

Automated Recording of Human Movement Using an Artificial Intelligence Identification and Mapping System

David Barbarash, Moiz Rasheed, Apoorva Gupta, Tong Wang

Purdue University, Indiana/USA · dbarbara@purdue.edu

Abstract: Few designers engage in post-occupancy evaluation of built works and those that do have limited tools available for capturing robust and unbiased data. This paper proposes a system in development to record continuous and unbiased site usage inventories with high accuracy and rich metadata for further analysis. To date the system can use multiple cameras with time synced video feeds to re-identify people in the landscape and project their locations onto a two-dimensional map. The unique multi-camera system avoids accuracy and identification issues inherent in single camera systems while recording unbiased data without privacy concerns. The inventory programming system is discussed, and preliminary outputs and data are presented along with issues encountered and future goals for the project.

Keywords: Post-occupancy evaluation, behavioral mapping, artificial intelligence, site inventory

1 Introduction

Designers often use case study to analyze interventions in the built environment during conceptual design phases. Design firms, however, rarely investigate the outcomes of their own built works. One reason is that post-occupancy evaluation (POE) and associated behavioral mapping (ZIMRING & REIZENSTEIN 1980) is resource intensive. Despite the costs, it is critical for design firms to invest resources into POE so the profession at large can realize the impacts of design on environmental, human structural, and socio-cultural processes and flows. As few firms perform POE or publish their findings, there is a lack of quantitative data describing how elements move through and occupy an environment, along with a lack of understanding in how a proposed intervention affects larger systems beyond project property lines.

The current project uses artificial intelligence and multi-camera-based computer vision that records and maps human use of defined spaces in the built environment. This represents the first steps in a proposed multi-stage inventory and data gathering toolset.

1.1 What is Wrong with Traditional POE?

Post-occupancy evaluation and behavioral mapping are methods whereby a trained observer spends time within a built space and records what they see using various coding and marking systems on a scaled base map. These observations have been key factors in design inspiration and analysis for decades. The modern incarnation of behavioral mapping began with Wilkel and Sasanoff's studies on interior architecture (WINKEL & SASANOFF 1966) followed by Jan Gehl's urban streetscape studies in Denmark (GEHL 1968, GEHL & SVARRE 2013). These efforts fostered collaboration with non-design professions that investigate how the built environment impacts human behavior, opposite of previous approaches (ITTELSON et al. 1970). These studies showed that observing how people use the built environment is a strong correlate with the social and physical success of designed space. Behavioral mapping as a design

practice continued to develop, incorporating recording the traces people leave behind when using a space (CLAY 1980, SANOFF & COATES 1971), and was presented through the landmark video and book; *The Social Life of Small Urban Spaces*, by William Whyte (WHYTE 1980).

Behavioral mapping theories were developed into a series of generalized guidelines by Claire Cooper Marcus and Carolyn Francis (MARCUS & FRANCIS 1997a, 1997b). Post-occupancy evaluation and its associated behavioral mapping went digital by the turn of the 21st century, though methods still relied on human input of visually observed activity, replacing pen and paper with a geo-referenced digital map (BAHILLO et al. 2017, COSCO et al. 2010, GOLIČNIK & NIKŠIĆ 2009, GOLIČNIK & THOMPSON 2010, MARUŠIĆ & MARUŠIĆ 2012). While POE has contributed greatly to design methods and theories across multiple scales, the process itself is inherently flawed due to human limitations, recording error, and both intentional and unintentional biases (BARBARASH & LU 2020), limitations that are discussed below.

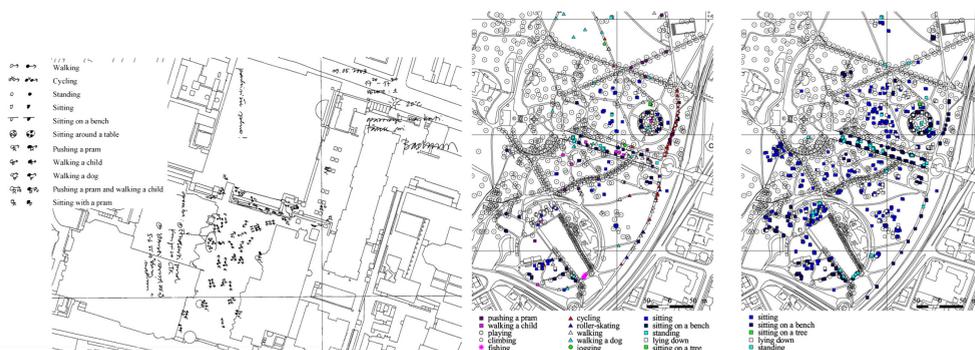


Fig. 1: Traditional and digital manual behavioral mapping outputs (GOLIČNIK & THOMPSON 2010)

Non-automated POE methods rely on a designer’s physical presence in situ or utilizes capture methods that fail to provide uninterrupted views of site activity. Firms rarely include POE in their contract budgets or scope of work because it is a time-consuming process that captures incomplete snapshots of site usage. The physical act of recording, whether performed digitally or via pen and paper, disassociates an observer from the happenings of the current moment, meaning that a person is always trying to “catch up” with what they see. The simple fact that a recorder must look away from the site under study to record observations results in missed behaviors and incomplete data. Time-lapse photography or crowdsourced video feeds results in missed site areas, occlusion, and incomplete movement data.

A trained observer cannot cover the entirety of a site themselves, nor can they remain in place without pause for hours, days, or longer. Observers performing POE are typically recording “snapshots” of site usage, using limited timeframe data to inform multi-process systems that fail to take various events and environments into account (i. e., festivals and scheduled events, weather, day/night cycles, and time of day variations). Viewing a single moment in time does not give insight into the richness of possible behaviors, uses, and patterns that the

built environment can encompass. Making site specific design decisions based on these limited viewsheds and timeframes sets a dangerous precedent that can result in system inefficiencies, safety and security concerns, and issues in access or population inequities.

Human judgement is commonly affected by errors in judgement and bias; whether intentional or not (KRUGLANSKI & AJZEN 1983, WILSON et al. 2017). The impacts of these errors and biases can result in inaccurate data or in extreme cases, data that is purposefully incorrect, leading to decisions that can greatly affect design decisions in negative ways. These decisions can then cascade from site design and planning projects to intersecting systems, resulting in dramatic fallout from poor validity data or biased assumptions and agendas. Issues of data bias exist in digital products as well (SYMONS & ALVARADO 2016) though the construct of this study is designed to minimize or remove such concerns.

Visual evaluation of larger or complex site configurations requires multiple people and may result in data recording inconsistencies (double counting, overlooked areas, differences in interpretation, etc.). Differences in recording speed, visual acuity, and related cognitive processing abilities also impact traditional visual recording efficacy.

The design firm SWA/Balsley used a single fixed position camera to record short videos of small plaza spaces in New York City (SCHLICKMAN & DOMLESKY 2019, SWA 2018). These videos were then processed by artificial intelligence (AI) and machine learning software to inventory occupation counts and locations and to highlight movement corridors. Since the firm was analyzing spaces that they created, they were able to link spatial programmatic intent with actual use, generating valuable POE data for their specific site design. Results from the study fostered a series of general design guidelines for public spaces. It is unknown whether these findings are specific to New York cultures and patterns, whether they can apply across various seasons and weather events, or if they would be successful under different programmatic systems or requirements. The establishment of this style of POE analysis should be tested across various locales and conditions on a more granular scale.

Agent based modelling has been demonstrated to be a viable means towards understanding use of the built environment, with evolving innovations in artificial intelligence and simulation engines increasing realism and accuracy. (ALMAHMOOD & SKOV-PETERSEN 2020) The simulation of people in an as-yet unbuilt project requires a robust behavioral database that can only be built through a multi-cultural, multi-climactic, and unbiased inventory system.

1.2 Complexities of POE Inventory Data

In many existing studies of POE, people are treated as particles in fluid dynamics models which simulates herd behaviours instead of allowing for individual differences and desires. Additionally, the built environment is designed with experiential programming in mind, with spaces for people to move, linger, gather, and rest. The social structure of a site is further complicated by entries and exits from off-site and from architecture, attractors (e. g., art, signage) and repulsors (e. g., trash dumpsters, back alleys), and dedicated movement lanes (e. g., separate pedestrian and bicycle lanes). These realities of the site-scale built environment make planning-scale dynamic studies inappropriate, or even misleading, for analysing site-scale location spaces in the built environment.

A data collection toolset that addresses issues with traditional human recording methods while embedding rich metadata into the process can help designers inventory existing spaces

and use patterns. It can also record trends, behaviours, and events as they occur over time. An example with impacts at both the micro- and macro-scale is the unplanned for introduction of e-scooters to urban environments and the resulting disruption to existing systems and flows. Automated POE can be used to identify human and vehicular behavioral shifts that result from disruptive technologies and trends to inform design recommendations and guidelines. This project aims to address such inventory and analysis issues.

2 Methods and Preliminary Outputs

The system under development and described in this paper differs from traditional POE processes and other existing digital tools in its ability to record unbiased data without pause, across an entire area of interest, without gaps in site coverage. It leverages a variety of computer vision and mapping techniques to fully automate the POE process, using a fixed multi-camera network to cover an observable space for unbroken periods of time. Using multiple cameras to observe a defined space allows for higher confidence in recorded data due to removal of visual occlusion (BARBARASH, 2020), distance and pixel density issues, or variations in lighting impacting identification when compared with single camera systems. Fixed cameras are preferable to drones or handheld video as camera movement creates severe slowdowns in video analysis, as homography matrices can not be re-used across frames.

The main challenge of developing a robust and data-rich site inventory system is the ability to identify (initial recording of a person) and re-identify (recording of a previously identified person) unique persons within the network and then transforming their pixel coordinates into geographical coordinates. The re-identification process poses a difficult problem for real world applications compared with single camera setups, as the system must compare camera outputs to decide whether a detected person is a novel ID (LENG 2020).

The users of a space in the built environment represent an undefined population, with the possibilities consisting of every person on earth. Even then the sample set remains open as identification is affected by daily clothing choices, modes of transportation, and other variables. The difficulty of open-set re-ID can be circumvented by capturing a person's likeness and saving it in a database (referred to as the "gallery") with a unique ID as they cross predetermined entry thresholds (BARBARASH 2020). This effectively removes the open-set problem, taking advantage of more developed technologies in the field of closed-set re-id, where all the IDs to be assigned are known ahead of time.

Open set issues coupled with the limited resolution of cameras in large exterior spaces forces human identification methods to rely on larger visible features instead of other forms of location-based service (LBS) data, facial or gait ID methods, or other data fusion network tracking. The proposed ID process does not know who a specific person is; instead, it uses pixel data to develop a unique code identifier that is only valid while a person remains within the observable site. Upon reaching a set interval of time after an ID is not found within a scene, that ID and associated gallery data are removed from the working set. If the same person were to re-enter the site later or time, they would receive a new ID code as there would no longer be a positive gallery match, removing long-term tracking and privacy concerns.

Software code is written using various Python languages and libraries and follows a multi-stage and recursive flow series of steps as illustrated in Figure 03. The re-ID process uses

YOLO v3 and Faster R-CNN to identify people and objects, generate their bounding boxes, and assign ID codes to an output .csv file. After each camera's homography matrix (a formula calculated to correct perspective distortion created by associating camera elements with real world coordinates) is set up by the end user, mapping and projection reads the output file for pixel x,y coordinates and translates them through Google Maps API, appending lat/long locations to the .csv for the highest confidence location for each ID in each time frame.

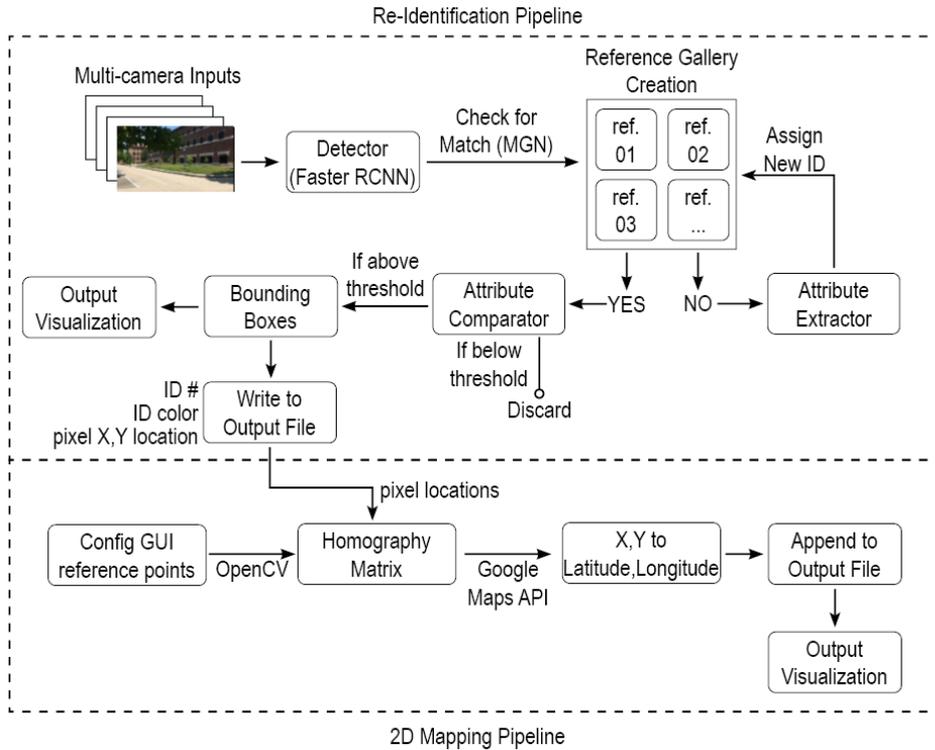


Fig. 2: The re-ID and 2D mapping system pipeline

The system runs using an artificial intelligence driven series of machine learning algorithms that must be trained to prove it will output viable data when processing unknown entities. It uses the Market1501 dataset to train and test the re-identification process (ZHENG et al. 2015) as it contains multi-camera examples of pre-cropped people and has a strong support network of users and projects. These example images are introduced to the system and outputs are compared against expected data described in the Market1501 documentation. If results are like that of the reference Market1501 data set, then the system can be considered as operating correctly and can begin to process new data. The machine learning a system is capable of is only as good as the training dataset provided through the system. New training sets and even more importantly, iterative and evolving sets are critical to the viability of these algorithms.

To identify people in the built environment, the system operates in two phases custom coded for this project. In the first phase, a triggering point, line, or area is created for each camera

view at places where the user determines things will enter frame. Once a change in RGB pixel values is detected at a trigger threshold, the detector process tries to find people or objects and draw a bounding box around them as shown in Figure 2. This avoids false positives from elements within a scene that would otherwise confuse the detection process (such as statuary, odd lighting contrasts, or other unpredictable elements). The pixel content is saved to the gallery as a new individual ID. If an object is detected away from the trigger area(s), bounding box pixel content is saved as a query for re-identification.

In the second phase, the query boxes are cropped out and run through a feature extractor to develop an attribute extraction score which becomes the primary means of re-identifying elements in scene. Cropped images marked as queries are compared against features in the gallery to generate a similarity score (WANG 2018) to decide which existing ID to assign. To facilitate re-identification, a multi-granularity network (MGN) model averages visible multidimensional features within a bounding box. This is compared to images already in the gallery to determine if an object or individual is unique or has already been identified (and therefore assigned an existing ID).



Fig. 3: A single camera frame showing trigger lines for new IDs (green), newly ID'ed individuals to be added to the gallery (cyan), and people re-ID'ed from the gallery (red)

The gallery is purged at regular intervals to avoid prolonged storage of images and the reported features are the only tracked element from the system, removing privacy issues as people are no longer directly associated with their numerical identifier. This process happens for each camera in isolation (See Table 01 for a single timeframe analysis), and then outputs are compared for ID conflicts and confidence scores to determine ultimate ID assignments.

No two objects in a scene can have the same ID number, which means that in some cases identification is more of a “best-fit” process than an absolute determination. The camera view that generates the highest confidence match score for a specific ID in a single timeframe

(above a user-set threshold) writes the pixel location and ID to the output file while the other camera's data is recoded (for testing and evaluation purposes) but rejected. This additional complexity is critical in addressing issues inherent in single camera systems such as occlusion, lighting inconsistencies, and partial identification (BARBARASH 2020).

Table 1: Sample output showing frame, ID codes (re-id and track_id), x,y coordinates in two separate planes, and lat/long coordinates (rounded here for legibility)

line	frame	re-id	track_id	x1	y1	x2	y2	lat	long
8710	1822	3	214	1043	291	1059	325	40.430951	-86.916567
8718	1822	5	219	2410	767	2475	936	40.431302	-86.916276
8716	1822	6	12	963	520	985	600	40.431162	-86.916372
8712	1822	10	200	644	650	676	714	40.431263	-86.916331
8722	1822	13	221	1820	799	1876	987	40.431276	-86.916298

Outliers are ignored by the system if their attribute extraction score falls below a user determined threshold to maintain high data validity. While this means that some IDs may disappear for a given analyzed timeframe, accuracy is not compromised as frame sampling rates remain high enough to smooth out any omissions over time. Additionally, safeguards are placed within the re-ID module so that no object can move more than a certain number of pixels (as a percentage of the total frame resolution) between timestamps to avoid large jumps in AI determined ground truth location for unique IDs.

To date the multi-camera system has been able to operate with a 97% success rate in identifying individuals across a variety of times and weather conditions, validated through random frame spot checks by the research team. Longer multi-camera video data has proven difficult to gather during COVID-affected times, but new public/private partnerships are being established to solve this issue.

The 2D mapping and projection system primarily implements two modules, OpenCV and Google Maps. OpenCV is used to perform homography transformation, which requires a user to perform some initial setup for each camera before projection can begin. This setup involves selecting pairs of reference points from two different perspectives to be matched between known points on Earth (via Google Maps) and visible points in an image (through recorded camera frames). OpenCV then generates a corresponding homography matrix that can be used to convert x,y pixel locations into lat/long coordinates. The Google Maps API is then used to generate a detailed satellite image of the site under study, along with plotted location points for identified people at various times (see Figure 4).

The OpenCV module for coordinate translation solves the mathematical details of homography transformation though this can potentially be a drawback as well. The module does not provide process feedback, so observation of empirical output data must provide insight into troubleshooting needs, though errors typically stem from user inaccuracy or choosing undesirable reference points. For best results, reference points should be as far apart as possible and prioritize the "top" of an image despite these points being further away from the camera.

An advantage of the Google Map API is that it can output a satellite image, which is more detailed than API's like OpenStreetMaps, which only output a map with landmarks. The disadvantage is that Google does not currently have a native Python API for Google Maps and the function of the Python wrapping in use is limited.



Fig. 4: Google API projection output for a three-camera study on Purdue University's campus. Output data can also be viewed through GIS software for analysis.

3 Discussion

Observational inventory data of this type is critical for designers and planners if they are to make informed design decisions in an increasingly complex built environment. The goal of the current project is to generate a toolset that can provide current and robust patterns of use inventories; the impacts of having actionable data can improve efficiencies in movement and flows through and around transportation networks and site designs.

Analysis of the data can contribute to optimization of spaces for site programming specific gathering, resting, meeting, or recreating, while basing decisions around local social, micro-climatic, and structural conditions. Disruptive technologies and design interventions can be observed in real time, allowing for fast adjustment or response planning in future projects and renovations. The data can provide rich insight into evolutions of human interaction with

the built environment and within social systems, fostering new understandings of sociology, environmental psychology, network theory, and design standards. Designers can test hypotheses about site use and behavioral changes that design interventions might foster.

Lessons learned from automated POE systems can save firms and municipalities time and money as deeper understandings of human use of space can lead to more efficient and successful designs. Site programming and related socio-cultural benefits of improved design can then foster healthier communities and more efficient networks and maintenance practices.

The proposed POE tool has use cases beyond understanding existing patterns in the built environment. If enough data can be captured to build a comprehensive list of multi-cultural and multi-environmental human behavioral traits and tendencies, then AI-driven actors can be programmed with random behavioral attributes from the dataset, placed into a 3D model of a not yet built intervention, and simulated use of a space can be evaluated. This could save designers time, effort, and money in avoiding programmatic and structural mistakes in their design work while refining performance-based data to be presented to clients and governmental agencies. An AI behavioral testing system could simulate the richness of the built environment and associated human behaviors over time, adding to the discoveries of single variable or limited scope agent-based explorations of today.

There are issues and barriers in adopting this AI and camera-based POE system in both physical access and data security realms. Design firms and their clients, and municipalities, will need incentives to install fixed cameras in new and renovation projects to generate data towards understanding and projecting user safety, maintenance, and access equity demands. These cameras and their associated data storage will require both electrical and internet access, resulting in higher operating costs outside of initial installation budgets. Placement of cameras should be unobtrusive to avoid impacting people's behaviors, and field of view cannot be blocked by signage or vegetation. As the addition of a new camera to a system seems to increase processing complexity by more than a fixed increment, maximum study area size and camera density are yet to be determined. The system under development is written to be device agnostic so as faster distributed and onboard processing options become available it can be deployed with minimal configuration and troubleshooting.

Data ownership and access is an area with little legal or historical precedent. While the American 4th Amendment and subsequent Supreme Court case *Katz vs. United States* ("Katz v. United States," 1967) established the concept of "reasonable expectations of privacy" in the public realm, the general population remains wary of being observed in their daily lives (POWER 2016). This issue is exacerbated by the lack of oversight and management precedent over who has the right to see recorded data of public spaces, how that data is stored, and what rights camera owners have in displaying and reproducing image-based content. The lack of public trust in such systems may make adoption of automated inventory processes difficult to include in design contracts and masterplans. Instead, these decisions have been left to software developers and end users of recorded data (DE SOUZA E SILVA & FRITH 2010, EDWARDS & URQUHART 2016). If the data systems are private, can public entities access or subpoena content? How often should visual data be purged from the system to maintain public anonymity? What rights do system owners have as off-site development occurs, impacting fields of view and data validity? What protections can be put into place to avoid cross-linking of visual data with deep meta-data generated from smartphone and similar devices via LBS data?

Real time data capture, analysis, and output is not possible with current consumer grade technologies, making AI and camera driven systems more experimental than readily applicable. Multi-camera systems require more processing power than a simple multiple of single camera setups. The largest bottleneck lies in the re-identification pipeline and comparing results across cameras. As a result, higher complexity and larger scale sites require larger processing demands. It is not known what the most efficient system for complete site camera coverage should be, as every site has its own unique spatial organization and fields of view. Studies in ideal camera resolution, ID success and confidence as distance increases, appropriate 2D mapping points for generating homography matrices, and camera overlap requirements are currently underway with hopes of improving both speed and data validity.

Benchmarking for speed and data validity thresholds continues as video is processed through the system. At present the team has rough guidelines and recommendations towards camera site-coverage and maximum expected ID distance. Recommended sample frame rate and installation angles are going through processing trials to be published in future works.

The system is comprised of plug and play software modules and libraries that should be relatively quick to exchange as new technologies and opportunities arise. The system is designed purely as an inventory system. Few necessary changes are predicted in the short term, other than adding new modes of transport or additional object identifiers made visible via higher resolution camera vision data. Before robust data can be gathered by POE systems, work needs to be done to permanently install cameras in the built environment along with associated data transmission, storage, and processing costs and methods.

4 Conclusion

The multiple-camera and AI system described here can identify human figures through video camera vision without recording personally identifiable information beyond what is visible within the public realm. Locations of identified individuals are then projected onto a map along with time-stamped metadata so that location, movement speed, and directionality can be interpreted. Identification of non-human elements (cars, bicycles, animals, shopping bags, etc.) is under development to add breadth and depth to the behavioral database. Additionally, self-sustaining micro-environmental sensors are being distributed throughout areas under investigation to determine climatic impacts on human behaviors and movement.

These efforts have resulted in an effective automated inventory system that generates rich metadata for analysis via GIS and parametric modelling applications. Processing efficiencies, data validity, and additional data targets (non-human objects) are constantly being refined, tested, and added to the system in hopes of generating a complete open course product that is ready for deployment by both public and private entities in the near future.

References

- ALMAHMOOD, M. & SKOV-PETERSEN, H. (2020), Public Space Public Life 2.0: Agent-based Pedestrian Simulation as a Dynamic Visualisation of Social Life in Urban Spaces. *Journal of Digital Landscape Architecture*, 5-2020. <https://doi.org/10.14627/537690032>.

- BAHILLO, A., GOLIČNIK MARUŠIĆ, B., MARUŠIĆ, D. & DÍEZ, L. E. (2017), A dependents' cooperative location network for behaviour analysis in public spaces. In: ZAMMIT, A. & KENNA, T. (Eds.), *Enhancing Places Through Technology*. Lisbon: Edições Universitárias Lusófonas.
- BARBARASH, D. & LU, Y.-H. (2020), Artificial intelligence systems for automated site analytics and design performance evaluation. *Landscape Research Record*, 9, 192-203.
- CLAY, G. (1980), *Close-up: How to read the American city*. University of Chicago Press.
- COSCO, N. G., MOORE, R. C. & ISLAM, M. Z. (2010), Behavior mapping: a method for linking preschool physical activity and outdoor design. *Medicine & Science in Sports & Exercise*, 42 (3), 513-519.
- DE SOUZA E SILVA, A. & FRITH, J. (2010), Locational Privacy in Public Spaces: Media Discourses on Location-Aware Mobile Technologies. *Communication, Culture and Critique*, 3 (4), 503-525. <https://doi.org/10.1111/j.1753-9137.2010.01083.x>
- EDWARDS, L. & URQUHART, L. (2016), Privacy in public spaces: what expectations of privacy do we have in social media intelligence? *International Journal of Law and Information Technology*, 24 (3), 279-310. <https://doi.org/10.1093/ijlit/eaw007>.
- GEHL, J. (1968), *People on Foot*. In (pp. 429-446): *Arkitekten*.
- GEHL, J. & SVARRE, B. (2013), *How to study public life*. Island press.
- GOLIČNIK, B. & NIKŠIĆ, M. (2009), Geographic information system behavioral and cognitive mapping in the city centre revitalisation process. *Journal of Urban Regeneration & Renewal*, 3 (2), 161-179.
- GOLIČNIK, B. & THOMPSON, C. W. (2010), Emerging relationships between design and use of urban park spaces. *Landscape and Urban Planning*, 94 (1), 38-53.
- ITTELSON, W. H., PROSHANSKY, H. M., RIVLIN, L. G. & WINKEL, G. H. (1970), The use of behavioral maps in environmental psychology.
- KATZ V. UNITED STATES, 389 U.S. 347 (U.S. Supreme Court 1967).
- KRUGLANSKI, A. W. & AJZEN, I. (1983), Bias and error in human judgment. *European Journal of Social Psychology*, 13 (1), 1-44. <https://doi.org/https://doi.org/10.1002/ejsp.2420130102>.
- LENG, Q., YE, M. & TIAN, Q. (2020), A Survey of Open-World Person Re-Identification. *IEEE Transactions on Circuits and Systems for Video Technology*, 30 (4), 1092-1108. <https://doi.org/10.1109/TCSVT.2019.2898940>.
- MARCUS, C. C. & FRANCIS, C. (1997a), *Design Guidelines for Urban Open Space*. In: John Wiley & Son Inc, New York.
- MARCUS, C. C. & FRANCIS, C. (1997b), *People places: design guidelines for urban open space*. John Wiley & Sons.
- MARUŠIĆ, B. G. & MARUŠIĆ, D. (2012), Behavioural maps and GIS in place evaluation and design. *Application of geographic information systems*, 115-138.
- POWER, D. J. (2016), "Big Brother" can watch us. *Journal of Decision Systems*, 25(sup1), 578-588. <https://doi.org/10.1080/12460125.2016.1187420>.
- SANOFF, H. & COATES, G. (1971), Behavioral mapping: An ecological analysis of activities in a residential setting. *International Journal of Environmental Studies*, 2 (1-4), 227-235.
- SCHLICKMAN, E. & DOMLESKY, A. (2019), *Field Guide to Life in Urban Plazas*. SWA. https://live-swa-2019.pantheonsit.e.io/wp-content/uploads/2019/08/Field-Guide-to-Life-in-Urban-Plazas_digital1.pdf (January 13, 2020).
- SWA (2018), *Plaza Life Revisited*. <https://www.swagroup.com/idea/plaza-life-revisited/> (13.01.2020).

- SYMONS, J. & ALVARADO, R. (2016), Can we trust Big Data? Applying philosophy of science to software. *Big Data & Society*, 3 (2), 2053951716664747. <https://doi.org/10.1177/2053951716664747>.
- WANG, G., YUAN, Y., CHEN, X., LI, J. & ZHOU, X. (2018), Learning Discriminative Features with Multiple Granularities for Person Re-Identification. Proceedings of the 26th ACM international conference on Multimedia, Seoul, Republic of Korea. <https://doi.org/10.1145/3240508.3240552>.
- WHYTE, W. H. (1980), The social life of small urban spaces.
- WILSON, J. P., HUGENBERG, K. & RULE, N. O. (2017), Racial bias in judgments of physical size and formidability: From size to threat. *Journal of personality and social psychology*, 113(1), 59.
- WINKEL, G. H. & SASANOFF, R. (1966), An approach to an objective analysis of behavior in architectural space.
- ZHENG, L., SHEN, L., TIAN, L., WANG, S., WANG, J. & TIAN, Q. (2015), Scalable Person Re-Identification: A Benchmark. Proceedings of the IEEE International Conference on Computer Vision (ICCV), Santiago, Chile.
- ZIMRING, C. M. & REIZENSTEIN, J. E. (1980), Post-Occupancy Evaluation: An Overview. *Environment and Behavior*, 12 (4), 429-450. <https://doi.org/10.1177/0013916580124002>.