

Next Generation Physical Analytics for Digital Signage

Mateusz Mikusz, Anastasios Noulas, Nigel Davies, Sarah Clinch and Adrian Friday
School of Computing & Communications

Lancaster University, Lancaster, UK

m.mikusz | a.noulas | n.a.davies | s.clinch | a.friday @ lancaster.ac.uk

ABSTRACT

Traditional digital signage analytics are based on a display-centric view of the world, reporting data on the content shown augmented with frequency of views and possibly classification of the audience demographics. What these systems are unable to provide, are insights into viewers' overall experience of content. This is problematic if we want to understand where, for example, to place content in a network of physically distributed digital signs to optimise content exposure. In this paper we propose a new approach that combines mobility simulations with comprehensive signage analytics data to provide viewer-centric physical analytics. Our approach enables us to ask questions of the analytics from the viewer's perspective for the first time, including estimating the exposure of different user groups to specific content across the entire signage network. We describe a proof of concept implementation that demonstrates the feasibility of our approach, and provide an overview of potential applications and analytics reports.

Keywords

Pervasive displays; analytics; simulation

1. INTRODUCTION

Web analytics have helped drive forward significant innovation in the design and deployment of web content. Using sophisticated analytics site owners are able to understand how visitors navigate their site and can design controlled experiments to measure the effectiveness of changes made to the site. In contrast, the owners of physical digital signage networks must rely on relatively crude analytics tools that typically provide a record of the content that is shown on displays and, in the case of video analytics systems, data on the number and demographics of viewers of content items and signs (e.g. [11]).

The key weakness of existing digital signage analytics systems is that they focus on reporting data relating to a sign and its audience (analogous to early website hit counters)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WPA'16, June 25-30, 2016, Singapore, Singapore

© 2016 ACM. ISBN 978-1-4503-4328-2/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2935651.2935658>

yet fail to capture and report data based on a viewer's perspective of the content. While reporting the number of times content is shown, or the number of viewers in front of a display when a content item is shown, is interesting in terms of advertising penetration and impression counting, much more important in terms of improving user experience is the ability to report on the content that an individual viewer sees as they walk through a physical display environment. For example, understanding that viewers see the same content multiple times on different displays, or miss important pieces of content, can help improve content scheduling and placement. Equally, knowing the order in which viewers see content may help provide insights into particular viewer behaviours (e.g. are viewers more likely to purchase items based on the content they saw most recently when entering a shop). Generating such viewer-centric analytics is beyond existing sign analytics systems.

In this paper we report on the design, development and evaluation of a novel approach to physical sign analytics that can help provide viewer-centric analytics data. Our prototype system combines traditional sign analytics data with simulations of viewer mobility to generate *synthetic traces* of viewer sightings of content. As a result we are able to offer sign owners an approximation of viewer-centric analytics without requiring any additional hardware or software within the signage system, or invading users' privacy by tracking and identifying them. Overall, we make three important contributions:

1. we propose a novel approach for generating and reporting viewer-centric analytics,
2. we describe the design and first implementation of a system that demonstrates the viability of this approach, and
3. we present results obtained from our system based on over 18 months of real-world sign analytics data.

We believe this approach has the potential to revolutionise the field of sign analytics—providing the owners and operators of display networks with new tools that enable them to optimise their content schedules for the benefits of viewers.

2. EXPERIMENTAL CONTEXT

Our analytics work is conducted in the context of a large-scale signage deployment at Lancaster University. The campus public display deployment consists of over 50 displays and is one of the world's largest signage research testbeds.

Displays are located in both academic buildings (such as lecture theatres, the university reception and student services) and colleges, and typically show a mix of local content (e.g. departmental or college news) and university-wide news (e.g. produced by the university press office). Displays and content are managed by a web-based application that allows display owners to create individual content schedules [3, 8].

The e-Campus deployment uses PHEME [13] for capturing and processing traditional signage analytics data. In particular, the PHEME system keeps track of each content change including the display identifier, content identifier (consisting of the URI to the content item and its file hash), and the current display power status. The PHEME datastore consists of a detailed log of every content item shown on displays on campus during the past 18 months (over 30 million entries resulting in over 20 GB of data).

Potential viewers of content shown on the signage network are chiefly the 12,965 students (9,235 undergrads and 3,735 postgrads), 1,490 academic and 3,025 administrative members of staff who are based on campus. Most of the students live on campus in one of nine colleges. These residential buildings are physically interleaved with departmental buildings and lecture theatres. Students and staff who live off campus typically arrive in the centre of the campus at a large underground bus stop and walk to their destination.

3. SUPPORTING NEW FORMS OF PHYSICAL ANALYTICS

The content that a viewer actually sees ultimately depends on their mobility pattern and whether they are physically proximate to a given display showing a piece of content. Thus, being able to describe viewer mobility and proximity to displays is critical for the next generation of viewer-centric analytics. However, acquiring such detailed mobility data through user tracking can be costly and represents a potentially significant privacy violation. Instead, we propose the use of simulation to generate synthetic traces of viewer mobility and combining these with data produced by a lightweight conventional sign analytics system that provides information on content shown on displays within the network.

3.1 Mobility Simulation

To support the generation of synthetic traces, we designed and built a prototype simulation system. Our goal is to simulate synthetic mobility traces and compare content visibility given movement and behaviour patterns of different groups of people. We currently model three major groups of viewers to enable us to illustrate how different mobility patterns can influence the analytics results (such as viewer and content exposure across the signage system):

- **On-Campus Student.** Models a person who lives on the university campus and moves between random colleges, departments and lecture theatres once an hour.
- **Off-campus Student.** This model aims to represent members of the campus population who live off-campus and arrive by bus at the bus station located in the centre of campus, and subsequently exhibit similar movement patterns to On-Campus Students.
- **Random Building Navigator.** This model represents theoretical viewers who constantly move between

randomly chosen buildings throughout the simulation. The model is important as it explores a wide spectrum of mobility scenarios without the limitation of fixed schedules and start and destination patterns.

To perform the actual simulation and generation of synthetic traces we need a model of our campus’ spatial layout. This includes key navigation elements such as specific types of buildings, pathways on which people are allowed to move, entrance and exit points, etc. Taking a map as input, our simulation system is able to generate a graph-based representation of the spatial layout in a semi-automated process. For our specific use case in the signage domain, the graph representation further consists of all display locations and their individual *view probability* values for any given location—for example, some displays are more likely to be seen over others due to having a more prominent location.

For the generation of synthetic traces, the simulator creates a pre-defined number of person instances based on our mobility models. Each instance contains a unique identifier and an individual trace history of locations visited. In each step person instances are moved along the graph and the simulator can perform a range of tests—for example, if a display is in proximity, its view probability value is used to determine whether the person instance glanced at the display. Each glance is saved to an individual person instance allowing us to build up a very detailed display view trace.

In order to generate new kinds of insights and analytics reports based on elements of real-world data, the simulation framework combines individual location traces with our historical dataset of actual display content logs. To facilitate this, we created an API to PHEME allowing us to easily determine which content items were shown on a specific display for a given date and time. We are able to create datasets that describe for each simulated person every content impression they would have viewed together with an associated timestamp, display name and content identifier (URI).

3.2 Implementation

Our simulator is written entirely in Python. To represent the spatial layout of the campus, we automatically extracted features from a map of campus into a graph data structure provided by NetworkX. The coded campus map (2,119x5,122 pixels) resulted in 11,225 nodes (6,827 nodes representing buildings, 4,398 nodes representing paths) and 29,539 edges. Each node translates to a square on a map of about $5m^2$ in area. The described mobility models were implemented as individual classes each with their own movement algorithm. The 18 month historical content dataset from PHEME (display, content item, time shown, duration etc.) was transferred into a local MySQL database for performance reasons, resulting in a database of about 30 million entries.

4. EXAMPLES OF NEW INSIGHTS GAINED

We ran the simulation with 2,000 viewer instances of each of the three classes resulting in a total of 6,000 people for a duration of 62 days¹. By combining our mobility simulation data with content impressions from PHEME, we can show

¹Corresponding with a two month period (01 Oct - 01 Dec inclusive) in which students were active on campus and for which we had relevant analytics data.

the potential of our approach to firstly reconstruct complex sign analytics insights (e.g. those traditionally produced by video analytics) and, secondly, provide novel viewer-oriented physical analytics that describe how the display network as a whole, and its associated content, are experienced from the perspective of a viewer/passers-by. This final class of viewer-experience analytics goes beyond the capabilities of systems in current use.

4.1 Effectiveness of displays

Combining content presentation data with the mobility simulation offers the potential to report the reach or effectiveness of displays based on the geographic location and the number of actual views that display receives.

To initially quantify the visibility of displays, we use the results based on *Random Building Navigators* to give us an overall sense of the connectedness of campus; across the network the mean number of views for a display is 73,299 (s.d. 81,067).

Top ranked displays include three screens outside a centrally located lecture theatre (Faraday Centre 270,625, Faraday Left 259,685, and Faraday Right 250,820)—this building has many entrances and exits and is often used as a cut through to many closely located buildings, hence these can indeed be considered well-connected locations. The 4th and 5th most highly-viewed displays in this simulation are located in a very central outdoor plaza area (Alex Square West 201,744) and on the exterior of a building towards the North of campus (Nuffield External 152,750), most likely it is the external nature of these screens that makes them more prominent.


By contrast, the five poorest ranked displays include two residential buildings at the most Southerly part of campus (Cartmel 1 College 463, Cartmel 2 College 568), and three embedded deep inside office buildings with a single entrance/exit (i.e. those that cannot be used as a pass through to other spaces; Human Resources 1,763, Student Services 7,528 and ISS 11,521).

Whilst this measure provides an indication regarding the prominence of displays based simply on the connectedness of the graph, by adding more accurate user models we can gain further data that more accurately describes the prominence of particular displays. For example, only considering *On Campus Students* and *Off Campus Students*, we found a total of 5,536,291 glances at the 36 displays across the deployment² (mean 87,878 views per day, s.d. 26,068) with each display attracting an average of 178,590 views overall (s.d. 278,327) equating to 3,700 per day (s.d. 5,973).


We note that the results of these simulations are largely comparable to those from our *Random Building Navigator* reflecting the popularity of many central locations.

Being able to recognise viewers in our simulation across multiple displays allows us further to report the proportion of unique viewers to the total number of views of a display to identify displays with high viewer coverage (fig. 2). We observe that displays with a lower ratio (displays that are revisited often) are dominated by displays located outside (Nuffield and Alex Square) or at key locations on campus


²Over the course of our deployment, the number of actual displays varies from 29 to 41 as a result of new screens being added and others undergoing maintenance. This figure is therefore the mean number of displays in the deployment over time (36.27).



(a) *Off Campus Student*



(b) *On Campus Student*



(c) *Random Building Navigator*

Figure 1: Cumulative Distribution Function of content view frequencies for all three signage viewer mobility models.

(LEC and Faraday) that are consequently passed while walking between parts of campus. Displays with a high ratio are located in colleges or rather niche locations (e.g. Graduate, Student Services and Learning Zone).

Together these different measures (connectedness, total visits, unique visits, duration of glances) have significant potential for understanding a physical deployment of displays that is useful to both display owners (e.g. to plan alterations to hardware such as adding a display or deploy a new interaction modality), and to content providers (e.g. for planning future content campaigns).

4.2 Network Visibility of Content

A common focus in signage analytics is the measurement of *content impressions*, similar to the concept in the web in which it describes the number of times a particular piece of content has been accessed or viewed (on the web these two

things are typically synonymous). However, unlike the web, content shown on a public display is not guaranteed to be viewed – there may be no people currently in the viewing area of the display or those people may simply be attending to other things.

Using the same simulations as above (*Random Building Navigator*, *On Campus Student*, *Off Campus Student*) we can report the average number of viewers for each content item. In our *Random Building Navigator* each content item yields an average of 217 views (116 unique views) per content item per day (s.d. 507 and 190 respectively). The most frequently seen content items include *Lancaster Arts* and *eCampus Sticky*. The former is scheduled exclusively on Nuffield External (previously reported as one of the most-viewed displays) while the latter is scheduled globally across all displays.

The least frequently seen items include content that is only scheduled on single displays that are located on less-accessible locations (e.g. ISS).

For content scheduled multiple times onto a single display, or indeed across multiple displays, there arises the potential for a single viewer to see a content item more than once (as reflected by the differences in the total and unique viewer numbers reported above). By utilising the simulation’s capacity to track unique simulated users across the deployment, we can easily start to identify the frequency with which a content item is seen repeatedly by the same user (this is useful both for content providers who wish to maximise exposure to a particular message, and to those who aspire to ensure users see different content each time).

Figure 1 looks to answer the question “*how many views does a single content item receive from the same viewer?*” by plotting the Cumulative Distribution Functions of content view frequencies for all three viewer groups. We note that the *Off Campus Student* group is exposed to fewer content items with lower frequencies per item while having a single very popular item (almost 90 views). This is probably a result of the students’ specified starting point and the exposure to specific displays (e.g. Alex Square East, Alex Square West) that are situated near the main bus stop and configured to show just one content item (e.g. the bus timetable).

Finally, we can approximate the duration that viewers have spent looking at displays in our deployment by combining the above figures with other display studies on display attention. A recent eye-tracking study by Dalton et al. indicated that passers-by attended to displays for a mean of 0.318 seconds (s.d. 0.261) [4], whilst an earlier paper from Huang et al. describing observation studies suggested a longer period attending to the display (1-2 seconds) [10]. Using the Dalton figure as a minimum and the midpoint from Huang’s range (i.e. 1.5) as a maximum, we find that the 217 mean views our content attracts in a day equates to between 1.15 minutes, and 5.43 minutes (average total glance duration per content item 20.33 minutes – 95.88 minutes).

4.3 Visibility of Content to Viewers

Considering first the issue of display impressions, we can use combined analytics and mobility data to identify the total number of content items seen and the number of distinct content items seen for each of the viewer models. For example, instead of asking the question “*How many unique users see this display?*”, we can instead provide approximations to

answer questions such as “*How many unique content items does an average user see?*”.

Overall we find clear similarities in the way *On Campus Students* and *Random Building Navigators* experience content across the network (approx 10 views of 7 unique content items per day). By contrast, *Off Campus Students* typically view a lower number of unique content items but see them much more frequently (approx 14 views of 4 unique content items per day). Large standard deviations across all the *On Campus Student* and *Random Building Navigator* populations reflect the high variation in content schedules across different display locations [Table 1], whilst this is again smaller for *Off Campus Students*. The smaller selection of content seen by *Off Campus Students* is a result of a more constrained mobility model that has students arriving and leaving from a specific point on campus (the main bus terminal).

A second question we can attempt to answer with this combination of mobility and analytics data relates to content repetition – i.e. how many times does a single user see repeated content? For a viewer in our *On Campus Student*, *Off Campus Student* and *Random Building Navigator* groups, the average number of times a content item is seen is 86.77, 54.63 and 94.86 respectively (s.d. 22.82, 12.47, 26.02) across the whole simulation. These numbers are likely to be a result of infrequent content changes.

Finally, if we consider the glance duration statistics from earlier, we can estimate the amount of time one viewer spends looking at content in our deployment throughout the entire simulation. For a viewer in our *Off Campus Student*, *On Campus Student* and *Random Building Navigator* groups, the total duration spent looking at content in this two month period is 4.52-21.33, 5.26-24.82 and 5.26-24.83 minutes respectively. This amounts to just 5-25 seconds per day, i.e. less than 0.714 seconds per content item for *On Campus Student* or *Random Building Navigator* and 1.25 seconds per item for an *Off Campus Student*.

4.4 Limitations

In this section we have endeavoured to provide examples of the new types of insights that can be delivered through viewer-centric physical analytics for digital signage. These insights cannot be generated without traces of individual viewer mobility. Such traces can be provided through tracking individual users but not without a substantial risk to personal privacy. As we have discussed in this paper, we solve this problem, and obviate the need for new tracking hardware, through use of synthetic traces of human mobility.

To demonstrate the principle, we have chosen simplistic mobility models. We recognise that this choice has a direct bearing on the utility of the exact numbers produced by our simulator. However, our intent in this paper is not to provide detailed results of exactly how many screens or content items an individual viewer sees, but rather to illustrate the types of results and reports that can be generated using a combination of physical analytics and simulated viewer mobility. Clearly additional factors can be introduced to improve the simulation model. For example, student timetables, bus arrival times and room occupancy counts could all be used to improve the overall simulation. Similarly, traces representing real student mobility (e.g. wifi associations) could also be used to improve the models. This is a subject of future work.




Figure 2: Plot showing the mean ratio of unique visitors to number of total visits for each class of student (ordered by ratio for the *Off Campus Student*). Low ratio values (i.e. those on the left side of the graph) represent displays with more recurring visitors, whilst high values represent displays with a higher proportion of unique visitors.

	unique views			total views		
	mean	median	std	mean	median	std
<i>Off Campus Student</i>	4.49	4.00	2.70	13.76	14.00	8.16
<i>On Campus Student</i>	8.77	7.00	5.91	16.01	10.00	20.06
<i>Random Navigator</i>	8.14	7.00	5.36	16.02	11.00	17.75

Table 1: Aggregated count of unique and total content views per day per viewer for each mobility model.

5. RELATED WORK

The described work brings together two distinct domains: signage analytics and simulations of human mobility.

5.1 Sign Analytics

The emergence of video analytics systems in the last decade has enabled researchers to gather new analytics around event detection, surveillance and crowd control [12]. Williamson et al. developed a system that measures the impact of a public display in a space by tracking pedestrian traffic and gathering information about direction, speed and location of pedestrians [20]. To take sign analytics one step further, Tian et al. developed a framework that combines video analytics with points-of-sale information to provide targeted advertising on digital signs using data mining [19]. Gillian et al. developed a system that combines viewer location data with on-sign sensors to support display applications that “follow” the viewer across multiple displays [9]. The system is able to provide analytics on the current number of viewers and their location in front of a display. With the aim of enhancing advertising campaigns, Farniella et al. proposed a framework that used video analytics and face recognition to identify revisiting viewers in front of a display allowing to run specific advertising campaigns [5]. Commercial products such as Intel Anonymous Viewer Analytics [11] and Fraunhofer AVARD [7] are designed to collect audience impressions, dwell times, and viewer demographics (e.g. age).

The systems described above depend on signs being equipped with video analytics capabilities and, in general, aim to report analytics relating to a single sign. In contrast, our

work requires no additional hardware to be associated with a sign and focuses not on individual signs and their surroundings but rather on understanding viewers’ experiences as they navigate environments with multiple signs.

5.2 Mobility Simulation

Human mobility modelling has become a popular area of research due to its importance in physical and societal processes and the emergence of new location-based mobile applications. In this paper we design human mobility models on a graph aiming to capture aspects of crowd mobility in a university campus. Simulations of human mobility correspond to a branch of human mobility studies. Their application spans a number of contexts and application scenarios, including simulations for routing in mobile ad-hoc networks [1, 6], modelling of crowds in disaster and evacuation scenarios [16, 21], or human travelling patterns in the urban domain which include vehicle traffic modeling [2], home to work and resource flow simulations [17] or agent-based models which aspire to capture abstract human mobility patterns in urban or wider geographical areas [18, 14].

In an initial approach Ostkamp and Kray explored the use of simulation techniques for evaluating potential locations for new display deployments in an environment using “augmented panoramic imagery” [15]. In contrast, our approach focusses on showing the potential of viewer-centric analytics that can be obtained by simulating people’s mobility throughout the network and mapping these onto historical physical analytics datasets.

6. CONCLUSION

In order to achieve a significant change in the perceived value of digital signs we need to increase our understanding of the relationship between viewers and the content that is presented. Current sign analytics fail to provide this understanding because *they cannot capture the viewer experience across multiple signs*. Our early experiences suggest that by combining easy to capture sign analytics with synthetic traces of viewer mobility we can address this fundamental shortcoming and provide detailed reports that offer *viewer-centric analytics*, providing us with new tools to improve the end-user experience. In future work in this area we will explore more detailed mobility models and seek to validate our findings through observational studies.

7. ACKNOWLEDGMENTS

This research is partially funded through the Future and Emerging Technologies (FET) programme within the 7th Framework Programme for Research of the European Commission, under FET grant number: 612933 (RECALL) and was made possible with the support of a Google Faculty Research Award and Google Cloud Credits.

8. REFERENCES

- [1] ASCHENBRUCK, N., FRANK, M., MARTINI, P., AND TÖLLE, J. Human mobility in manet disaster area simulation—a realistic approach. In *29th Annual IEEE International Conference on Local Computer Networks* (2004), IEEE, pp. 668–675.
- [2] BEHRISCH, M., BIEKER, L., ERDMANN, J., AND KRAJZEWICZ, D. Sumo—simulation of urban mobility. In *Proceedings of SIMUL '11* (2011).
- [3] CLINCH, S., MIKUSZ, M., GREIS, M., DAVIES, N., AND FRIDAY, A. Mercury: An application store for open display networks. In *Proceedings of UbiComp '14* (2014), ACM, pp. 511–522.
- [4] DALTON, N. S., COLLINS, E., AND MARSHALL, P. Display blindness?: Looking again at the visibility of situated displays using eye-tracking. In *Proceedings of CHI '15* (2015), ACM, pp. 3889–3898.
- [5] FARINELLA, G. M., FARIOLI, G., BATTIATO, S., LEONARDI, S., AND GALLO, G. Face re-identification for digital signage applications. In *Video Analytics for Audience Measurement* (2014), Springer.
- [6] FEELEY, M., HUTCHINSON, N., AND RAY, S. Realistic mobility for mobile ad hoc network simulation. In *Ad-Hoc, Mobile, and Wireless Networks*. Springer, 2004, pp. 324–329.
- [7] FRAUNHOFER IIS. Fraunhofer AVARD. <http://www.iis.fraunhofer.de/en/ff/bsy/tech/bildanalyse/avard.html>.
- [8] FRIDAY, A., DAVIES, N., AND EFSTRATIOU, C. Reflections on long-term experiments with public displays. *Computer, IEEE* 45, 5 (May 2012), 34–41.
- [9] GILLIAN, N., PFENNINGER, S., RUSSEL, S., AND PARADISO, J. Gestures everywhere: A multimodal sensor fusion and analysis framework for pervasive displays. In *Proceedings of PerDis '14* (2014).
- [10] HUANG, E. M., KOSTER, A., AND BORCHERS, J. Overcoming assumptions and uncovering practices: When does the public really look at public displays? In *Proceedings of Pervasive '08* (2008), Springer.
- [11] INTEL CORPORATION. Intel AIM Suite. <https://aimsuite.intel.com/>.
- [12] LIU, H., CHEN, S., AND KUBOTA, N. Intelligent video systems and analytics: A survey. *IEEE Transactions on Industrial Informatics* 9, 3 (Aug 2013), 1222–1233.
- [13] MIKUSZ, M., CLINCH, S., JONES, R., HARDING, M., WINSTANLEY, C., AND DAVIES, N. Repurposing web analytics to support the iot. *Computer* 48, 9 (Sept 2015), 42–49.
- [14] NOULAS, A., SCCELLATO, S., LAMBIOTTE, R., PONTIL, M., AND MASCOLO, C. A tale of many cities: universal patterns in human urban mobility. *PloS one* 7, 5 (2012), e37027.
- [15] OSTKAMP, M., AND KRAY, C. Supporting design, prototyping, and evaluation of public display systems. In *Proceedings of the 2014 ACM SIGCHI Symposium on Engineering Interactive Computing Systems* (New York, NY, USA, 2014), EICS '14, ACM, pp. 263–272.
- [16] PAN, X., HAN, C. S., DAUBER, K., AND LAW, K. H. A multi-agent based framework for the simulation of human and social behaviors during emergency evacuations. *Ai & Society* 22, 2 (2007), 113–132.
- [17] ROBINSON, D., HALDI, F., KÄMPF, J., LEROUX, P., PEREZ, D., RASHEED, A., AND WILKE, U. Citysim: Comprehensive micro-simulation of resource flows for sustainable urban planning. In *Proceedings of IBPSA '09* (2009), Citeseer, pp. 1083–1090.
- [18] SIMINI, F., GONZÁLEZ, M. C., MARITAN, A., AND BARABÁSI, A.-L. A universal model for mobility and migration patterns. *Nature* 484, 7392 (2012), 96–100.
- [19] TIAN, P., SANJAY, A. V., CHIRANJEEVI, K., AND MALIK, S. M. Intelligent advertising framework for digital signage. In *Proceedings of KDD '12* (2012), ACM, pp. 1532–1535.
- [20] WILLIAMSON, J. R., AND WILLIAMSON, J. Analysing pedestrian traffic around public displays. In *Proceedings of PerDis '14* (2014), ACM, pp. 13:13–13:18.
- [21] ZHONG, M., SHI, C., TU, X., FU, T., AND HE, L. Study of the human evacuation simulation of metro fire safety analysis in china. *Journal of Loss Prevention in the Process Industries* 21, 3 (2008), 287–298.