. 9 revisits the example from section 5.1, this time additionally displaying interaction data. Map interactions (“pan” and “zoom”) tend to occur in clusters, while instructions are switched throughout the session, in this case only ever forward and never back to a previous one. Combined with acceleration data and information about the application state (grey background periods), this completes the picture about when wayfinding and locomotion occurs.

Looking at the whole data set, panning or zooming the map happens at least once in a majority of navigation sessions (89.6%), with a median of 18.5 of those interactions. Instructions are changed significantly less frequently in only 52.6% of all navigation sessions (Mdn = 8.5 interactions / session). It is worth noting that 9 out of 10 times, the instruction is moved forward, suggesting that users generally want to plan ahead.

In terms of additional interactions, in 1 out of 10 sessions a nearby landmark is clicked on in order to display a picture that provides aid during wayfinding. Other options such as landmark upload, proactive position correction or the feedback functionality occur even less frequently.

6. Discussion and Future Work

The analysis has shown that the main concepts known from the literature can be identified in our naturalistic data set.

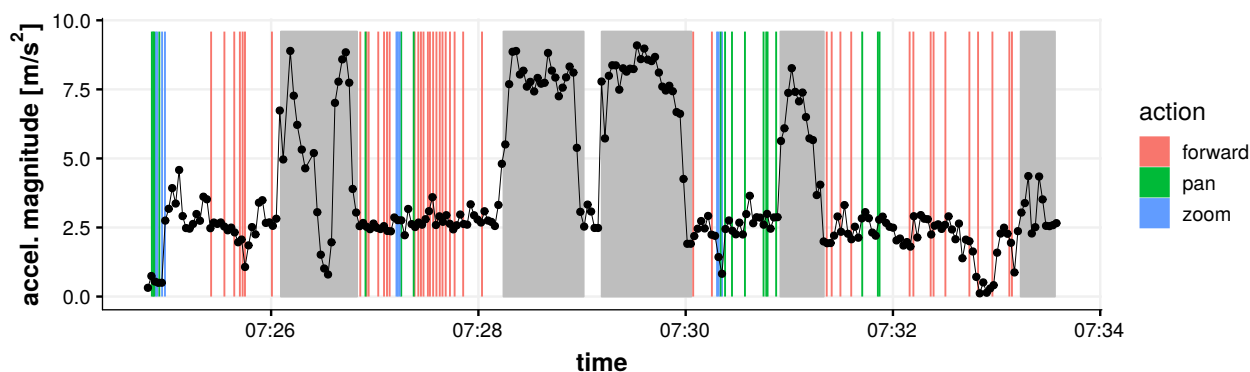


Figure 9. The same session as in Fig. 5, additionally displaying interactions with the user interface. Background periods are marked in grey.

On the one hand, we observed users slowing down and interacting with the app for wayfinding, while at other times the device is put away and the focus lies solely on locomotion.

Some of our initial assumptions about the general use of a navigation application have to be revised however: Firstly, users do not primarily use it to actually navigate, but more often to plan or look up routes, sometimes while not even being located near their intended starting point. Design recommendations for mobile pedestrian navigation have been formulated before in the literature, usually for the case where a user is already following a route (May et al., 2003, Delikostidis et al., 2016). We argue that one has to take a step back and make sure that a navigation app is well suited to support the steps leading up to the actual navigation process, e. g. by including dedicated modes for both strategies. Also, in our specific case, some features such as position tracking or compass calibration are unnecessary unless navigation takes place, and taking pictures of nearby landmarks requires the user to actually be on site.

Secondly, during wayfinding most users hold or carry the device in multiple ways throughout the session, sometimes putting it away in their pocket. This has consequences for the design of the navigation instructions as well as the positioning system:

It is self-evident that graphical or text-based instructions can only be effective if the user actually pays attention to the screen; in other cases, additional information channels such as auditory or haptic feedback need to be added. With knowledge about the current device position, this decision can be made automatically, thereby only using the more intrusive methods when required. Furthermore, when it can be inferred that the user does not need assistance to reach a certain intermediate goal (e. g. because he or she put the phone in the pocket), the system can remember this information and adapt future routes and navigation instructions accordingly.

To provide proper instructions at any time, accurate positioning is required. The strategy of using device orientation as a proxy for walking direction is only reasonable when the device is carried in a specific way, which — as we know now — is less likely than initially thought. Therefore, techniques such as Principal Component Analysis that have been proposed in the literature (Combettes and Renaudin, 2015, Deng et al., 2015) should be used to

determine a user’s walking direction independently of the device orientation.

A limitation of the present study is that user behaviour is analysed entirely based on quantitative data. The incorporation of many different sensors, combined with interaction data and knowledge about the environment, does help in painting a coherent picture; but further inquiry is needed to find the reasons for certain observations, e. g. the information needs that lead users to open the app in the first place, or why most navigation processes are aborted prematurely.

In the subset of actual navigation sessions we were looking at, most persons do make use of their possibilities to interact with the map and instructions, even though there theoretically is no need. While without qualitative data we cannot determine the exact reasons, certain explanations make sense: An incorrectly estimated position or direction can cause difficulties to match the displayed map to the real environment and create the need for interaction. The same holds true for unclear or untimely navigation instructions. Since the majority of instructions were switched forward, the practice of previewing the next one in advance is one possible improvement.

From a design perspective, the goal should be to eliminate these “unnecessary” interactions since they do not provide reliable information about the user state. Instead, explicit communication about the navigation state along the lines of “I do not understand this instruction” or “I don’t think my position is displayed correctly” could enable a more targeted assistance.

The results presented in this naturalistic study of real navigation data are based on a snapshot of the log data generated by our app several months into its public availability. Data collection is still ongoing, with more users expected at the start of each semester. There are many aspects left to be explored, e. g. temporal and spatial patterns, or the identification of user groups with different information needs.

References

Böhmer, M., Hecht, B., Schöning, J., Krüger, A. and Bauer, G., 2011. Falling Asleep with Angry Birds, Facebook and Kindle: A Large Scale Study on Mobile Application Usage. In: *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile*

Devices and Services, MobileHCI '11, ACM, New York, NY, USA, pp. 47–56.

Brajdic, A. and Harle, R., 2013. Walk Detection and Step Counting on Unconstrained Smartphones. *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '13*. ISBN: 9781450317702.

Brown, B., McGregor, M. and McMillan, D., 2014. 100 Days of iPhone Use: Understanding the Details of Mobile Device Use. In: *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services - MobileHCI '14*, ACM Press, Toronto, ON, Canada, pp. 223–232.

Combettes, C. and Renaudin, V., 2015. Comparison of Misalignment Estimation Techniques between Handheld Device and Walking Directions. In: *2015 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, IEEE, Banff, AB, Canada.

Delikostidis, I., van Elzakker, C. P. and Kraak, M.-J., 2016. Overcoming challenges in developing more usable pedestrian navigation systems. *Cartography and Geographic Information Science* 43(3), pp. 189–207.

Deng, Z.-A., Wang, G., Hu, Y. and Wu, D., 2015. Heading Estimation for Indoor Pedestrian Navigation Using a Smartphone in the Pocket. *Sensors* 15(9), pp. 21518–21536.

Jackermeier, R. and Ludwig, B., 2018. Exploring the Limits of PDR-based Indoor Localisation Systems under Realistic Conditions. *Journal of Location Based Services* 12(3-4), pp. 231–272.

Jones, S. L., Ferreira, D., Hosio, S., Goncalves, J. and Kostakos, V., 2015. Revisitation Analysis of Smartphone App Use. In: *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '15*, ACM Press, Osaka, Japan, pp. 1197–1208.

Klein, I., Solaz, Y. and Ohayon, G., 2018. Pedestrian Dead Reckoning With Smartphone Mode Recognition. *IEEE Sensors Journal* 18(18), pp. 7577–7584.

May, A. J., Ross, T., Bayer, S. H. and Tarkiainen, M. J., 2003. Pedestrian navigation aids: information requirements and design implications. *Personal and Ubiquitous Computing* 7(6), pp. 331–338.

Montello, D. R., 2005. *Navigation*. Cambridge Handbooks in Psychology, Cambridge University Press, p. 257–294.

Oberg, T., Karsznia, A. and Oberg, K., 1993. Basic Gait Parameters: Reference Data for Normal Subjects, 10-79 Years of Age. *Journal of Rehabilitation Research and Development* 30, pp. 14.

Ruggiero, L., Charitha, D., Xiang, S. and Lucia, B., 2018. Investigating Pedestrian Navigation in Indoor Open Space Environments using Big Data. *Applied Mathematical Modelling* 62, pp. 499–509.

Webber, E., Burnett, G. and Morley, J., 2012. Pedestrian Navigation with a Mobile Device: Strategy Use and Environmental Learning. In: *Proceedings of the 26th Annual BCS Interaction Specialist Group Conference on People and Computers*, British Computer Society, pp. 286–291.