Design Considerations for Multi-Stakeholder Display Analytics

Mateusz Mikusz[1], Sarah Clinch[2] and Nigel Davies[1]

[1]School of Computing & Communications
Lancaster University
Lancaster, UK

m.mikusz | n.a.davies @lancaster.ac.uk

[2]School of Computer Science
University of Manchester
Manchester, UK

sarah.clinch@manchester.ac.uk

ABSTRACT

Measuring viewer interactions through detailed analytics will be crucial to improving the overall performance of future open display networks. However, in contrast to traditional sign and web analytics systems, such display networks are likely to feature multiple stakeholders each with the ability to collect a subset of the required analytics information. Combining analytics data from multiple stakeholders could lead to new insights, but stakeholders may have limited willingness to share information due to privacy concerns or commercial sensitivities. In this paper, we provide a comprehensive overview of analytics data that might be captured by different stakeholders in a display network, make the case for the synthesis of analytics data in such display networks, present design considerations for future architectures designed to enable the sharing of display analytics information, and offer an example of how such systems might be implemented.

Author Keywords

public displays; open display networks; analytics

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

Traditionally public displays are operated as closed networks in which the content shown is controlled by a small number of stakeholders and is typically focused on advertising. Recent research has suggested that future public display networks will be open to a wide range of content from multiple stakeholders—helping to drive new levels of innovation in the sector [14]. In particular, such networks are likely to see a shift towards highly personalised content, and will enable viewers to interact with displays—for example through their mobile phones, gestures or on-sign sensors.

Analytics will be critical for improving the effectiveness of such open display networks enabling, for example, display

providers to gain deep insights into viewer interaction patterns and helping developers deliver high quality interactive experiences for users. As a point of comparison, for conventional web applications, systems such as Google Analytics already provide such comprehensive insights and can be easily integrated into new and existing websites. Such systems have helped drive significant improvements in the experience offered to web users and enable site owners to constantly fine-tune their offering to maximise their return on investment.

Current signage analytics systems provide records of content shown on public displays and audience information. In particular, video analytics can provide audience demographics [29] as well as tracking behaviour such as gaze [27, 31]. However, signage analytics remain extremely limited in comparison to traditional web analytics as they are typically unable to report on individual user behaviour (i.e. they can not track users over multiple engagements) or activities that might take place following interaction with a sign (e.g. making an off-line purchase of an advertised product).

The emergence of open display networks and personalised and interactive applications raise analytics challenges, but also offers new opportunities. In particular, the increase in the number of stakeholders in such networks is important since each stakeholder may have opportunities to collect distinct sets of analytics data. In this paper we explore the concept of *analytics synthesis*, i.e. combining analytics data from multiple stakeholders to provide new insights into display network effectiveness and viewer behaviour.

Our contributions are: (1) a comprehensive description of analytics data that might be captured by different stakeholders and the identification of synthesis opportunities; (2) a set of design considerations for future architectures that facilitate analytics data synthesis across multiple stakeholders; and, (3) an example of how support for analytics synthesis could be implemented in the context of an existing analytics system.

Analytics synthesis relies on stakeholders being willing to share subsets of their analytics data and we explicitly explore precedents that exist for data sharing and consider how support for different levels of sharing can be provided.

STAKEHOLDERS AND ANALYTICS DATA

Prior research has identified four stakeholders in typical public display deployments (figure 1): display owners, space owners, content providers and viewers [3, 13, 54]. In traditional display systems many of these stakeholders are repre-

Stakeholder	Data Category					
	Anonymous Counting	**Anonymous Tracking**	**Gesture Recognition**	**Behaviour Analysis**	**Pseudo-anonymous Tracking**	**Contextual Events**
Display Owner	Number of viewers facing a display [25, 29, 45]; interacting with a display [25, 29]; and viewer demographics (gender, age, etc.) [19, 29]	Orientation of viewers' heads [31, 56, 57]; viewer distance from a display [5, 31, 56, 57]; and viewer dwell time [19, 31, 56]	Identification of pre-configured gestures [25, 27, 31, 56]; gesture logging [45]		Re-identifying people across multiple screens [18]	Content change patterns [42]; display state and failure monitoring [42]; presentation of personalised content based on proximity [5, 56, 57]
Space Owner	Number of potential viewers in display area [44, 50, 55]; time potential viewers spend in area [50]; and entrance and exit counts [50]	Physical characteristics of viewers (hair type, eyewear, clothing colour) through surveillance cameras [26, 52]		People tracking within video [19, 44, 49]; re-identifying people [28, 44]; pedestrian models [4]; frequent visitors [23]; fraud detection [17]; item interaction [33, 49]	Direction, speed and location [58]; multi-camera tracking [59]; shopping cart tracking [47]; WiFi fingerprints [1]	Number of sales [50, 55]; conversion rate [50]; PoS events (refund, cancellation) [55]
Content Provider			Gestures for content selection [27]; touch input [30, 56]	Navigation patterns [8]; content interaction [2]		Displays showing content [42]; QR code scans [21]; user engagement [16, 38]
Viewer				Viewer walking patterns [10, 48]; step counts [10]; dwell times [48]; gaze tracking [48]; gesturing [48]	Data from intertial sensors [52]; indoor location with other viewers [37]	

Table 1. Overview of data that can be captured by stakeholders within open display networks, as grouped by category. Green cells indicate a rich set of data is held by a stakeholder in that category. Red cells indicates that stakeholders are typically unable to capture or hold information in the category.
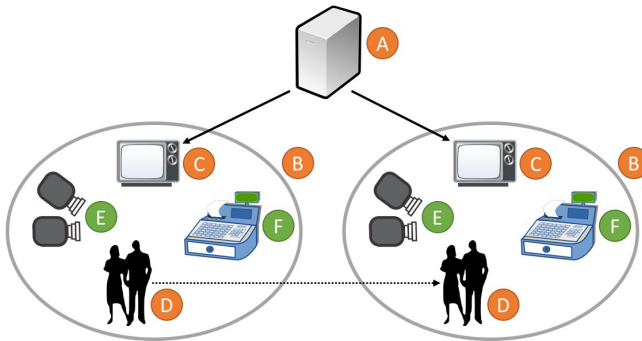

Figure 1. Overview of stakeholders and sample sensing devices in an open display network: (A) content provider; (B) space owner; (C) public display owner; (D) viewers who may move across multiple spaces; (E, F) CCTV cameras and point-of-sales terminals as examples of additional data-sources in the space.

sented by a single organisation (e.g. an advertising company may simultaneously act as both the display owner and content provider). However, display networks are increasingly featuring a more diverse set of stakeholders. For example, in shopping complexes many hundreds of displays may exist in close proximity yet be situated in spaces owned by different stakeholders (shops or concession owners), operated by different signage companies and show content from multiple distinct content providers. A review of the literature reveals the breadth of analytics data that each of the four stakeholder groups may capture:

Display owners. In general, display owners have access to data captured by the display presentation software and associated sensors mounted to the display. These sensors typically enable the display owner to capture detailed information for a narrow field of view in front of the display and can supply data such as anonymous viewer demographics [29] and interaction tracking [25].

Space owners. In contrast to display owners, space owners have a broader view of the surrounding environment. Fixed infrastructure mounted in the space enables behaviour analysis of potential viewers and other people in a predefined area [44]. Pseudo-anonymous tracking can be performed to retrieve comprehensive insights into movement patterns [58, 59]. Contextual events may also be captured by space owners, e.g. in retail environments space owners might have access to sales numbers and targets [50], and point-of-sales events [55].

Content providers. Content providers typically have detailed information about which displays are showing their content [42] and the interaction [27] and navigation patterns [8] within their applications but have limited knowledge about the context and environment of displays.

Viewers. Viewers typically hold a very rich set of information about their own activities including physical location traces, online activities and data from a wide range of mobile sensors including pictures, accelerometer data and in some cases even information on emotional states [39]. Moreover, viewers inherently are able to identify themselves. In addition, viewers may also use technology such as life-logging cameras to capture knowledge about other viewers in the environment and tasks being performed [48].

In table 1 we categorise the type of analytics data that different stakeholders may have access to—focusing particularly on viewer-related analytics (c.f. analytics relating to system performance). Anonymous counting refers to the ability of stakeholders to produce an estimate of the number of viewers (or potential viewers) for a given display, while anonymous tracking and gesture recognition provide more detailed insight into how viewers physically interact with the display. In the majority of cases the stakeholders concerned are unable to identify individual viewers and cannot link records of viewings over time. Behaviour analysis and pseudo-anonymous

tracking cover analytics that typically try and provide more detailed information on individual viewer patterns. Such systems usually track specific events (e.g. mobility patterns or fraud detection) and use video analytics or Wi-Fi fingerprinting techniques for re-identifying or tracking individuals. The contextual events category covers analytics data relating both to displays themselves (e.g. their power state) and information from other systems in the environment, e.g. point-of-sale information to help understand purchasing patterns.

THE CASE FOR ANALYTICS SYNTHESIS

Creating New Insights

While individual stakeholders can capture considerable quantities of analytics data none can approach the richness of insights that can be obtained by combining data from multiple stakeholders. Consider the following examples of how analytics from multiple sources can provide new insights:

Display Owners & Space Owners. Display owners typically capture detailed information from displays and their immediate surroundings but have no broader view of the space in which displays are located. In contrast, space owners may have access to a rich set of information about the space in which a display is located (e.g. how viewers navigate through the space and purchasing records from nearby retail outlets). Combining data from both sets of stakeholders enables us to answer questions such as *how does the activity of the display influence viewer movement patterns?* and *what percentage of potential viewers in a space actually interact with a given set of content?* By using additional data from the space owner such as purchasing activity *we can begin, for example, to directly correlate signage activity with purchasing patterns.*

Display Owners & Content Providers. In the context of open display networks, display owners often show applications and content that is provided by third parties. As a result, display owners may have limited information on the nature of content that is being shown (e.g. a display owner may know their display is showing a url for the BBC but not have access to the metadata behind the content or the details of user interactions with the page). In contrast, content providers have highly detailed knowledge about their own content but very little information regarding the display on which it is shown. By combining their datasets both stakeholders can accrue benefits. For example, *understanding which content has caused viewers to dwell for long periods in front of a display*, and *which demographic groups have seen content items?* Such insights could be used to *automatically develop schedules of content to match measured audience engagement* – improving both the use of the display and the experience for viewers.

Display Owner & Viewer. Viewers can provide extensive data on a given individual. In contrast display owners very rarely know the identity of specific viewers. By combining data from both stakeholders, new insights can be gained such as *how many returning viewers does a display have*, *how many displays do viewers' location traces suggest they pass-by throughout a day?* and *how does a display impact on the physical activity or emotional state of a viewer after they*

have left the immediate vicinity of the display. New technology such as wearable cameras offer additional possibilities by combining video analytics data from both the perspective of the display and the viewer enabling us to, for example, *build a picture of the objects viewers look at immediately before and after interacting with a display.*

Further combinations of the stakeholders described above can provide additional insights, e.g.:

Space Owner & Content Provider. Content providers typically have no way of determining the broader context in which their content is displayed. Combining data from content providers and space owners would help address this shortcoming and enable, for example, us to *understanding the relationship between content interaction and subsequent actions in the real-world.* In particular, combining content display records with sales statistics could lead to new data on the effectiveness of signage advertising campaigns.

Space Owner & Viewer. While space owners are able to capture detailed viewer movement traces, the ability of space owners to recognise individual viewers is limited. Viewers hold a rich set about their activities including those that happened beyond the boundaries of the space that a space owner is able to track. Synthesising data from both stakeholders would allow us to answer questions such as *where did viewers come from and go to immediately before and after they entered the space.*

Viewer & Content Provider. Content providers are typically not able to recognise and identify viewers that are interacting with their content. However, by combining analytics data from viewers and content providers we will able to answer questions such as *how many unique viewers did a content item have across multiple displays*, and allow content providers to *show content targeted to the predicted makeup of an audience (based on viewer mobility patterns).*

All of the examples above focus on combinations of two stakeholders. Bringing together 3 or 4 stakeholders provides yet further opportunities for new insights. For example, analytics from space owners, display owners and content producers would enable detailed data on *how interactions with display content impact on behaviour in a space, segmented by viewer demographic.*

Additional Benefits

In the discussion above we have focused on insights that can be obtained from combining analytics data from multiple stakeholders. Additional benefits that might be accrued include cost savings through the sharing of data collection overheads between stakeholders and an increased potential for user control of analytics data since viewers become an important stakeholder – supplying information directly rather than simply being the subject of observations.

More generally we note that research in other domains has often shown that fusing information from multiple datasources leads to more valuable results (e.g. [46, 53]). The synthesis of analytics data opens up opportunities for stakeholders to cap-

ture the same kinds of data from multiple sources, potentially increasing the veracity of the information.

DESIGN CONSIDERATIONS

Creating systems for analytics synthesis will require a number of significant design considerations to be addressed.

Protecting Commercially Sensitive Data

While there are clear benefits to be derived from analytics synthesis each stakeholder data set is likely to contain commercially sensitive information. For example, while a content provider might know which displays their content is shown on, they should not be able to access records of other content shown on the same display as these could reveal potentially sensitive information from competitors such as advertisement statistics or strategies. Equally, a display owner may have the overall picture of the content shown on their own display but will not be able to see where else in the world the content they have selected has been shown as this may provide unintended insights into competitor's display networks. Indeed, it is likely that *all analytics data is commercially sensitive* to at least one of the stakeholders because it provides insights into either a physical space, a display network, digital content being shown, or viewer interactions.

The presence of commercially sensitive data might suggest insurmountable obstacles to sharing. However, precedents for sharing analytics data exist in many other areas of business. For example, in business-to-business relations (e.g. along a supply chain) it is common to share information with business partners – providing insights and benefits to all those involved [35, 36]. Inventory levels and point-of-sales statistics can be automatically transmitted to suppliers with the goal of better demand prediction for certain products and performance improvements [34]. Such information sharing typically requires written agreements and contracts between all involved parties [22] and as a most basic requirement, benefits from data sharing must be clear for all partners [34]. Finally, we observe that while it might be expected that less commercially sensitive data would be more likely to be shared—for example, visitor statistics for shopping malls – we found that, even for this type of data, stakeholders typically only publish aggregated figures (e.g. [20]) suggesting that the need for contracts and agreements is inevitable.

In summary, we believe that there are clear benefits to be obtained by combining analytics data from multiple stakeholders and clear precedents exist for data sharing in commercial environments but *providing appropriate appropriate controls will be critical to success*.

Architectural Models

Initially it may appear that a data sharing architecture in which each stakeholder can control access to their own analytics data would be appropriate to provide support for analytics synthesis. A simple example with our stakeholders illustrates the drawbacks in this approach. Consider the case in which a display owner has access to the schedule of content shown on a screen and a space owner has access to a count of the number of potential viewers in the space. The aim is to produce a report that shows the number of potential viewers for a specific content item without risking disclosure of the full data sets of either stakeholder. It is not possible for either of the stakeholders to independently decide which data they need to release—only by merging both data sets can such a report be produced. As a result, *any future analytics platform will most likely need to include a logically common aggregation component across multiple stakeholders*.

We note that a logically common component does not necessarily imply a centralised implementation—a distributed implementation could be developed if required. In addition, we note that it is not necessary for there to be one single analytics service—multiple such services could exist but a single service would need to be employed by any group of stakeholders that wished to collaborate. Deploying an analytics systems as a logically centralised service with common components potentially introduces new trust concerns, particular as analytics data may be commercially sensitive. However, there is increasing confidence in the use of cloud services backed by appropriate service level agreements and we do not consider that this will be a significant issue. For example, companies already trust providers such as Google to maintain clear boundaries between data from different customers.

Data Ingestion

An analytics systems must make the ingestion of raw data as easy as possible as each stakeholder will typically maintain a large number of data sources. To simplify processing and storage of such information, and to allow easy integration of future data sources, a standard reporting format such as the Universal Measurement Protocol (UMP) [24] could be used. Such formats typically provide a wealth of appropriate data models (e.g. page impression, location information), and can be extended to support domain-specific data types such as viewer proximity or personal content preferences. *However no existing analytics reporting protocols are able to cope with the diversity of data types presented in table 1 and suitable extensions will be required.*

Aggregation Rules and Policies

Future analytics aggregation platforms will need to provide functions that allow users to control which data is aggregated, how the data is aggregated and with whom the results are shared. The first of these functions, i.e. controlling which data is aggregated, is likely to be amenable to standard access control and filtering approaches in which stakeholders can specify how data should be shared and any redaction (e.g. reduction in the sampling rate) that is needed prior to release. Indeed, it is crucial in sensor-rich environments that "users should be able to control the release of their own data" [15]. Providing the capability for data aggregation itself is more complex. *Our initial work suggests that aggregation is likely to be extremely specific to domain and data-type, and that support is likely to be needed for custom code fragments that perform the aggregation (rather than, for example, through generic queries).*

Control of the aggregation performed will also be necessary and may cover a range of options such as specifying the ana-

lytics data feeds involved, the level of aggregation, combinations of policies across datatypes and stakeholders, additional sharing limitations and permissions. For example, stakeholders could specify that their anonymous people counts should be correlated with sales statistics and display impressions to understand the effectiveness of display content but that no personally identifiable data should be released. Finally, stakeholders will wish to control the sharing and distribution of processed datasets, e.g. by specifying rules that describe the future use of their data and combined datasets, and stating access and permission rights.

Viewer Privacy

In traditional sign analytics no personalised information about the audience is captured [12] and hence such systems are typically compliant with legislation governing personal data collection (e.g. EU data protection rules) and raise few concerns amongst viewers. However, with the increased focus on analytics relating to individual viewer behaviour and delivering personalised content to public displays, new privacy related concerns emerge that are compounded when analytics data from multiple stakeholders is aggregated. e.g.:

Location tracking. Face recognition can be combined with video analytics to track an individual's movement between pervasive displays, potentially leading to comprehensive location traces in areas with a dense deployment of displays.

Purchasing patterns. The inclusion of additional contextual analytics data such as point of sales information, when combined with viewer tracking from space stakeholders, potentially allows detailed tracking of viewer purchasing patterns.

Susceptibility to advertising. Through analytics synthesis it may be possible to profile an individual viewer's susceptibility to display-based advertising. Such profiling of individuals is likely to be particularly contentious for viewers.

As a result, we believe that *a substantial challenge for future analytics platforms will be the need to address issues of viewer privacy and that fundamental work is required to understand the level of guarantees that will be possible in future environments that support synthesis of analytics data.*

IMPLEMENTATION CONSIDERATIONS

In order to explore how such a platform for synthesising analytics might be implemented in practice we have considered the modifications that would be necessary to Pheme [42], an existing platform for signage analytics. By extending Pheme we avoid having to create a new analytics system and are able to assess the additional complexity support for multi-stakeholder synthesis might entail.

Overview of Pheme

Pheme is a cloud-based service that provides a flexible way of storing and processing analytics data from a range of display networks and IoT devices [42]. In particular, Pheme provides features to map streaming analytics data to one or more alternative analytics services enabling, for example, display analytics data to be used in the context of a web analytics service. The architecture consists of four components (figure 2): an
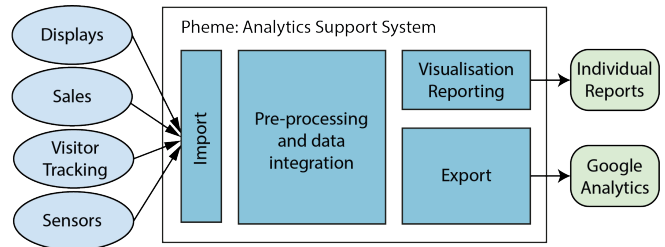

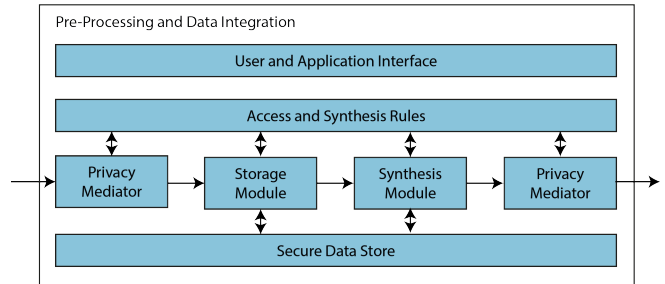**Figure 2. Pheme systems architecture (originally reported in [42]).**

Figure 3. Extended Pheme *pre-processing and data integration* module to support data synthesis from different stakeholders.

import module, a *pre-processing and data integration* module, a *visualisation and reporting* engine, and an *export* module. Client devices and sensors report their data to the *import* module using an extended version of the UMP that provides a richer set of data types for describing interaction with IoT components. Each incoming dataset triggers *pre-processing and data integration* modules that parse the incoming data and create objects that can be read by other components of the system. These modules can, for example, carry out data validation or filtering. The modular design of this component allows Pheme to be extended with additional pre-processors by plugging in new modules.

Following pre-processing and integration, data may then be automatically *exported* and injected into third-party analytics engines (e.g. Google Analytics) or stored locally for future processing. The *export* component allows developers to plug in multiple "injectors" – each providing a specific mapping to make the data compatible with a third-party service.

Pheme is implemented in Python on the Google AppEngine cloud service and has been in daily operation for around 30 months and has processed over 80 million analytics events.

Extensions to Support Synthesis of Analytics Streams

To support the synthesis of analytics streams that are owned by a number of different stakeholders we designed extensions to the original Pheme systems architecture located within the *pre-processing and data integration* module (figure 3).

Incoming analytics data would initially flow through an *inbound privacy mediator module*. This module builds on ideas first proposed in [15] and allows stakeholders with analytics data to deploy (or have deployed for them) software components that can redact data prior to it being used in the synthesis process. In many cases we would expect stakeholders to redact the data before sending the data stream to our proposed

service but the inclusion of a privacy mediator component enables them to also delegate this role to Pheme.

Depending on the rules specified, incoming data would then be written to secure storage via the *storage module*. This module is necessary because analytics data may need to be stored for a considerable period of time to support the full range of synthesis, anonymisation and redaction policies while Pheme is currently optimised for processing streaming data and has limited support for data storage. The *synthesis module* is responsible for subsequently combining the analytics data sets according to the rules specified by the relevant stakeholders. Once the synthesised data stream has been created, the synthesis module would submit the new combined dataset recursively back to the storage module for potential further synthesis.

Finally, the resulting data stream passes through an *out-bound privacy mediator module* that provides a further opportunity for stakeholders to mediate data prior to release. Data streams can then follow the original Pheme architecture flow: being passed in to the export module in which associated Injectors are run for reporting and visualisation either through internal or third-party services. These new data processing modules are supported by a secure storage facility and components for providing appropriate interfaces to enable applications and users to communicate with Pheme for management purposes.

RELATED WORK
Algorithms and techniques for extracting user behaviour patterns with the goal of improving user experience have been extensively studied in the context of the web (e.g. [6, 9, 51, 32]. With the increasing use of social media platforms, research is now focused on mining information published in these networks such as identifying current events and trends [41, 60]. In the digital signage domain video analytics have been used to provide detailed insights in user behaviour [11, 40] and many commercial systems are now available, e.g. [19, 28, 29, 49]. In retail environments, video analytics techniques are used to provide people counts and sales statistics [50], track customer locations based on the cart [47], and determine which places are most frequently visited and how customers navigate through the store [23]. In the context of public display research Williamson et al. have deployed a video analytics system to analyse changes in movement patterns of pedestrians around a display [58].

While there has been extensive research and commercial development of individual analytics technologies there has been significantly less research into combining multiple analytics datasources. Tian et al. have developed a framework that combines video analytics retrieved through cameras mounted on displays with data mining techniques to provide targeted advertising [53]. Venetianer et al. combined CCTV streams with point-of-sales events (such as sales, returns and cancellations) for cross-verification and fraud detection [55]. Beach at al. have developed a prototype system that combines social data from Facebook with sensor data captured on the smartphone, allowing provision of context-sensitive videos on a nearby public display [7]. A similar approach for combining data from social network and smartphone sensors was taken by Phan et al. who developed a framework for creating context-sensitive applications [46]. Combining simulated traces of user mobility with real-world signage data has been proposed by the authors in [43] to generate synthetic (artificial) analytics traces.

The systems and frameworks described mostly assume that different stakeholders roles are adopted by a single individual or organisation (for example, in the case where a shop owner acts as space owner, display owner and content provider). In contrast, this paper explicitly addresses the fact that each of the stakeholders are likely to be represented by different organisations, increasing the level of complexity in designing and implementing such a system. Understanding display stakeholder eco-systems has been briefly explored in [3, 13, 54] but we are not aware of any prior work that has focused on the relationship between stakeholders and display analytics.

CONCLUSIONS
Comprehensive analytics are likely to play a key role in enhancing the effectiveness of future display networks. While traditional display and signage analytics systems operate in stakeholder-defined silos, the move towards open networks of displays is likely to lead to new opportunities to synthesise analytics data for deeper insights into viewer behaviour. However, aggregating analytics data must balance the tensions that exist between a desire for more detailed insights and the need to maintain control of commercially sensitive data and to protect viewer privacy.

In this paper we have described the characteristics of this emerging analytics eco-system and presented a number of important design considerations for the developers of systems to support analytics synthesis.

While we believe that the realisation of a system to support multi-stakeholder analytics synthesis is viable, we note that conducting a full-scale trial of such a system would be extremely challenging. Specifically, such a trial would entail (i) the identification of a large multi-stakeholder display network, (ii) negotiation with a wide range of competing stakeholders to ensure that the trial is able to capture viewpoints from multiple stakeholders, (iii) establishment of levels of trust (and associated contracts) that would enable stakeholders to pass commercially sensitive data to a research organisation for hosting, and (iv) a longitudinal study that provided stakeholders with the opportunity to understand the business benefits that they might accrue from such synthesis.

As a result, having shown the conceptual feasibility of systems to support analytics synthesis our plan for future work is to produce a small number of highly detailed exemplar reports from a real-world display eco-system to further assess the benefits for stakeholders before attempting to conduct a significant multi-stakeholder, multi-year trial of the sort that would be required to demonstrate experimental validity.

Acknowledgements

REFERENCES

1. AirTight Networks, Inc. 2013. *AirTight Social Wi-Fi and Analytics for the Retail Store of the Future*. Whitepaper. AirTight Networks, Inc.

2. Florian Alt, Thomas Kubitza, Dominik Bial, Firas Zaidan, Markus Ortel, Björn Zurmaar, Tim Lewen, Alireza Sahami Shirazi, and Albrecht Schmidt. 2011. Digifieds: Insights into Deploying Digital Public Notice Areas in the Wild. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia (MUM '11)*. ACM, New York, NY, USA, 165–174. DOI:
`http://dx.doi.org/10.1145/2107596.2107618`

3. Florian Alt, Albrecht Schmidt, and Jörg Müller. 2012. Advertising on Public Display Networks. *Computer* 45, 5 (May 2012), 50–56. DOI:
`http://dx.doi.org/10.1109/MC.2012.150`

4. Gianluca Antonini, SantiagoVenegas Martinez, Michel Bierlaire, and JeanPhilippe Thiran. 2006. Behavioral Priors for Detection and Tracking of Pedestrians in Video Sequences. *International Journal of Computer Vision* 69, 2 (2006), 159–180. DOI:
`http://dx.doi.org/10.1007/s11263-005-4797-0`

5. Till Ballendat, Nicolai Marquardt, and Saul Greenberg. 2010. Proxemic Interaction: Designing for a Proximity and Orientation-aware Environment. In *ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*. ACM, New York, NY, USA, 121–130. DOI:
`http://dx.doi.org/10.1145/1936652.1936676`

6. Matthias Baumgarten, AlexG. Büchner, SarabjotS. Anand, MauriceD. Mulvenna, and JohnG. Hughes. 2000. User-Driven Navigation Pattern Discovery from Internet Data. In *Web Usage Analysis and User Profiling*, Brij Masand and Myra Spiliopoulou (Eds.). Lecture Notes in Computer Science, Vol. 1836. Springer Berlin Heidelberg, 74–91. DOI:
`http://dx.doi.org/10.1007/3-540-44934-5_5`

7. Aaron Beach, Mike Gartrell, Xinyu Xing, Richard Han, Qin Lv, Shivakant Mishra, and Karim Seada. 2010. Fusing Mobile, Sensor, and Social Data to Fully Enable Context-aware Computing. In *Proceedings of the Eleventh Workshop on Mobile Computing Systems & Applications (HotMobile '10)*. ACM, New York, NY, USA, 60–65. DOI:
`http://dx.doi.org/10.1145/1734583.1734599`

8. Bettina Berendt and Myra Spiliopoulou. 2000. Analysis of Navigation Behaviour in Web Sites Integrating Multiple Information Systems. *The VLDB Journal* 9, 1 (March 2000), 56–75. DOI:
`http://dx.doi.org/10.1007/s007780050083`

9. José Borges and Mark Levene. 2000. Data Mining of User Navigation Patterns. In *Web Usage Analysis and User Profiling*, Brij Masand and Myra Spiliopoulou (Eds.). Lecture Notes in Computer Science, Vol. 1836.
Springer Berlin Heidelberg, 92–112. DOI:
`http://dx.doi.org/10.1007/3-540-44934-5_6`

10. Agata Brajdic and Robert Harle. 2013. Walk Detection and Step Counting on Unconstrained Smartphones. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, New York, NY, USA, 225–234. DOI:`http://dx.doi.org/10.1145/2493432.2493449`

11. Joshua Candamo, Matthew Shreve, Dmitry B. Goldgof, Deborah B. Sapper, and Rangachar Kasturi. 2010. Understanding Transit Scenes: A Survey on Human Behavior-Recognition Algorithms. *Intelligent Transportation Systems, IEEE Transactions on* 11, 1 (March 2010), 206–224. DOI:
`http://dx.doi.org/10.1109/TITS.2009.2030963`

12. Ann Cavoukian. 2011. *White Paper: Anonymous Video Analytics (AVA) technology and privacy*. Technical Report. Information and Privacy Commissioner of Ontario, Canada.

13. Sarah Clinch, Mateusz Mikusz, Miriam Greis, Nigel Davies, and Adrian Friday. 2014. Mercury: An Application Store for Open Display Networks. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. ACM, New York, NY, USA, 511–522. DOI:`http://dx.doi.org/10.1145/2632048.2636080`

14. Nigel Davies, Marc Langheinrich, Rui Jose, and Albrecht Schmidt. 2012. Open Display Networks: A Communications Medium for the 21st Century. *Computer* 45, 5 (May 2012), 58–64. DOI:
`http://dx.doi.org/10.1109/MC.2012.114`

15. Nigel Davies, Nina Taft, Mahadev Satyanarayanan, Sarah Clinch, and Brandon Amos. 2016. Privacy Mediators: Helping IoT Cross the Chasm. In *Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications (HotMobile '16)*. ACM, New York, NY, USA, 39–44. DOI:`http://dx.doi.org/10.1145/2873587.2873600`

16. Florin Dobrian, Vyas Sekar, Asad Awan, Ion Stoica, Dilip Joseph, Aditya Ganjam, Jibin Zhan, and Hui Zhang. 2011. Understanding the Impact of Video Quality on User Engagement. In *Proceedings of the ACM SIGCOMM 2011 Conference (SIGCOMM '11)*. ACM, New York, NY, USA, 362–373. DOI:
`http://dx.doi.org/10.1145/2018436.2018478`

17. Quanfu Fan, A. Yanagawa, R. Bobbitt, Yun Zhai, R. Kjeldsen, S. Pankanti, and A. Hampapur. 2009. Fast detection of retail fraud using polar touch buttons. In *Multimedia and Expo, 2009. ICME 2009. IEEE International Conference on*. 1266–1269. DOI:
`http://dx.doi.org/10.1109/ICME.2009.5202732`

18. GiovanniMaria Farinella, Giuseppe Farioli, Sebastiano Battiato, Salvo Leonardi, and Giovanni Gallo. 2014. Face Re-Identification for Digital Signage Applications. In *Video Analytics for Audience Measurement*, Cosimo

Distante, Sebastiano Battiato, and Andrea Cavallaro (Eds.). Lecture Notes in Computer Science, Vol. 8811. Springer International Publishing, 40–52. DOI: http://dx.doi.org/10.1007/978-3-319-12811-5_3

19. Fraunhofer IIS. 2015. *AVARD - Anonymous Video Analytics for Retail and Digital Signage*. Technical Report. Fraunhofer IIS.

20. Retail Gazette. 2012. Westfield Stratford attracts 47m shoppers in first year. (Sept. 2012). https://www.retailgazette.co.uk/blog/2012/09/20444-westfield-stratford-attracts-47m-shoppers-in-first-year

21. Matthias Geel, Daniel Huguenin, and Moira C. Norrie. 2013. PresiShare: Opportunistic Sharing and Presentation of Content Using Public Displays and QR Codes. In *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays (PerDis '13)*. ACM, New York, NY, USA, 103–108. DOI: http://dx.doi.org/10.1145/2491568.2491591

22. Anupam Ghosh and Jane Fedorowicz. 2008. The role of trust in supply chain governance. *Business Process Management Journal* 14, 4 (2008), 453–470. DOI: http://dx.doi.org/10.1108/14637150810888019

23. Andreas Girgensohn, Frank Shipman, and Lynn Wilcox. 2008. Determining Activity Patterns in Retail Spaces Through Video Analysis. In *Proceedings of the 16th ACM International Conference on Multimedia (MM '08)*. ACM, New York, NY, USA, 889–892. DOI: http://dx.doi.org/10.1145/1459359.1459514

24. Google Inc. 2015. Measurement Protocol Parameter Reference. https://developers.google.com/analytics/devguides/collection/protocol/v1/parameters. (September 2015).

25. Kazjon Grace, Rainer Wasinger, Christopher Ackad, Anthony Collins, Oliver Dawson, Richard Gluga, Judy Kay, and Martin Tomitsch. 2013. Conveying Interactivity at an Interactive Public Information Display. In *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays (PerDis '13)*. ACM, New York, NY, USA, 19–24. DOI: http://dx.doi.org/10.1145/2491568.2491573

26. Arun Hampapur, Russell Bobbitt, Lisa Brown, Mike Desimone, Rogerio Feris, Rick Kjeldsen, Max Lu, Carl Mercier, Chris Milite, Stephen Russo, Chiao-Fe Shu, and Yun Zhai. 2009. Video Analytics in Urban Environments. In *Advanced Video and Signal Based Surveillance, 2009. AVSS '09. Sixth IEEE International Conference on*. 128–133. DOI: http://dx.doi.org/10.1109/AVSS.2009.100

27. John Hardy, Enrico Rukzio, and Nigel Davies. 2011. Real World Responses to Interactive Gesture Based Public Displays. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia (MUM '11)*. ACM, New York, NY, USA, 33–39. DOI: http://dx.doi.org/10.1145/2107596.2107600

28. IBM Corporation. 2013. *IBM Intelligent Video Analytics*. Technical Report. IBM Corporation.

29. Intel Corporation. 2011. *A Report on a Field Trial of Anonymous Video Analytics (AVA) in Digital Signage*. Technical Report. Intel Corporation.

30. Giulio Jacucci, Ann Morrison, Gabriela T. Richard, Jari Kleimola, Peter Peltonen, Lorenza Parisi, and Toni Laitinen. 2010. Worlds of Information: Designing for Engagement at a Public Multi-touch Display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 2267–2276. DOI: http://dx.doi.org/10.1145/1753326.1753669

31. Marko Jurmu, Masaki Ogawa, Sebastian Boring, Jukka Riekki, and Hideyuki Tokuda. 2013. Waving to a Touch Interface: Descriptive Field Study of a Multipurpose Multimodal Public Display. In *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays (PerDis '13)*. ACM, New York, NY, USA, 7–12. DOI: http://dx.doi.org/10.1145/2491568.2491571

32. Raymond Kosala and Hendrik Blockeel. 2000. Web Mining Research: A Survey. *SIGKDD Explor. Newsl.* 2, 1 (June 2000), 1–15. DOI: http://dx.doi.org/10.1145/360402.360406

33. N. Krahnstoever, J. Rittscher, P. Tu, K. Chean, and T. Tomlinson. 2005. Activity Recognition using Visual Tracking and RFID. In *Application of Computer Vision, 2005. WACV/MOTIONS '05 Volume 1. Seventh IEEE Workshops on*, Vol. 1. 494–500. DOI: http://dx.doi.org/10.1109/ACVMOT.2005.17

34. Susan Cohen Kulp, Hau L. Lee, and Elie Ofek. 2004. Manufacturer Benefits from Information Integration with Retail Customers. *Management Science* 50, 4 (2004), 431–444. DOI: http://dx.doi.org/10.1287/mnsc.1030.0182

35. Hau L. Lee, V. Padmanabhan, and Seungjin Whang. 2004. Information Distortion in a Supply Chain: The Bullwhip Effect. *Management Science* 50, 12_supplement (2004), 1875–1886. DOI: http://dx.doi.org/10.1287/mnsc.1040.0266

36. Hau L. Lee and Seungjin Whang. 2000. Information sharing in a supply chain. *International Journal of Manufacturing Technology and Management* 1, 1 (2000), 79–93. DOI: http://dx.doi.org/10.1504/IJMTM.2000.001329

37. Sungnam Lee, Hyojeong Shin, and Hojung Cha. 2012. A Pedestrian Tracking System using Group Mobility Information. In *Military Communications Conference, 2012*. 1–6. DOI: http://dx.doi.org/10.1109/MILCOM.2012.6415710

38. Janette Lehmann, Mounia Lalmas, Elad Yom-Tov, and Georges Dupret. 2012. Models of User Engagement. In

Proceedings of the 20th International Conference on User Modeling, Adaptation, and Personalization (UMAP'12). Springer-Verlag, Berlin, Heidelberg, 164–175. DOI: http://dx.doi.org/10.1007/978-3-642-31454-4_14

39. Christine Lætitia Lisetti and Fatma Nasoz. 2004. Using Noninvasive Wearable Computers to Recognize Human Emotions from Physiological Signals. *EURASIP Journal on Advances in Signal Processing* 2004, 11 (2004), 929414. DOI: http://dx.doi.org/10.1155/S1110865704406192

40. Honghai Liu, Shengyong Chen, and N. Kubota. 2013. Intelligent Video Systems and Analytics: A Survey. *Industrial Informatics, IEEE Transactions on* 9, 3 (Aug 2013), 1222–1233. DOI: http://dx.doi.org/10.1109/TII.2013.2255616

41. Michael Mathioudakis and Nick Koudas. 2010. TwitterMonitor: Trend Detection over the Twitter Stream. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data (SIGMOD '10)*. ACM, New York, NY, USA, 1155–1158. DOI: http://dx.doi.org/10.1145/1807167.1807306

42. Mateusz Mikusz, Sarah Clinch, Rachel Jones, Mike Harding, Christopher Winstanley, and Nigel Davies. 2015. Repurposing Web Analytics to Support the IoT. *Computer* 48, 9 (Sept 2015), 42–49. DOI: http://dx.doi.org/10.1109/MC.2015.260

43. Mateusz Mikusz, Anastasios Noulas, Nigel Davies, Sarah Clinch, and Adrian Friday. 2016. Next Generation Physical Analytics for Digital Signage. In *Proceedings of the 3rd International on Workshop on Physical Analytics (WPA '16)*. ACM, New York, NY, USA, 19–24. DOI: http://dx.doi.org/10.1145/2935651.2935658

44. Dennis Mitzel, Esther Horbert, Andreas Ess, and Bastian Leibe. 2010. Multi-person Tracking with Sparse Detection and Continuous Segmentation. In *Computer Vision – ECCV 2010*, Kostas Daniilidis, Petros Maragos, and Nikos Paragios (Eds.). Lecture Notes in Computer Science, Vol. 6311. Springer Berlin Heidelberg, 397–410. DOI: http://dx.doi.org/10.1007/978-3-642-15549-9_29

45. Gonzalo Parra, Joris Klerkx, and Erik Duval. 2014. Understanding Engagement with Interactive Public Displays: An Awareness Campaign in the Wild. In *Proceedings of The International Symposium on Pervasive Displays (PerDis '14)*. ACM, New York, NY, USA, Article 180, 6 pages. DOI: http://dx.doi.org/10.1145/2611009.2611020

46. Thomas Phan, Swaroop Kalasapur, and Anugeetha Kunjithapatham. 2014. Sensor fusion of physical and social data using Web SocialSense on smartphone mobile browsers. In *Consumer Communications and Networking Conference (CCNC), 2014 IEEE 11th*.

98–104. DOI: http://dx.doi.org/10.1109/CCNC.2014.6866555

47. Harikrishna G.N. Rai, Kishore Jonna, and P. Radha Krishna. 2011. Video Analytics Solution for Tracking Customer Locations in Retail Shopping Malls. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '11)*. ACM, New York, NY, USA, 773–776. DOI: http://dx.doi.org/10.1145/2020408.2020533

48. Swati Rallapalli, Aishwarya Ganesan, Krishna Chintalapudi, Venkat N. Padmanabhan, and Lili Qiu. 2014. Enabling Physical Analytics in Retail Stores Using Smart Glasses. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking (MobiCom '14)*. ACM, New York, NY, USA, 115–126. DOI: http://dx.doi.org/10.1145/2639108.2639126

49. SCALA. 2017. Advanced Analytics - Measurable data for Optimized Messaging. http://scala.com/advanced-analytics/. (April 2017).

50. A.W. Senior, L. Brown, A. Hampapur, C.-F. Shu, Y. Zhai, R.S. Feris, Y.-L. Tian, S. Borger, and C. Carlson. 2007. Video analytics for retail. In *Advanced Video and Signal Based Surveillance, 2007. AVSS 2007. IEEE Conference on*. 423–428. DOI: http://dx.doi.org/10.1109/AVSS.2007.4425348

51. Jaideep Srivastava, Robert Cooley, Mukund Deshpande, and Pang-Ning Tan. 2000. Web Usage Mining: Discovery and Applications of Usage Patterns from Web Data. *SIGKDD Explor. Newsl.* 1, 2 (Jan. 2000), 12–23. DOI: http://dx.doi.org/10.1145/846183.846188

52. Thiago Teixeira, Deokwoo Jung, and Andreas Savvides. 2010. Tasking Networked CCTV Cameras and Mobile Phones to Identify and Localize Multiple People. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp '10)*. ACM, New York, NY, USA, 213–222. DOI: http://dx.doi.org/10.1145/1864349.1864367

53. Phil Tian, Addicam V. Sanjay, Kunapareddy Chiranjeevi, and Shahzad Malik Malik. 2012. Intelligent Advertising Framework for Digital Signage. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '12)*. ACM, New York, NY, USA, 1532–1535. DOI: http://dx.doi.org/10.1145/2339530.2339773

54. Ville Valkama and Timo Ojala. 2011. Stakeholder Value Propositions on Open Community Testbed of Interactive Public Displays. In *Proceedings of the International Conference on Management of Emergent Digital EcoSystems (MEDES '11)*. ACM, New York, NY, USA, 107–113. DOI: http://dx.doi.org/10.1145/2077489.2077509

55. P.L. Venetianer, Z. Zhang, A. Scanlon, Y. Hu, and A.J. Lipton. 2007. Video verification of point of sale

transactions. In *Advanced Video and Signal Based Surveillance, 2007. AVSS 2007. IEEE Conference on*. 411–416. DOI:
`http://dx.doi.org/10.1109/AVSS.2007.4425346`

56. Daniel Vogel and Ravin Balakrishnan. 2004. Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology (UIST '04)*. ACM, New York, NY, USA, 137–146. DOI:`http://dx.doi.org/10.1145/1029632.1029656`

57. Miaosen Wang, Sebastian Boring, and Saul Greenberg. 2012. Proxemic Peddler: A Public Advertising Display That Captures and Preserves the Attention of a Passerby. In *Proceedings of the 2012 International Symposium on Pervasive Displays (PerDis '12)*. ACM, New York, NY, USA, Article 3, 6 pages. DOI:
`http://dx.doi.org/10.1145/2307798.2307801`

58. Julie R. Williamson and John Williamson. 2014. Analysing Pedestrian Traffic Around Public Displays. In *Proceedings of The International Symposium on Pervasive Displays (PerDis '14)*. ACM, New York, NY, USA, Article 13, 6 pages. DOI:
`http://dx.doi.org/10.1145/2611009.2611022`

59. Tao Yang, F. Chen, D. Kimber, and J. Vaughan. 2007. Robust People Detection and Tracking in a Multi-Camera Indoor Visual Surveillance System. In *Multimedia and Expo, 2007 IEEE International Conference on*. 675–678. DOI:
`http://dx.doi.org/10.1109/ICME.2007.4284740`

60. Daniel Zeng, Hsinchun Chen, R. Lusch, and Shu-Hsing Li. 2010. Social Media Analytics and Intelligence. *Intelligent Systems, IEEE* 25, 6 (Nov 2010), 13–16. DOI:`http://dx.doi.org/10.1109/MIS.2010.151`