Whither in-store analytics: How in-store behavioural analytics has changed and where it is heading

Received (in revised form): 12th July, 2021

Gary Angel
Chief Executive Officer, Digital Mortar, USA

Gary Angel is the founder and Chief Executive Officer of Digital Mortar. Digital Mortar provides comprehensive collection and measurement of the retail customer journey. Following Ernst & Young's acquisition of his previous venture, Semphonic, Gary led EY's Digital Analytics Practice. He is the author of 'Measuring the Digital World' (FT Press, 2016).

Digital Mortar, 1005 A Street, #211, San Rafael, CA 94901, USA
E-mail: gary.angel@digitalmortar.com

Abstract As organisations evolve, they want a more comprehensive view of customer behaviour. While digital channels are easily measured and have been studied extensively, for retailers with a physical presence, the store is somewhat of a black box: retailers know how many customers they have going in and what sales are coming out the other side, but what happens in between has long been a mystery. As this paper shows, however, improvements in both technology and analytics are shining fresh light on in-store activity. This paper describes recent advances in technology for monitoring in-store journeys and how such advances have enabled people-movement data to be used in more sophisticated ways, despite privacy limitations.

KEYWORDS: shopper analytics, store analytics, full-journey measurement, shopper measurement, store optimisation, customer experience

INTRODUCTION

E-commerce stores routinely track every aspect of the customer shopping experience: which campaigns attracted visitors, which products got searched, what filters were used, the detail pages viewed, and how long people spent on each feature or page. Because of this, e-commerce stores know the entire shopper journey from initial interest to checkout.

For traditional in-store marketers, however, the world is a lot less clear.

Naturally, all retailers measure at the point-of-sale (POS) — ie what got sold — as POS data can provide powerful insight into the in-store customer journey. For a grocery store, understanding a shopper's cart is almost like having a map of their path through the store. Grocery stores, however, benefit from two simple factors. Almost every shopper that goes into a grocery store buys something, with most buying quite a few items. As a result, POS data provide a pretty fair approximation of shopper journey and interest.

In most other retail situations, these two factors do not hold. For most retailers, less than half the people who enter the store buy something and, in most cases, the average basket is quite small. When this is the case, most of the in-store shopper journey cannot be inferred from POS data.

Thanks to door-counting, stores generally have a good idea about the number of

customers that have come through the door on any one day. By combining these entrance counts with POS data, stores can infer store conversion rates. However, as digital marketers know, high-level averages at the store or site level tend to hide all the interesting detail.

It would be nice to understand so much more about the shopper journey — from foot-traffic at every area of the store (down to every square foot), to where shoppers spent time, what engaged them along their journey, and the path they took between areas of interest. It would also be nice to tie all of this information to eventual conversion.

With modern people-measurement systems, this is all now possible.

PEOPLE-MEASUREMENT TECHNOLOGY

Overview

People-measurement technologies fall into two broad categories: electronic or camera. Electronic systems do not really track people at all — they track electronic signals. But if the electronic signal being tracked is from a device carried by a person, then they function as people-measurement devices. Camera systems, on the other hand, are looking for people. They work by analysing a captured image and identifying people in the field of view. Much like radar, light detection and ranging (LiDAR) technology forms a third category of sensor that works differently from both video camera and electronic sensor. However, as the performance characteristics of LiDAR map quite closely to camera, the technology tends to be evaluated as a special kind of camera.

Electronic detection can be passive or active. In passive detection, sensors look for signals being sent by devices carried by shoppers. By far the most common signal detected is a Wi-Fi network probe — a signal sent by a phone looking for a potential Wi-Fi network. Other signals sent by phone include carrier signals and Bluetooth Low Energy (BLE) signals. Electronic sensors pick up these signals and then use various fingerprinting or triangulation techniques to calculate the location of the signal source. By tracking signals over time, electronic detection can map an entire journey and even track the device over time (to understand repeat visits or visits from one store location to another).

Although their collection methods are very different, the data generated by camera and electronic systems are surprisingly similar. In each case, the core of the data consists of a timestamp, a location (an x,y coordinate in space), and a randomly generated person identifier. With camera, some additional data may be generated (like gender or height), but that is pretty much it.

On the other hand, this similarity masks profound differences in the way the technologies perform, the quality of the data they generate and the applications they support. People-measurement technologies differ widely in the percentage of the shopper population they track, the positional accuracy of the location data they deliver, the frequency of the measurements they collect, the ability to distinguish between shoppers and associations, their ability to track across large spaces or without direct line-of-sight, their ability to track visitors over time (repeat visits), and their ability to add additional data to the tracking (like gender or age). To really understand how a technology works and to which applications it is best suited, every one of these factors must be considered.

Population capture

The goal of most measurement systems is to measure everyone in a particular space. This is not always possible. To the machine-learning in a video camera, a person in a wheelchair might be mistaken for a cart, or one person may block the camera's view of a second person. Electronic tracking,

meanwhile, records only those people with devices that are broadcasting a trackable signal. Population capture is usually measured as percentage of the population one should expect to capture, with 100 per cent being ideal.

Positional accuracy

Positional accuracy measures how well a measurement system does in gauging a location of a person within a space. It is usually measured in feet or metres with an implicit probability (for example, there is a 95 per cent chance that the person is within one metre of the assigned position). Some types of people-measurement (eg door counting) do not require positional accuracy. For broader journey tracking, however, positioning is vital. And the finer-grained the position, the more types of analytics a system can support. The smaller the measured space, the more important positional accuracy is. In a football stadium, knowing where someone is with five-metre accuracy is good enough, but for a mall retailer, that distance might the difference between men's jeans and lingerie.

Tracking rate

Tracking rate indicates how often and how consistently the system measures the shopper journey. In other words, if a shopper spends ten minutes in a store, how many measurements will the system capture? The higher the tracking rate, the more detailed the journey. A system that captures the shopper position every second will miss almost nothing. A system that sees where the shopper is every minute or two may miss significant detail.

Associate detection

Differentiation between shoppers and associates is important. Knowing which people are shoppers and which are associates not only ensures accurate shopper counts, but also makes it possible to study how associates spend their time and how they navigate store processes. Key performance metrics like shopper-to-associate ratios can only be calculated when a system can reliably distinguish between customers and workers.

Full journey tracking

It is one thing to measure people as they move through a confined space, but technologies differ in their ability to expand that tracking across very large spaces. Some technologies do this seamlessly, while others require some extra work. Furthermore, in many cases large spaces create special data quality challenges that are important to understand. There is no single metric for measuring journey tracking capability; it simply encompasses a set of issues that must be considered on a space-by-space basis.

Repeat visitor tracking

Repeat visitor tracking is the time-based extension of full journey tracking. Full-journey tracking tells us everything a shopper did in a single visit to the store; repeat visitor tracking tells us whether the customer came back. No technology is perfect for this, but some are better than others.

Data enhancement

People measurement is behavioural measurement. However, analysts know that every kind of data can be valuable, and that demographic information is often a powerful supplement to core behavioural data. Gender is the most common additional data point, but in addition to demographics, some technologies can also create powerful join strategies to POS or customer data.

Electronic sensors

There are significant limitations with electronic tracking. Without using unfair

(and probably illegal) means, electronic capture is restricted from fingerprinting phones. It must rely on phones pinging out publicly and identifying themselves in a stable manner while doing so. Unless connected to Wi-Fi (which is declining in use), modern phones will not do this. So, electronics usually track somewhere between 20–30 per cent of the population. Vendors often mask this fact by 'truing up' numbers, but this is more make-believe than measurement. Another common trick is to count every device regardless of whether or not it has stable identification. This is deceptive when it comes to high-level space usage and flat-out useless when it comes to full-journey analytics.

Electronic tracking also presents significant limitations in terms of positional accuracy and tracking rate. Electronic signals bounce around a lot, and using their strength to position shoppers is a dicey business. Really good electronic tracking can get positional error rates down to a radius of about three metres, but typical electronic tracking generates a positional radius of about ten metres. This is not bad for airports or stadiums, but it is problematic in retail stores. A less well understood issue with electronics has to do with tracking rates. It is only possible to track electronic devices when they choose to broadcast their presence. For devices that one controls oneself, this is not an issue. Shoppers' phones, however, tend to ping only every minute or two. This can (and does) leave significant gaps in the journey.

On the other hand, electronics excel at full journey tracking. With a stable device identifier, one can confidently track a device over very large spaces (whole cities) and over substantial periods of time.

Camera sensors

People take for granted their ability to distinguish between people, shelves, grocery carts, boxes and even mannequins, but for a computer, doing these things is not easy, and the more person-like the object, the harder it is for the computer to classify it correctly. Fortunately, machine-learning techniques to identify people in images have improved dramatically in the last five years. This means camera technologies can be relied on to capture nearly every person in a space with high positional accuracy.

Cameras also do a superb job with tracking rates. Most camera sensors are capturing at multiple frames per second and are outputting data at least once every second. This means there are no gaps in the shopper journey.

On the other hand, camera sensors can struggle with full journey tracking and repeat visitor tracking. While modern camera systems can track shoppers across very large fields of view (by matrixing multiple cameras into a single array), they can be defeated by very large or very cluttered spaces. They lose track of people when line-of-sight gets blocked either by physical obstructions or other people. This typically results in breakage across the journey, making some kinds of journey metrics unreliable.

Most camera variants do provide at least some demographic enhancement, particularly with respect to gender. Keep in mind, however, that measurement cameras are typically ceiling-mounted for a top-down view (which maximises their ability to track people). This is not ideal for demographic classification, so both gender and age classification are less reliable than one would expect from cameras mounted with face-on views.

Figure 1 presents a functional comparison of camera and electronic sensors for people-measurement.

It is worth noting that what camera is good at, electronics does very poorly — and vice versa. For this reason, there is often a strong case for deploying both technologies

Why Multiple Technologies?

Individually video and electronic sensors have capabilities gaps. Together they create a complete, comprehensive measurement solution.

Capability	Camera	Electronic	Notes
Full Population Reporting	✓		Camera will pick up 98% of visitors, electronics about 30%
Display Tracking	✓		Electronics lacks the positional accuracy to measure display
Intra-Area Tracking	✓		Electronics lacks the positional accuracy to track inside areas
Occupancy Management	✓		Electronic population capture is far too imprecise for occupancy
Queue Management	✓		Electronic population capture and tracking rates are insufficient
Full Journey Tracking	✓	✓	Cameras will generally lose track of a visitor at least once or twice during a visit. Electronics hold visitors but aren't as detailed.
Repeat visitor Tracking		✓	Camera cannot track people across visits
Staff Tracking	✓	✓	Camera has to guess at staff identification and cannot tie out at the individual level.

Figure 1: Functional comparison of people-tracking technologies

to answer a full range of measurement questions.

Figure 2 compares the performance of the main people-tracking technologies in greater detail.

Finally, it is worth noting that measurement cameras are the most popular choice for full-store analytics because they do the fundamental measurement tasks (population capture, position accuracy and tracking rate) really well. However, measurement cameras require good, stable lighting in addition to overhead mounting points at a reasonable height; they are also quite expensive to deploy in very large areas. Until recently, that meant that for very large stores and other complex spaces (malls, airports, museums, etc), electronic technologies were the only viable option for people-measurement. Because LiDAR provides indoor/outdoor, all-weather tracking, flexible mounting, and significantly improves coverage per sensor, it has opened up high-accuracy journey measurement to a number of previously impossible people-tracking locations.

APPLICATION AREAS

The combination of high-accuracy and good tracking rates across very large fields-of-view has created new analytics opportunities including occupancy management, queue management and merchandising analytics — none of which were possible with either electronics or single camera systems.

Occupancy management

COVID-19 transformed the market by making people-measurement an operational necessity. Existing door-counting systems were quickly converted to occupancy management — not always with sterling results. The two keys to successful occupancy management are high accuracy and real-time monitoring. While door-counting accuracy has improved markedly in the past few years, the accuracy requirements for occupancy are much higher. A door count system with 2–3 per cent error rate is fine for trending door-count numbers but may produce occupancy counts that are off by 30–40 people. This is simply unworkable.

Technology Capabilities

Camera and Sensor have very different (and largely opposite) performance characteristics for store measurement:

	Camera	LiDAR	Electronic
Positional Accuracy	1ft	1-2 inches	10-20ft
Population Capture	99%	99%	20-30%
Capture Rate	0.5 Seconds	0.5 Seconds	15-120 Seconds
Display Interaction	Possible with Extra Work	Possible with Extra Work	Not Possible
Associate Detection	Inferred	Inferred	Yes
Largest Practical Area	100,000 square ft.	500,000 square ft.	Any
Over Time Tracking	No	No	Yes
Outdoor/Low-Light	No	Yes	Yes
Data Enhancement	Gender, Height, Age, Cart, PoS	Height, Cart, PoS, Vehicle	

Figure 2: Detailed capability comparison of people-tracking technologies

Fortunately, the best camera and LiDAR systems today deliver people-counting accuracy of around 99 per cent. This is probably better than most people would do by hand-counting, and it supports digitally driven occupancy management applications. Occupancy management also illustrates one of the main trends in modern people-counting; increasingly accurate people-measurement that is sufficiently precise to drive specific interventions and operational change — not just high-level analytics and reporting.

Queue management

Queues at cash-wrap are one of the highest friction points in the shopper experience — and one of the most tunable. By adjusting registers, staffing and checkout strategies, stores can significantly improve queue management. Analytics can drive such change, but as with occupancy management, queue management is a great place for direct operational interventions.

By coupling occupancy and predictive queue management, it is possible to forecast the state of queues in the next 5, 15 or 30 minutes. That means store managers can make staffing decisions to head-off queue problems or reassign staff to more productive work. Precise queue measurement also allows the integration of virtual queuing techniques into queue management. Instead of having customers stand in line at customer support desks, they can use SMS to enter a virtual queue. Queuing systems can track their position and automatically ping them when it is their turn. Not only does virtual queuing provide a better customer experience, it maximises shopping time while in store.

Merchandising analytics

Display is the heart of the store experience, yet merchandising has been almost unmeasured. With camera and LiDAR, however, population capture, positional accuracy and continuous tracking are all

good enough to support very detailed display analytics. Camera and LiDAR systems can count how many shoppers passed by a display, how many stopped, how long they spent, and even where they stood. This provides a powerful platform for measuring merchandising effectiveness in the store. LiDAR technologies can even support specialised analytics of key interactions like product touches or product takes.

Full journey

The advent of matrixed cameras and LiDAR has dramatically improved the accuracy and detail of full journey tracking. In even the largest spaces, it is now possible to get step-by-step journey records. This provides remarkable opportunities for improving shopper flow at every level of the store, from individual displays to departments or even whole floors. Full journey tracking even supports segmentation of specific flows in terms of broader shopper behaviours. One can look at the flow of shoppers in an area, then isolate buyers, or compare flows for people who came from one direction or area vs shoppers who entered from a different direction. By moving the level of layout measurement from the store down inside sections and all the way to specific displays, the power and utility of modern people-measurement systems have expanded dramatically.

People-process optimisation

Associates are a key part of the store experience. They are also typically the largest variable cost for most stores. For this reason, good associates executing good processes can be a massive competitive differentiator. With COVID-19, many stores had to make significant changes to how associates operate. This is hard to do well at scale. Fortunately, people-measurement techniques work on associates as well as shoppers. Measurement systems that span back-office operations and front-of-house make it possible to understand how associate time is allocated, how processes are being executed, and where different processes are going wrong. With so many new processes integrated into the workflow and with even higher levels of churn, finding ways to measure and optimise people processes has become easier to do and far more important.

PRIVACY

The privacy story in people-measurement analytics is surprisingly good. Most video and LiDAR systems do not collect or save *any* personally identifiable information (PII) so as to comply with privacy standards such as the General Data Protection Regulation. Electronic detection often collects basic phone identifiers that may be considered PII. Unless an opt-in is obtained from the consumer, this information should be discarded. This is particularly important if the geo-location data are tied to the customer record.

For camera and LiDAR systems set up in the standard fashion, there is no need to post a shopper-facing notification about measurement; likewise there is no need for an opt-in or opt-out system. In fact, neither opt-out nor opt-in is possible because no PII is ever collected and it is impossible to tie an opt-out to the collected data.

JOINS

Given that most people-tracking technologies do not collect PII, joining in-store data to electronic data is problematic. With the exception of mobile app tracking, there is seldom any way to tie shopper behavioural records to a known individual or to their digital exhaust.

However, there is one technique that is available to stores. With full-journey camera and LiDAR measurement, it is possible to track an individual to a specific register

and time of day. Using the timestamps embedded in register data, it is easy to join register data capturing the exact purchases a shopper made to the complete track of the shopper journey leading to that register at that time. Doing this join makes it possible to analyse which specific parts of a journey resulted in a sale — a powerful extension of the basic journey data.

If, an individual's identity is known or captured at the point-of-sale, then it is also possible to tie that identity to the rest of the shopper journey. This is a powerful analytic technique, but it is usually problematic from a privacy standpoint. In Europe and in a growing number of US states, making this tie would require informed opt-in from the shopper. This might be achievable as part of a broader loyalty or customer service initiative, and it provides remarkably detailed shopping behaviours at the customer relationship management profile level.

SUMMARY

In-store analytics has benefitted from the rapid advancement of technologies for sophisticated people-measurement. Camera and LiDAR technologies, in particular, have dramatically improved their accuracy, coverage and matrixing capabilities, while LiDAR further extend the range of environments that can be measured. However, even full-matrixed camera systems have some gaps when it comes to shopper journey measurement. Camera systems cannot do repeat visitor tracking, often lose sight of people in longer journeys, and still struggle to identify associates. Because of this, it makes sense to deploy additional electronic measures to collect the data necessary to fill in some of these gaps — particularly around associates.

As the precision and coverage of people-measurement technologies has improved, new applications have emerged — often with significant operational implications. Door counting has evolved into real-time occupancy management. Queue measurement has evolved into predictive queue management and virtual queuing. On the analytics front, the higher journey resolution of camera systems has enabled merchandising and display analytics programmes that are remarkably similar to common A/B testing regimes. When it is possible to measure shopper movements at the square-foot level, it enables detailed studies of display effectiveness. Similarly, full-journey tracking via camera has enabled in-store flow and layout analytics to tackle much finer-grained areas of the store. Finally, the high-precision, full-journey tracking available from camera has been repurposed for associate process optimisation. People process optimisation provides a new level of analytics and compliance monitoring to labour allocation and process design — a very good thing in the post-COVID world of high-churn and increasingly complex in-store processes.