

Engagement as visual attention: A new story for publishers

David Bassett (Lumen Research) and Andrew Green (Ipsos Connect)

Introduction

This paper is about measuring reader engagement. It is not a new topic. A total of 96 papers have been delivered at previous symposia on the subject of either ‘quality of reading’ or ‘engagement.’ WARC lists 1,373 research papers on the topic. 82 of the 107 surveys listed in the most recent Summary of Audience Research Worldwide (2013) feature engagement questions of one sort or another.

Despite all this, defining and measuring engagement in a way that is commercially relevant has proved elusive. Numerous engagement metrics have emerged, but to the frustration of both publishers and advertisers, none have stood out as being universally applicable for trading. In this paper, we examine the potential for eye-tracking to provide a tradable engagement metric by measuring *engagement as visual attention*.

Using eye tracking to measure engagement has a natural appeal. After all, advertisers regularly talk of “buying eyeballs”, so actually trading them is a logical next step.

Eyeballs are also media-neutral. Unlike other engagement metrics, which tend to be tied to particular media and platforms (clicks, swipes, video plays, shares, likes, etc.), visual attention can be measured across channels using the same tools, allowing advertisers to compare the full gamut of visual media, from press, to magazines, TV, Out of Home¹, desktop, tablet and mobile.

Until recently, the cost of eye tracking prohibited its use in readership measurement. But technological developments have dramatically improved its scalability. In future it is likely that everyday devices will have eye trackers built into them, further expanding the technology’s potential. Quite how ubiquitous the technology will become is difficult to predict but, even with today’s technology, a lot more is possible than was the case just a few years ago.

Of course like any methodology, eye tracking has limitations. Even though more scalable than it has been in the past, it would still be cost-prohibitive to employ eye tracking with a very large sample or using a strict random-probability approach. At present, the most cost-effective way to employ the approach is via central location studies – which restrict the research in practice to urban areas. And though the experience of participants (as we will show) is much improved, it is still inevitably somewhat of an artificial environment we place them in, as we are asking them to read in our time on our premises.

But in our view the advantages outweigh the disadvantages and ways can doubtless be found of linking eye tracking data with survey data in ways which are outside the scope of this paper, but which may be considered by those interested in the approach.

To demonstrate what eye tracking offers, we discuss its scalability and draw on some data collected by Lumen Research. Lumen runs a weekly omnibus eye tracking study to help UK brands optimise the design of their advertising. In doing so Lumen has built up a database that reveals how people read newspapers. Below we present some key facts from this database that suggest eye tracking’s potential for measuring engagement:

- Eye tracking measures “eyes on” dwell time *with advertising*, as well as editorial: on average 73% of press exposures are viewed, usually for around 2-3 seconds of dwell time.

¹ In fact, eye tracking is already used by the UK’s out of home JIC, Route.

- Dwell time is correlated with ad effectiveness: just 2 seconds dwell time can drive recall
- Eye tracking quantifies the contribution publishers make to advertisers: editorial engagement is strongly correlated with attention to advertising
- Eye tracking quantifies differences between titles and platforms, across different audiences; the data varies strongly across titles and audiences

Accordingly, we suggest that eye tracking is an option worth serious consideration for measuring audience engagement. By measuring engagement as visual attention, advertisers can know how much an audience actually engages with their advertising and publishers can be rewarded for their contribution in providing engaged readers.

Before we go on to present these data though, let us step back for a moment to assess the state of play with engagement metrics today. What metrics are currently available, and why have they struggled to find a place in day to day trading?

The State of Play: Existing Engagement Metrics

A vast array of engagement metrics have been employed over the years by content publishers. In the print area, examples include two *Quality of Reading* surveys undertaken in the UK between 1997 and 2000 and the emma engagement study from Australia. emma, for example, asked an extensive range of questions to readers around five ‘axes’ of engagement – Source, Loyalty, Motivation, Connection and Actions taken.

A key challenge with survey-based metrics is that only a very limited number of questions can reasonably be asked of a survey participant and it can be difficult to get the kind of nuanced perspective that a publisher might like. For example, when we ask people about how much they ‘trust’ a newspaper brand or how much they regard a magazine as ‘a friend’ – are the results substantively different to those we might get from looking at reader loyalty, the level of subscriptions or time spent reading? And how many different questions can we ask about peoples’ reading of a single title before they become frustrated and bored – especially if they are part of a much longer survey?

In the digital space, various metrics are collected and collated. A recent report by the IAB in the United States² listed thirty-three definitions of digital engagement, including cognitive (e.g. recall), emotional (e.g. impact on brand perceptions), and behavioural (clicks, likes, shares, etc.). Meanwhile, an extensive survey by the UK’s communications regulator Ofcom³ outlined even more behavioural metrics, subcategorising them further between loyalty, time-based and social metrics.

Figure 1: Examples of existing engagement metrics

Type	Metrics
Survey-based	Source, Loyalty, Motivation, Connection, Actions, Recall, Impact on brand perceptions, Impact on purchase, Trust, Closeness
Interactions	Clicks, Swipes, Roll-over time, Video plays, Shares
Loyalty	Page views per user per month Visits per visitor per month
Time-based	Average time spent on site Average time spent per page
Social	Comments, Likes, Follows, Shares Site referrals from social networks

² IAB (2014). *Defining and Measuring Digital Ad Engagement in a Cross-Platform World*

³ Ofcom (2014). *Measuring Online News Consumption and Supply*

As noted in the Ofcom report, these metrics are put to a broad range of uses – editorial, commercial and academic. For editors and academics, a broad range of metrics honed to particular purposes brings welcome nuance and insight into audience behaviour. But commercially, this large set of possible engagement measures can often produce more heat than light. Advertisers are keen to trade on the basis of engagement, but the metrics currently available are too disparate and difficult to compare across publishers and platforms (let alone to other media) to become a part of day to day trading in the way reach figures are.

Part of the problem has been that different measurement systems at times yield wildly different results for the same metrics⁴. But even if all the metrics available – survey-based measures of engagement, page views, time spent on site, and so on – were all available through harmonised measurement systems, there are more fundamental, conceptual issues that would still need resolving. The key question that needs to be asked of a *tradable* engagement metric is - *how are advertisers going to price this in?* For the most widely collected engagement metrics, the answer is far from obvious or forthcoming. How much more should an advertiser pay to a publisher with a loyal readership? What is the value of a like or share?

The real obstacle to commercial adoption, we believe, is that these metrics do not usually measure the level of *engagement with advertising itself*. This is arguably what advertisers really care about. In the digital space in particular, knowing whether an audience actually sees an ad has become *the* critical issue facing advertisers today. Keith Weed, CMO of Unilever, recently told *The Times* that he believes as much as \$6bn to \$10bn is being wasted on ads that are never seen⁵, as a result of ad fraud and viewability problems.

Existing engagement metrics don't satisfactorily address these issues. Many simply refer to how engaged an audience is with a publisher overall, but this does not have a straight line relationship with a reader's likelihood of engaging with the advertising. Even if – on the balance of probabilities – a more engaged reader is likely to be more attentive to the advertising they are exposed to, it cannot be quantified within these measures.

And while interaction-based metrics (click through rates, roll-overs, video playbacks, etc.) are concerned with interactions with the advertising itself, the actual volumes here are often vanishingly small and getting smaller by the day. With the average click through rate at just 0.1% (according to DoubleClick), one has to wonder about the remaining 99.9%. To ignore them would be to pass over significant value. Moreover, as the volumes fall, interactions-based metrics become increasingly unreliable and noisy, even as proxies.

The *Financial Times* (and others) have gone down different paths, introducing time-based inventory. Working with ChartBeat, inventory on FT.com can now be bought on a cost-per-hour, rather than cost-per-impression basis. While time-on-screen based metrics are arguably a step forward in the quest to address advertisers' viewability concerns, they still don't measure engagement with the advertising itself. Time-on-screen, after all, remains an 'opportunity to see' measure.

Eye-tracking, on the other hand, measures actual eyeballs on screen. Attention has always been the commodity being exchanged in media markets. Hitherto we've made do quite well with proxies through the form of OTS measures, but through eye tracking, there is the prospect of going beyond the O and measuring the S.

But what, exactly, is eye tracking, and how does it work?

⁴ <http://fivethirtyeight.com/features/why-we-still-cant-agree-on-web-metrics/>

⁵ <http://www.thetimes.co.uk/tto/business/industries/media/article4488659.ece>

What is eye tracking?

Eye tracking involves the measurement of eye movements as someone looks at visual stimuli. Although our perception gives the impression that the eye moves reasonably smoothly across our visual field, in reality it proceeds through series of '*fixations*', whereby the eye focuses on a single point, before moving on to the next fixation. Fixations vary in duration, but are approximately 250ms long (i.e. one quarter of a second), while the transition between fixations is called a '*saccade*'.

The eye gathers the majority of information about the world from fixations and, accordingly, measuring the aggregated dwell time of fixations is a reasonable measure of how much attention someone is paying to something (although, of course, it does not *guarantee* that the brain is engaged with the visual stimuli).

Eye tracking has a surprisingly long history. As far back as the 1800s, scientists were making direct observations of the eye and discovered the phenomenon of fixation. The first technologies to measure eye movements arrived in the early twentieth century, initially using a sort of contact lens and then latterly by projecting light on to the eye while recording eye movements on film. The field expanded greatly in the 1970s, with some famous studies of reading⁶.

This early technology was expensive, and looked somewhat frightening:

Figure 2: Yarbus eye tracker from the 1960s




Fig. 21. The apparatus used in recording eye movements.

Thankfully, the sophistication of modern cameras has meant they are no longer so intrusive and scary to use. Modern eye tracking now has two different forms: glasses-based and screen-based.

⁶ See Rayner (1998), "Eye Movements in Reading and Information Processing: 20 Years of Research", *Psychological Bulletin* 124 (3), pp. 372-422.

Eye tracking glasses

Eye tracking glasses contain two types of camera: a normal video camera, pointing outwards, which records a film of the general area that respondents could have seen; and a pair of eye tracking cameras, pointing inwards, which record the respondents' eye movements. These two films are then automatically layered upon one another, so that the researcher can see which locations on the 'outward' film the respondent actually looked at.

This form of eye tracking has found widespread use in in-store, packaging and Out of Home audience research. However, while glasses-based eye tracking is extremely accurate - since each film is particular to each individual - it is also a time consuming approach. Each person's gaze path has to be coded manually, limiting the scalability and relevance of glasses-based eye tracking for publishers' audience research.

Figure 3: Eye tracking glasses



Screen-based eye tracking cameras

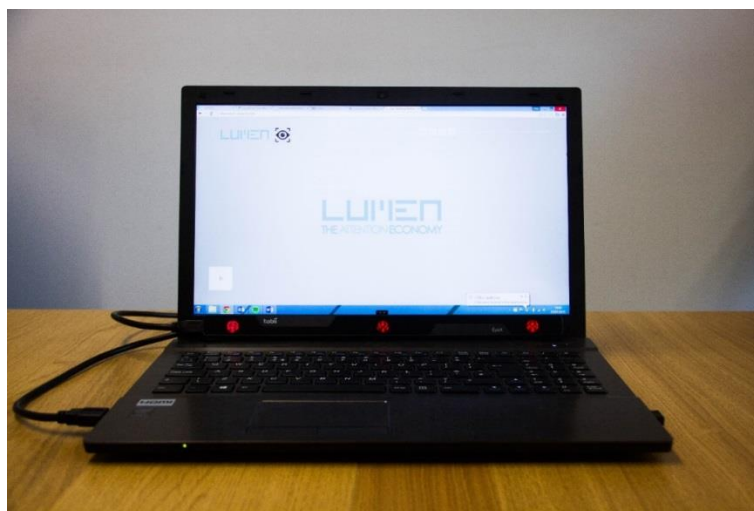
A second, and more scalable, means of eye tracking is with cameras that attach to computer displays, tablet screens and mobiles. These cameras (produced by Tobii, EyeTech and The Eye Tribe) record respondents' eye movements as they view stimuli on a screen, and then map it onto the stimuli automatically, negating the requirement for a human to analyse the data. This significantly increases the speed and scalability of the approach.

These cameras are highly accurate, capable of recording a fixation to within roughly 1cm. They are also very reliable – typically over 90% of participants in a study produce useful data.

Their accuracy and reliability is a product of several features that make these devices specialised: they process images at a very high frequency (Lumen operates satisfactorily at 30Hz, meaning the cameras refresh 30 times a second), and use multiple cameras and infrared light sources to enable depth perception and automatic adjustment of the geometry of fixation measurement.

It is these latter features that make modern screen-based eye trackers unobtrusive. Users can move naturally over quite a large area – both side to side and back and forth – and fixation measurement remains accurate.

Figure 4: Screen-based eye tracking



Webcam-based eye tracking

In recent years, attempts have been made to conduct eye tracking using webcams rather than specialist eye tracking cameras. This method has the advantage that webcams are already ubiquitous, but the use of single, low resolution cameras, without infrared light sources, results in inaccuracy and high failure rates.

The main problem is a lack of depth perception, meaning participants have to sit lock still in order for measurement to be accurate. Small movements either forward or backward result in inaccuracies or data loss altogether, as the geometry of the calibration is lost. As a result, recording sessions with webcams have to be extremely time limited to ensure the data remains robust.

Scalability

Eye tracking has been around for a while as market research technique. It has found extensive use in user experience and in-store research. It is also already an integral part of the UK's Out of Home audience research JIC, Route, where eye tracking has been used to create indices to adjust the prices of frames. Out of Home frames are now routinely valued not just according to the likely opportunity to see, but also the likelihood of actually being seen, based on formats, positions and orientations. The example of Route demonstrates that measuring how successfully media captures visual attention can form part of a media valuation and trading system.

Until recently, the cost of eye tracking has meant it has not been scalable enough to be applied to audience research for publishers. But the cost of eye tracking cameras has declined sharply in recent years, and is set to fall even further. Moreover, the leading eye tracker manufacturers are working with computer manufactures to integrate cameras into everyday devices, just as webcams have been⁷.

How ubiquitous eye tracking cameras become is somewhat uncertain at this point. Adoption will not be led by market research, of course. As the cameras become more available, more and more applications are being found that take eye tracking from a psychological research method to a control input. It's difficult to predict exactly how far this process will go, but the most obvious applications are in improving access to computing for those with disabilities and in gaming. For instance, *Call of Duty* and *Assassin's Creed* are already available to play using eye tracking, with more games on the horizon. But even if adoption of eye tracking cameras as control inputs turns out to be fairly limited, the falling cost of hardware means they are becoming available at an affordable scale for audience research.

To capitalise on this availability, audience researchers would need to develop systems to capture, process, analyse and report on the data in ways that are meaningful to publishers and advertisers. Lumen Research has already taken several important steps towards developing such systems. Founded in 2013, Lumen has taken eye tracking from its roots as a small scale qualitative research technique to a quantitative one, and is in the process of developing an online panel-based system.

Through 2014 Lumen was supported by a grant from the UK government's innovation arm, *Innovate UK* (part of the Department for Business, Innovation and Skills) to develop an ad scoring system for pre-testing print advertising. Aside from building the technological infrastructure, the company built a normative database of the performance of print ads against samples collected through street intercepts.

⁷ For example, MSI recently announced the launch of a laptop integrated with a Tobii eye tracker, targeted for gamers (<http://microsoft-news.com/msi-announces-gaming-laptops-with-tobii-eye-tracking-technology/>). Intel are also significant investors in Tobii (<http://techcrunch.com/2012/03/15/intel-capital-invests-21-million-in-swedish-eye-tracking-tech-company-tobii/>). Samsung have also recently patented eye-tracking enabled televisions, which aid reading of subtitles (<http://www.dailymail.co.uk/sciencetech/article-3213143/Never-struggle-read-subtitles-Samsung-patents-eye-tracking-technology-ZOOM-sees-squinting-TV.html>)

This database allows their clients to benchmark performance of ads across categories, for all the main print formats and for each page. Lumen are successfully commercialising this to offer a weekly pre-testing service to clients. They now have a database on how people pay attention to press adverts, based (in July 2015) on more than 60 studies, over 4,500 study participants and over 1,600 ads.

In 2015, again with support from Innovate UK, Lumen is building the infrastructure to take eye tracking out of the lab and into people’s homes. By the time of the conference, we expect that a trial panel will be up and running with around 500 households, each kitted out with a Tobii eye tracker, and a browser extension that gathers their data. Lumen aims to trace both on and offline sales effects amongst this panel back to attention to advertising, to demonstrate the brand effects of display advertising: “the display value of display advertising”.

All this is to say that eye tracking is in the process of becoming a scalable, cross-platform option for measuring engagement. Engagement with print titles (both press and magazines) can be measured in laboratory research at an affordable scale, while in-home panels are set to provide high frequency data on desktop browsing. For now, tablets and mobiles can be researched in laboratories, with panels a likely option over the coming years as the technology is further integrated into everyday devices.

So eye tracking is becoming a scalable proposition, but how could it be used to provide tradable engagement metrics?

Visual attention as a tradeable engagement metric

As discussed earlier, for engagement metrics to be tradable, they must be relevant to advertisers. This means they must have a few important features:


- They must measure engagement with the advertising itself
- They must be correlated with ad effectiveness
- They must show the contribution the publisher makes
- They must highlight differences between titles and platforms, across different audiences

Using data from Lumen’s omnibus research, below we discuss how eye tracking fulfils these criteria.

Engagement with advertising

Eye tracking’s great strength is that it reveals how readers divide their attention between editorial and advertising. Figure 5 shows a typical “gaze plot” of a reader, showing the fixations as someone reads across a double page spread.

Figure 5: Gaze plot of a typical reader



As one might expect, newspaper readers' attention is usually directed primarily at the editorial – after all, this is the main reason people read newspapers.

In fact, one of the biggest findings from Lumen's dataset is that people are very good at avoiding advertising. Sometimes, people don't view the advertising at all, as in figure 6.

Figure 6: Gaze plot of a reader avoiding advertising



Thankfully, this last gaze plot is fairly rare. Some readers are more inclined to avoid advertising than others, but most people view ads at least some of the time, and especially when they are relevant to them.

The gaze plots also show some other interesting features of how people read newspapers. People have well established routines for reading newspapers: they read most of the headlines, and use them to decide whether or not to read the article further. Having moved on to the lead, they will then continually assess whether to carry on reading, or skip to something else, as they proceed through the article. Quite often people read the first paragraph or two and then move on. And of course, quite often, people read the whole article.


The kind of “ad blindness” the reader displays above is a real phenomenon. Opportunity to see does not always equate to actual seeing. Across all the press ads Lumen have tested, only 73% are actually viewed at all (i.e. have at least one fixation). In other words, 27% of exposures go ignored altogether. The proportion that do look at ads – what Lumen call **standout** – varies substantially for different format, titles, advertised categories, and so on.

Even when people do look at an ad, they typically only engage with it for a short period of time. Mostly, people give ads a quick scan, assess whether they are relevant and interesting, and then carry on reading the editorial. A few go on to view ads for longer – maybe 5 seconds or so – but this is rare. A typical level of **engagement** with press ads – which Lumen defines as the average dwell time of those that view an ad – is around 2-3 seconds, but again there is much variation.

Even low levels of engagement with advertising has a positive impact for brands

One might wonder whether a few seconds of engagement is long enough to have any kind of impact, but Lumen's data show that it assuredly is. Alongside the eye tracking data, Lumen also routinely asks readers what ads they remember, providing a modest measure of advertising impact. In analysing the data it is clear that a few seconds attention is enough to have a genuine psychological effect.

Figure 7: The “2 second rule” of advertising recall.



In fact, there is a “two second rule” to ad recall: readers that view an ad for longer than 2 seconds are more likely to recall an ad than not. Interestingly, beyond 2 seconds recall doesn’t increase any further. Some ads do get much higher recall figures, but this is associated more with an ad’s design, and the salience of the brand.

That ads have impact at low levels of engagement is consistent with ideas such as Robert Heath’s “Low involvement” theory of advertising effectiveness⁸, as well as views expressed in Byron Sharp’s *How Brands Grow*⁹ We do not claim that visual attention captures everything though. Engagement is a multi-faceted concept after all. Emotions are also an important element to effective advertising and different media environments will provide advertisers with access to audiences in different moods and at different points in purchase cycles.

But the emotional side of engagement is conceptually nuanced and complex. The relevant metrics will be different for particular advertisers and campaigns, making it hard to see how a trading system could be built around them. On the other hand, some degree of visual attention is a definite, quantifiable and necessary precondition for any visual advertisement to work - a necessary but insufficient condition for it to be effective – with a proven link to ad effectiveness.

Attention to editorial drives attention to advertising

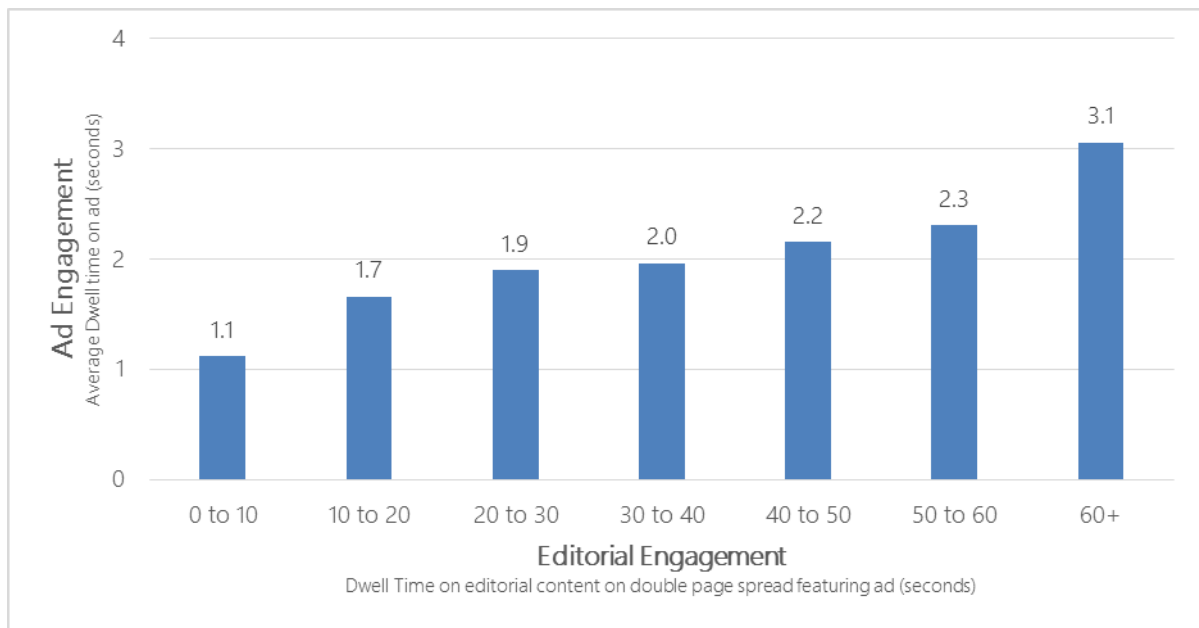
Engagement with advertising is influenced by many factors. Some of these publishers influence, some they don’t. Naturally, one of the biggest drivers is the quality of the creative. The advertised category also makes a difference: some categories (the latest blockbuster films) are simply more interesting than others (finance), or appeal to wider audiences. Other products are only interesting to people when they are in the market for them.

But certain factors are down to the publisher. Some formats naturally get more attention than others. And, as shown below, the appeal of editorial can make a big difference to engagement with the advertising. The longer someone reads the editorial surrounding an ad, the more likely they are to view it, and the longer they will engage with it for.

⁸ Heath, R. (2001). *The Hidden Power of Advertising: How Low Involvement Processing Influences the Way We Choose Brands*

⁹ Sharp, B. (2010). *How Brands Grow*

Figure 8: Attention to newspaper editorial correlates with attention to advertising



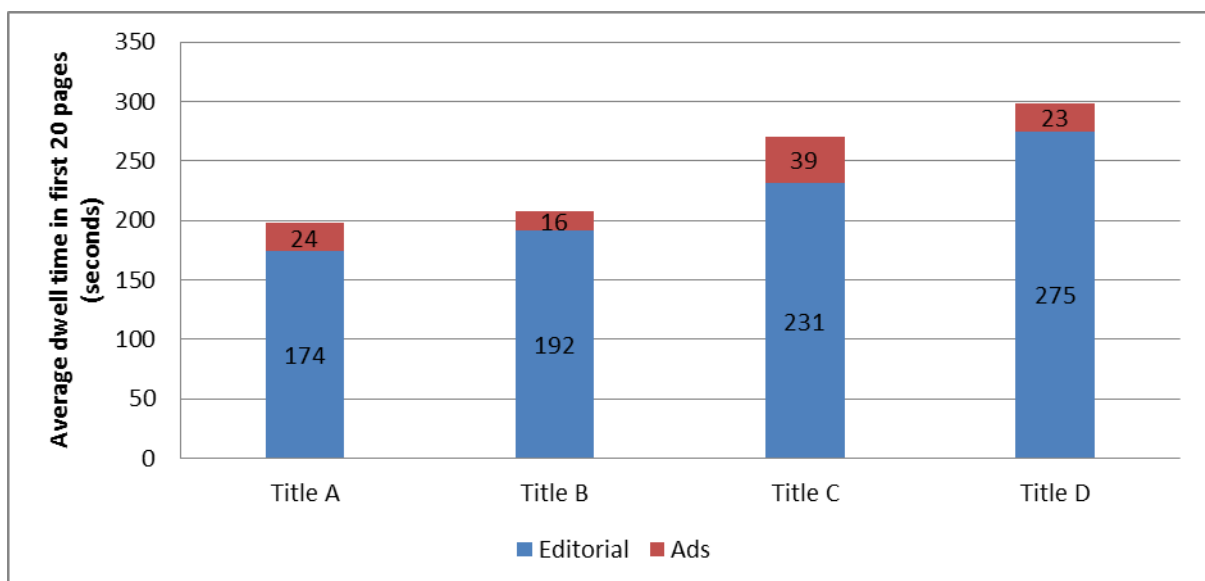
n = 27,918 observations, based on 390 25x4 ads

In other words, Lumen’s data demonstrates the basic interruptive model of advertising: advertisers effectively piggy back on the attention that publishers create. And as a result, there is a simple equation: more attractive editorial = better advertising impact.

Engagement varies by title and audience

The omnibus studies have used a selection of UK titles, revealing how different titles engage the nationally representative samples in the studies more or less strongly. Some titles have more engaging editorial to this broad audience, and as a result, tend to have more engagement with the advertising.

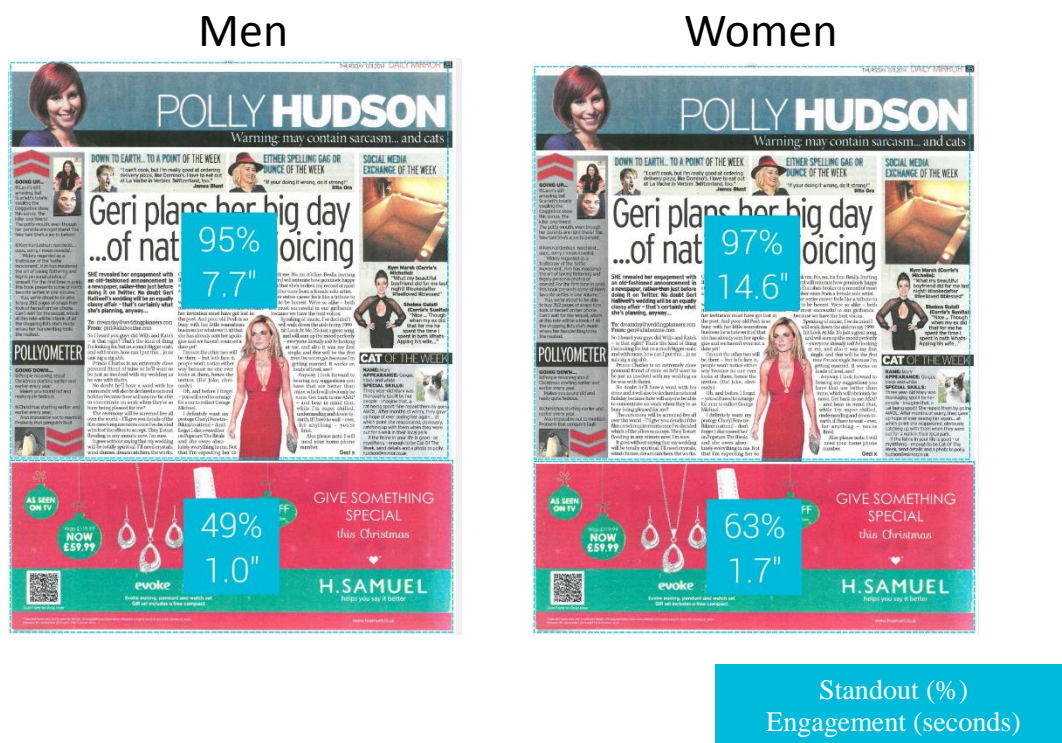
Figure 9: Titles have different levels of engagement with editorial and advertising



Of course, since this data is from nationally representative groups, it is inapplicable as a true measure for how well these titles engage their readers¹⁰. Including eye tracking in an audience measurement JIC would use samples of each publications readership to measure engagement with both editorial and advertising.

The equation of more engaging editorial = better ad impact also means that there is enormous value to matching ads to content relevant to the advertisers' target audience. Since different content appeals to different audiences, ad engagement will be more likely - and deeper – when it appears in the context of relevant content.

Figure 10: Engagement with content varies by audience – and this drives ad engagement



A deep understanding of a publications engagement would therefore come from studies which examined how effectively different titles engage different audiences, and accordingly how likely they are to engage with advertising. The resulting data would allow advertisers to refine buying strategies towards publications that provide high engagement among their target audience. Going even further, publishers could seek to maximise value by measuring engagement through different sections of their titles.

How might market participants use visual attention metrics?

So, we have seen how eye tracking can provide a means to measure engagement in a way that is relevant to advertisers. It provides metrics that are correlated with ad effectiveness, demonstrate the value that publishers make, and can show the relative strengths of different titles across different audiences. But how precisely would these metrics be used?

In a recent review of the publishing market, an Ofcom report quoted an unnamed publisher as calling for an engagement metric that would work as follows:

¹⁰ The nationally representative group is used to enable advertisers to optimise their creative against a broad audience.

“Ideally, we’d like one composite measure which was (some form of audience size) x (some form of engagement measure)” (National newspaper)

Measuring engagement as visual attention is set up to deliver exactly this kind of metric for the market to trade. We put forward the idea of **Gross Engagement Minutes**, or GEMs, defined as:

Gross Engagement Minutes (GEMs) =

Reach (Audience size) x Ad engagement (Av Dwell time per exposure/impression)

This will measure the number of minutes of attention an advertiser has bought, grossed up to the target audience being sought.

Exactly how granular and frequent measures of engagement can be made remains to be seen. It is conceivable that with a big enough panel, digital GEMs could potentially provide market participants with overnight data (analogous to BARB), and perhaps even become embedded into programmatic buying. For now, through a combination of Lumen’s more modest panel of 500 and in lab tests, they are able to give a view of likely engagement levels across publishers, platforms, formats and UK audiences, both digitally and in print.

What will be the impact of measuring engagement as visual attention?

To conclude, it is worth considering briefly what kind of an impact measuring engagement as visual attention might have.

The scale of the task for an effective engagement metric cannot be overstated. As audiences have moved online, the market has opened up to new competitors and a supply chain of ever greater complexity. Supply has swelled, often at very low prices. This cheaper inventory is usually understood to be of lower quality (because it’s more likely to be a bot, below the fold, unsafe, or “clickbait”), but because advertisers currently have no effective means of measuring quality in a comparable way, it is hard for media buyers to justify higher prices. As a result, quality publishers are struggling to justify their premiums.

The hope must be that a tradable measure of engagement can play a part in resolving this. The real indication that we have found a good measure of engagement will be when quality is able to be recognised and effectively priced in. It’s undoubtedly too early to promise that eye tracking can provide that measure. But it’s clear that something new is required.