

Measuring and Manipulating Audiences: A Personal Reflection

Dick C.A. Bulterman

CWI: Centrum Wiskunde en Informatica, and
Vrije Universiteit Amsterdam

ABSTRACT

Understanding the emotional reactions of audiences to a wide range of content types is an important area of research. In this article, I provide a personal reflection on various approaches to modeling, quantifying and understanding audience behavior based on a broad range of evaluation techniques. Using results from a study of the Heineken *Weasel* television commercial as a backdrop, I provide an overview of evaluation approaches and their impact in long-term and real-time evaluation. The main contribution is a personal reflection on audience evaluation based on multi-situation affinity with the area.

Author Keywords

Audience feedback; GSR; advertising impact.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI):

INTRODUCTION

Audiences are an important ingredient in creating successful performances. The audience is not only the target of performance content, but is also a vehicle that allows emotions and interest to be spread to a large community. Understanding audience reaction to content is important for content presenters, content producers, content distributors, and other content consumers. The content itself can be quite varied: it can be a play, a lecture, a sermon, a concert, the person across from you on a first date. Measuring and manipulating content recipient is often key to obtaining a desired goal. While the number of and nature of the goal(s) will vary, getting some feedback seems intrinsically more interesting than getting none.

Recently, as part of a research study conducted by a graduate student in our CWI group, I was asked to attend a

small-scale jazz concert in an intimate setting in my home city. About 50 people attended the performance, most of who were instrumented with a networked Galvanic Skin Response (GSR) sensor attached to two fingers on their left hand. I knew a few of the other audience members, but the venue was new to me, as were the performers. My wife, who was also wired for analysis, accompanied me.

The study was exploratory: given a small-scale concert setting with a set of tagged events, could audience feedback be obtained (in real time) that could be used to (a) characterize audience reaction to the event and (b) provide the basis for real-time feedback to the performers in a clear but unobtrusive (and non-threatening!) manner.

Each of our locally designed sensors had an ID number. After the concert, I asked one of our research assistants for a dump of the raw data for my wife and myself. I was interested to see if I, based on simple observation and personal intuition (rather than situational analysis and statistical prowess), could recognize some of my own reactions to the event and could correlate reactions that I gleaned from observing my wife to those represented by sensor data.

Figure 1 shows a graph of the GSR output for each of our sensors for the duration of the first half of the concert.

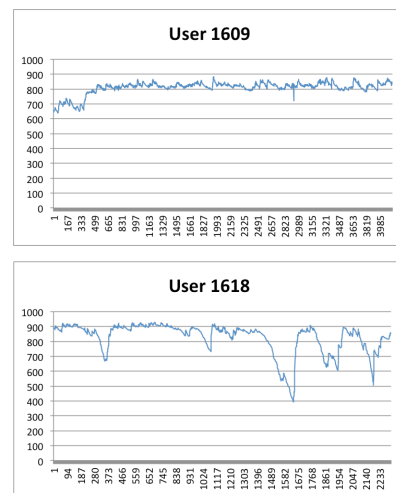


Figure 1: Two GSR Graphs

Note that the sensor with ID 1609 produced twice as many samples as sensor 1618, but that the host timestamps for the

first and last samples shown for both sensors was the same. When using these types of sensors, aligning data in the face of sampling frequency variances or occasional drop-outs is an important and non-trivial task.

For a trained data scientist, these data comparisons undoubtedly provide a wealth of inspiration for intuiting a broad range of significant correlations. As a subject-scientist (when it comes to this sort of analysis), I had more trouble understanding the baseline data and the meaning behind this data, even when adding my own temporal event markers (see Figure 2).

This lack of affinity with the reaction graphs of myself and my partner served as the motivation for a personal reflection on my own various experiences with understanding audience reactions and engagement. This article summarizes these thoughts and provides a vehicle for my own reflections on the value of analyzing the behavior of audience reactions, in the small and in the large.

This paper continues with an informal survey of interesting landmarks in gauging audience reaction to media content during the last century. I then reflect on a three-way study of audience reaction studied in the context of a television commercial. The paper concludes with some personal reflections on manipulating and understanding audiences.

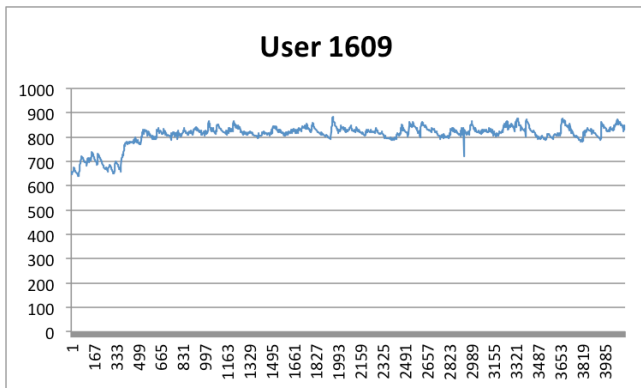


Figure 2: Informal Event Annotation Overlay

A SHORT HISTORICAL PERSPECTIVE

3-D Immersion, before 3-D

When I was a graduate student, I listened with fascination as a since-forgotten guest lecturer told of what she thought was the earliest documented study on audience engagement in the context of modern cinema. In 1896, the brothers Auguste and Louis Lumière produced and screened a short silent film that showed a train pulling into the station at La Ciotat, France [9]. Figure 3 shows a still image from this film.




Figure 3: A Frame From *L'Arrivée d'un train en gare de La Ciotat* [9].

The film was sensational, in a very literal use of the term. Reporting on the screening, an observer wrote that the screening “caused fear, terror, even panic” [7]. As I recall the story, the audience – not accustomed to seeing moving pictures – mistook the screening of the film for the arrival of an actual train, resulting in a panic run for the exits at the theatre. Now that’s audience engagement!

The measurement of audience engagement was, in a sense, quite informal: none of the participants were themselves photographed or connected to sensors. It is unclear if the ‘panic’ reaction was caused by event in the film or by cascading reactions to other participants [6]. Whatever the source, the fact that audiences were profoundly manipulated by the content seemed undisputable.

At the time, I can recall being puzzled by both the story of the reaction to the film and the ease with which the (otherwise critical) seminar audience accepted the premise of crowd panic. It was difficult for me to believe that a silent film in a probably otherwise quiet theatre could have such audience impact: a ‘real’ steam engines would have made lots of noise, have had a characteristic smell and would cause the ground to tremble under a viewer’s theatre chair. None of these characteristics were present. The people standing on the platform waiting for the train were no more excited than commuters waiting at the University Avenue station in Palo Alto, California for the arrival of an evening Caltrain Baby Bullet, where there is typically only panic if the train doesn’t arrive. As for the then modern-day audience listening to the seminar, it was interesting to reflect on the lasting impact that (this portion) of the presentation made, especially since none of us actually saw the film in question. An audience of scientists also can be easily manipulated, it seems.

As it turns out, the validity of the panic story has been called into question. The once-powerful example of audience engagement has degraded first to an urban legend [5], and later to a myth [3]. It now represents a general skepticism about a naïve audience’s sincere reaction to

unexpected content. I suspect, however, that the level of engagement was higher than we now are able to imagine. Early film screenings were often done at venues that also housed live performances: singers, comedians, jugglers, dancing dogs. The role of the audience was anything but passive. When early movies were first screened in New York, contemporary accounts talk of the electric atmosphere created by the *audience*, not the movie [11]! That some of this excitement exploded when as the recorded train arrived is at least plausible.

A Woman on 23th St.

Another oft-cited film work that highlights early audience reaction to events depicted on screen is Thomas Edison's 1901 film *What Happened on Twenty-third Street, New York City* [4]. A particularly famous fragment is shown in Figure 4, where we see the reaction of a young couple (unaware of the camera) when the woman's skirt is blown up by a blast of underground air. (The Marilyn Monroe remake of this image is significantly more popular.)

The conventional analysis of this fragment is that audiences, identifying with their on-screen counterparts, felt an embarrassment that consistently lead to audible shrieks and protective body movement (on the part of women) and a voyeuristic fascination with the content (on the part of the men viewing) [11]. The fact that these were real people caught in a real situation amplified the feeling of affinity of audience members.

Here, too, reality is a bit less powerful than the myth of the film. The two innocents on screen were actually actors and the film itself is probably more notable for being one of the first 'directed' productions in movie history. That the camera (which was significantly less unobtrusive than a smartphone) was visible can be seen from the reactions of 'real' real people in the film (particularly the young boy with the white shirt, who stares at the camera during the entire filming).



Figure 4: A Frame From *What Happened on Twenty-third Street, New York City* [4].

Still, the ability of an audience to knowingly engage itself with content (or to be influenced by the reactions of other audience members) was clearly demonstrated by this work.

Predictable reactions

During summers when I was in high school, I worked as an usher in Radio City Music Hall, a 5,000-seat theatre that offered patrons a variety show and a film screening as a package deal. The inside of the theatre is shown in Figure 5.




Figure 5: Radio City Music Hall, New York

During one period, I worked during the screening of the comedy *The Odd Couple*. In this movie, two men moved into an apartment together in New York (where even in 1968, rents were too high to occupy an apartment alone).

Being an usher is not a particularly intensive occupation. At the end of a show, people need to be moved out as quickly as possible so that, shortly thereafter, a new audience could be shown to the seats stamped on their tickets. During the actual screening, an usher would seat the occasional late-comer, but was otherwise free to view the film.

At one level, audience manipulation began before the film, as ushers chatted-up patrons in the (often unrealized) hope of receiving a tip. One particularly successful colleague always managed to inject the fact that that day was his birthday into a conversation during the short walk from the back door to the patron's seat. A more serious (and successful) attempt at audience manipulation came during the screening itself.

In the approximately 100 viewings of this film that I attended, I could sit outside the theatre door and track the run of the movie simply by listening to audience reactions to what was happening on screen. The durations and intensity of laughter were nearly identical every showing, independent of time of day, outside weather or even external events. It was this experience that started a personal fascination with understanding how audiences could be manipulated.

GAUGING AUDIENCE REACTIONS

In the previous section, largely anecdotal evidence for the presence of audience engagement was surveyed. The fact that audiences want to be engaged is, of course, nothing new. Airplane pilots routinely practice in on-the-ground flight simulators (even simple desktop ones), and yet can act and react in modes that are similar to when they are in the air. Audiences can be triggered to cry on command during dramas on screen or stage, and audiences allow themselves to be whipped up into a state of strident unity at political rallies (sometimes against their better judgment).

The literature is not particularly kind to these kind of emotional outbursts of audience interaction. Instead, there is a preference to more quantifiable measures of experience and engagement. In short: the intensity and frequency of applause is less important than the number of hands clapping in the crowd. Even these quantifiable measures are in transition, however. Table 1 reflects a change in how the appreciation for event ‘quality’ has evolved from an arts management perspective [13].

Audience Experience Measures	
<i>Traditional</i>	<ul style="list-style-type: none"> • Attendance numbers • #Showings/performances • # New works produced • Critical reviews • Peer assessment • Net income • Availability for audiences
<i>Proposed</i>	<ul style="list-style-type: none"> • Knowledge/information transfer/learning • Risk management • Authenticity and performer interaction • Collective engagement

Table 1: Measures for Quality in Performing Arts Based on Audience Experience [13].

While perhaps not definitive, the table does represent a more quality-based view of measuring content, which poses a problem if audience engagement is of interest.

From a personal perspective, I can recount at least one systematic attempt at gauging audience enthusiasm, which has positive and decidedly negative aspects. I play in a local jazz band, where enthusiasm and authenticity is in greater supply than technical talent. Every few months, our band invites another local band to participate in a Big Band Battle (BBB). Both bands play and the (supporter-rich) audience gets to decide who won, based on the reading of an audience applause meter. The setup is shown in Figure 6.



Figure 6: Two bands, one winner.

There are several aspects that influence the reliability of this direct form of audience measurement. First, the visiting BBB band often brings in fresh supporters who are (usually) positively influenced by a first meeting with the energy of the ‘battle’. Our own supporters, on the other hand, have heard *Come Fly With Me* 14 times previously and are less easily impressed. A second influence is that the audience measure is taken once, at the end of the BBB. Some of the fans have gone by then, others have arrived late and experienced only one band. This has the potential for skewing the results. These two factors play a role, but do not seem to dominate the result. During the last few battles, I have constructed an informal test to gauge audience behavior. It seems that the band that gets voted on second in the competition has the greatest probability to win. This has happened in four of the five recent events.

The explanation, I feel, has more to do with the audience members influencing each other than any inherent quality difference in the participants of the event [6]. For concerts like ours (and also for school plays, community theatre and half-time shows), the audience tends to listen with its heart rather than its ears. Thus, if the quality differences are small (which they often are), there is usually no strong artistic preference for one group above the other. What does seem to matter is the order in which contestants are presented. The first band can count on enthusiastic support from its supporter group and polite support from the others. The problem seems to be that this audience does not yet know

what ‘enthusiastic’ means: even if they applaud heartily, they have no idea if they are crossing some approval bound. Still, a baseline volume is set. When the second band is introduced, the audience seems to have a natural tendency to want to compete with itself. Since there is usually no strong artistic bias, the second band nearly always wins. (In the last 10 editions of the event, the second band has won 8 times.)

If the audience was asked to fill in a questionnaire upon exit the results might be more accurate, although the responses might be biased by the inherent politeness of our audiences. Using GSR sensors is an option, although since there is signification dancing and drinking during the event, it would be difficult to establish appropriate baselines for evaluation. It would also be very difficult to organize an evaluation structure that would help determine a reliable measure for user evaluation.¹

UNDERSTANDING WHAT TO MEASURE (AND WHEN)

One of the most complex aspects of performing any user evaluation is to know what to ask and when to ask it. Simple questions (which usually lead to wide participation) typically produce unverifiable results. Complex and repetitive question can insure robustness but often are a barrier to participation. Asking no questions but evaluating primary or second user responses may lead to less biased results, but correlating data from input sources with emotions for users (and user opinions) is daunting.

There are several popular approaches to measuring audience feedback. These include:

- Questionnaires
- Interviews
- Biometric Feedback
- Gesture, expression, posture evaluation
- Implicit action evaluation

(A combination of these is also possible.)

The use of questionnaires is time tested, but requires careful crafting for accurate results. Unlike more spontaneous measures, prompted thoughts appear to be less authentic [10]. Finally, questionnaires cannot be completed in real-time and thus has no potential to influence the event itself.

Conducting interviews provides the ability for a skilled interviewer to obtain deep results, but audience answers may be biased by social conventions or a lack of appropriate self-reflection. Again, real-time interpretation and integration into an event is impossible.

Biometric sensors can potentially provide a wealth of information that can be collected and (possibly) analyzed in

¹ After discussing these results, we decided to continue the policy of having visiting bands being introduced first.

real time. (In practice, real time collection and analysis is rare.) A significant challenge exists in attaching meaning to any sensor and to be able to filter out the larger number of irrelevant stimuli that can bias results.

Analyzing gesture, expression, posture, etc. seems a promising and non-obtrusive basis for evaluation, were it not that in many venues (such as our big band battle), audience members may sing, dance, walk and lounge during a presentation, each of which may introduce expressions not directly related to their quality assessment of the musicians. The venue itself may be dark (making capturing input impossible) and the sheer volume of audience members may make tracking difficult.

Up to now, interpreting implicit actions (such as applause, covering one’s ears or running for the exit) has been the principal indicator of quality or positive/negative engagement. Still, we see a tendency in the arts in Table 1 to move away from these measures.

At the high-tech end of the measuring scale, the direct monitoring of neurological signals is becoming popular, based on a belief that fMRI scans (or equivalents) can localize brain activity that can, in turn, be mapped to specific emotional responses. Figure 7 shows a scanner and one interpretation based on commercial preferences. Even if one believed that these measures were reliable, repeatable and representative, significant problems remain: hosting a jazz concert where all of the audience members were placed in fMRI cocoons would probably be a rather niche event.

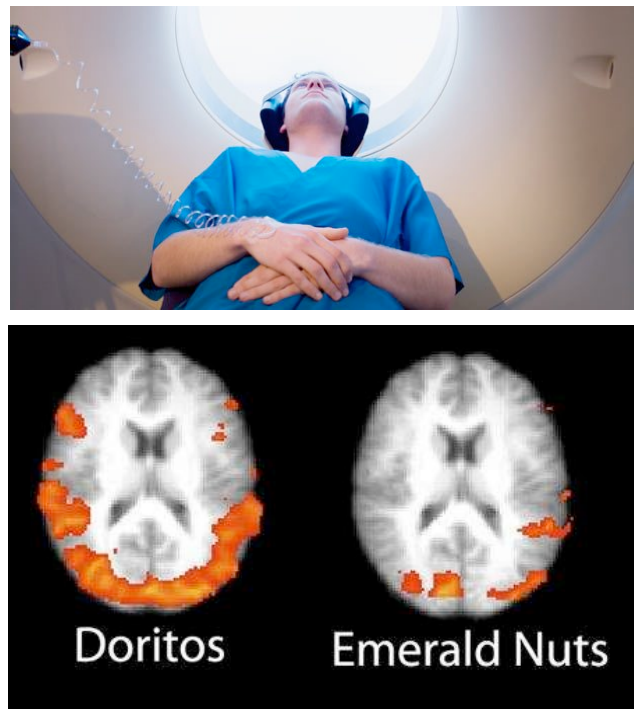


Figure 7: Using fMRI to measure and interpret preferences.

Understanding User Response to Commercials

The GSR sensors deployed at the jazz concert discussed in the Introduction section of this paper have proven to have potential for collecting networked responses that could be analyzed in real time. In 1996, one author wrote of the potential of using GSR measurements:

Empirical investigation of GSR revealed that there is a correlation between GSR scores and marketplace performance, that it is possible to pretest and rank alternative communications stimuli in terms of potential sales response before commercial production and that GSR scores can pinpoint insufficiently motivating communications stimuli. The study also demonstrated that GSR scores are better than consumer self-reported measures in predicting consumer marketplace behavior. Moreover, they can be used to accurately identify the more motivating and less motivating subelements. [8]

Still, there is an inherent problem of knowing what to look for, and when. Studies that collect massive amounts of values for multiple parameters that can then be analyzed and correlated off-line is a proven approach that is often used in long-term longitudinal studies (such as the Framingham Heart Study [2]). These studies, which operate on the principal that, given enough data, there will always be some correlation, probably are less suited to obtaining real-time feedback from small-scale events.

As a community, research on evaluating social signals for general-purpose networked application is in its infancy. We can all learn, however, from disciplines that have decades of experience in evaluating audience reaction – the world of television commercials. Much like the social interaction work performed within the multimedia and interactions communities, researchers in the field of advertising see great potential for measuring (and cashing in) on user feedback.

TV advertising research has long studied audience behavior in a wide range of stand-alone and embedded settings. In stand-alone setting, a commercial is presented to a focus group or to a monitored audience; data is collected across a well-understood set of parameters. In embedded settings, one or more commercials are inserted inside of a general content stream. The audience has no a priori knowledge of which item in the stream is of evaluation interest.

The Weasel Study

In this section, we will summarize a three-way comparison study conducted to evaluate audience reaction to a beer commercial. We summarize the report in [12].

Around 1970, the advertising agency Campbell Edward produced a television advertisement for the USA market for Heineken beer [12]. The storyboard is shown in Figure 8. In this ad, a young professional man strolls confidently into a party carrying a brown paper shopping bag. He exchanges casual glances with other partygoers as he enters.



Figure 8: Storyboard for *The Weasel* [12].

At one point, there is a more intense visual exchange with a striking female, who returns a flirtatious sign of interest. The man goes to the refrigerator, when he deposits a six-pack of Brand-X beer. Here he spots six bottles of Heineken beer – two of which he then takes out to the party (presumably to share with the woman with whom he exchanged glances). He first walks somewhat sheepishly away from the icebox, but then breaks into a confident stride. The commercial goes to black, then displays an *It's all about the beer* tag line, followed by the Heineken brand logo. The commercial lasts 30 seconds.

This advertisement was the subject of a study conducted under auspices of the Emotions in Advertising project of the American Association of Advertising Agencies and the Advertising Research Foundation [1]. In this study, three evaluation organizations conducted comparison research into the emotional engagement of audiences. One approach, conducted by Gallup & Robinson, used a measure of the contraction of facial muscles in subjects, a second approach, conducted by Ameritest, relied on comparative picture sorts via online interviews, and a third approach, conducted by Innerscope, used biometric monitoring via sensors embedded in subject clothing. An overview of the studies is given in Table 2.

The Weasel study provides an interesting comparison of three techniques to monitor audience engagement. A detailed summary of how engagement was experienced (and measured) is given in Figure 9.

	Facial EMG	Picture Sorts	Biometrics
Sample size	40 males; 21-35 years of age.	203 males; 21-35 years of age.	30 males; 25-35 years of age.
Mode of exposure	Embedded in a 15-minute video	Online clutter exposure	Embedded in 30-minute TV program
Method of study	FEMG; measurement of contraction of facial muscles	Visual picture sorts of key executional frames; respondents interviewed online	Biometric Monitoring System; sensors embedded in garment measuring respiration, motion, heart rate and skin conductance.
System (company providing it)	CERA (Gallup & Robinson)	PS (Ameritest)	BMS (Innerscope)

Table 2: Components used in the Weasel study [12].

DISCUSSION

There is something fascinating about the desire to predict audience reaction of an event. For performers (and speakers), gauging the reaction is often a critical component of fine-tuning a presentation. For commercial organizations (including advertising agencies), it is often a matter of maximizing return on investment or measuring impact. For all stakeholders, feedback can be used as a source of reflection or an agent of change.

Nearly all performers (academic and otherwise) have had an experience in which the reaction of an audience influences the pace, tone and depth of a presentation. In my own experience, I know that negative reactions (or, assumed negative reactions) are a much more powerful form of feedback than positive reactions. I naturally want to capture the mind (and heart) of the individual who is bored, dissatisfied or disengaged (often lost in his or her laptop).

Even if 98% of an audience is being swept along with the flow of a presentation, that 2% receives my attention.

Operating in a one-to-many personal performance setting is different than the mode in which my jazz band receives feedback. Here, the positive emotions of a dancing and active crowd can mask the (occasional) negative participant who is sitting quietly in the corner checking her e-mail.

For producers of everything from stage to film productions, predictive audience engagement (through the use of focus groups or the reliance of success-sequels) has proven to be more important than the feedback that can be provided by any particular audience on any particular day. Here the investment required before a production is audience-ready demands either a strong analytical justification or a finely-tuned producer’s ‘nose’ to motivate an investment decision. The analytical justification is often limited by the fact that audiences are good at reacting to things they know or imagine, but poor in reacting to content (or products) that they have never experienced.

The longitudinal approach to evaluating potential audience reactions based on a post-facto analysis of a wide range of measurement parameters has proved to be useful in detecting societal trends. Longitudinal studies help understand why smoking is bad, why eating eggs is unhealthy (and then to later justify why eating eggs is actually much healthier than assumed), and why carbon burned today may lead to climate change tomorrow. They

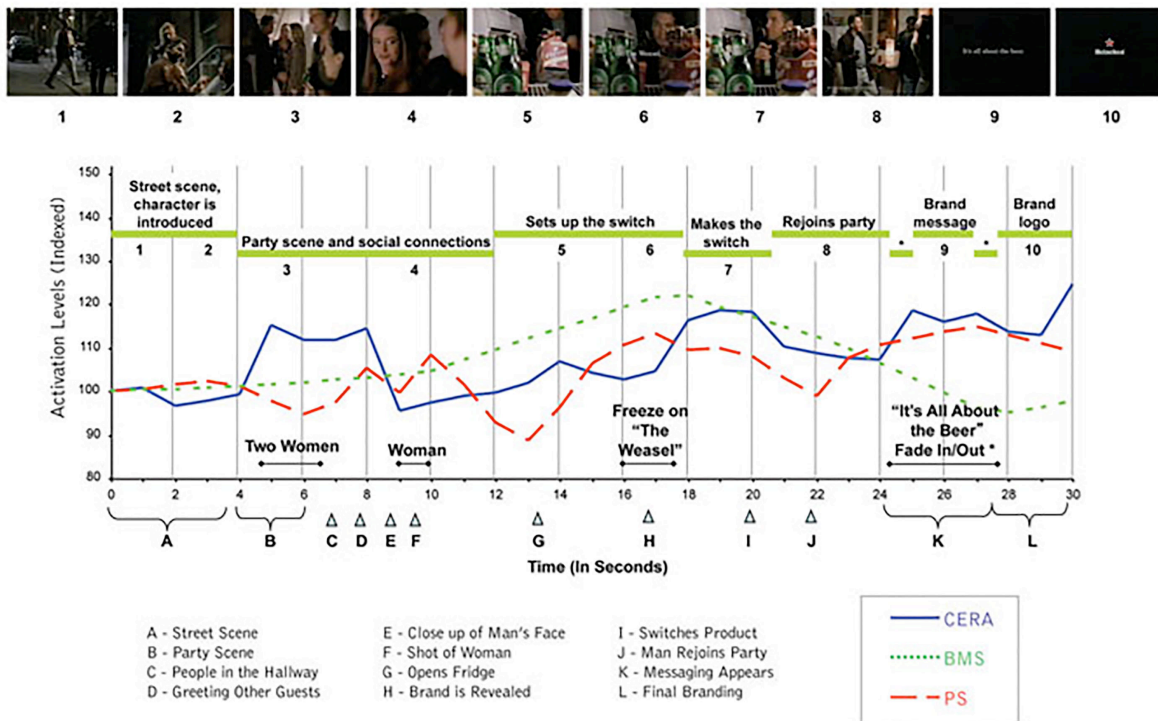


Figure 9: Evaluations based on three forms of emotion sensing [12].

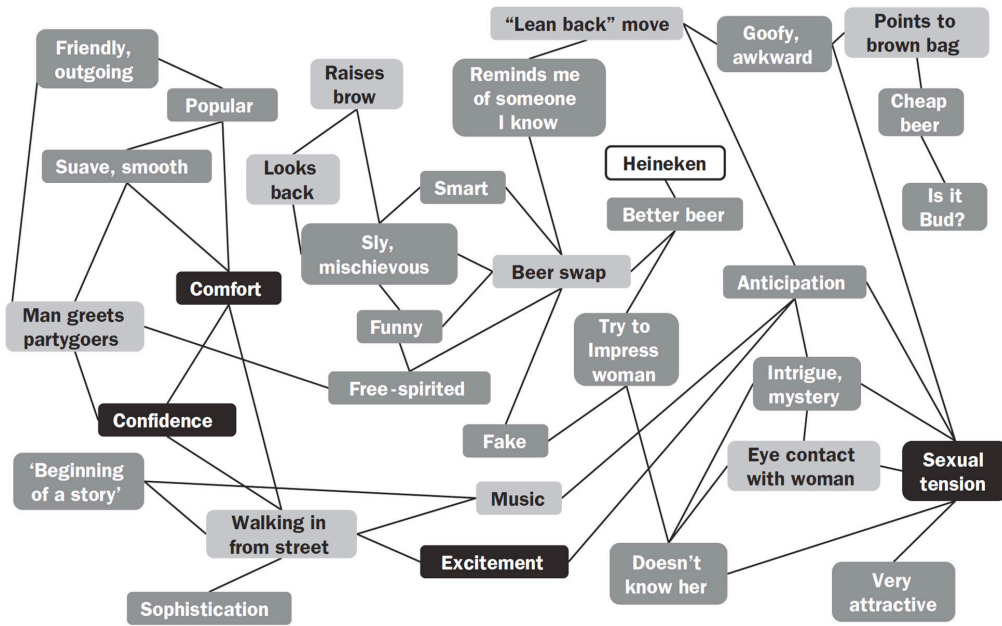


Figure 10: Concept maps for Weasel [10].

can also “prove”, however, that your chances of gaining weight increase significantly if a otherwise unknown friend of your friend’s friend gains weight easily [2].² In the same manner, simply wiring up the audience at a venue (as was described in the Introduction), without having any deep understanding of the audience members or the structure of the event, may lead to statistically correct but functionally absurd results.

Knowing whether someone likes jazz, if they were consuming an alcoholic beverage (or just had) or whether they were pre-occupied with problems that were orthogonal to the performance, are as essential to understanding the nature of feedback as recording their age and gender. More importantly, as illustrated by the *Weasel* analysis, knowing what you are look for during the presentation probably provides a more fruitful foundation for obtaining useful results that trying to overlay meaning on otherwise unstructured data.

One of the interesting aspects of the Weasel study (at least to me) was that there was little consistency and correlation between results based on biometric, anecdotal or visual analysis of an audience. The focus and structure (and the common language) used across all three approaches is particularly appealing. Often, however, even similar approaches to audience analysis remain locked in a battle of percentages rather than a battle of interpretation.

Still, even within the restricted domain of television advertising, with a known vocabulary of emotions and a

well-defined set of stimuli, there is tremendous room for deepening our understanding of audience behavior. In advertising, self-reported verbal reactions to ads remain the dominant method for obtaining audience feedback. A ‘concept map’ developed as part of an independent analysis of the Weasel commercial is given in Figure 10.

The second study on the Weasel also contained a GSR trace for audience emotional involvement, shown in Figure 11. Both the concept map and the GSR trace probably contain valuable information, although to the untrained eye, the main payback may be that GSR interest recovers (if only slightly) when the brand is shown on screen. (That this is a Heineken ad can hardly be a surprise, however, given the product’s prominent placement in the fridge and profiling in the content). I have no doubt, however, that a skilled marketing executive (or a data scientist) could obtain equally interesting explanations for the Weasel’s acceptance using the graphs in Fig. 1 as well.

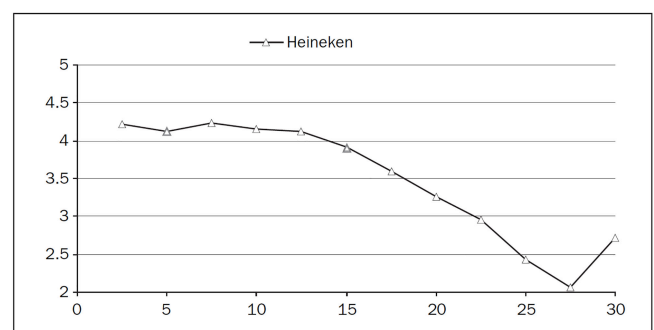


Figure 11: GSR trace for watchers of the Weasel commercial [10].

² It is unclear if Facebook friends exhibit the same properties.

It is difficult for me to articulate a ‘bottom line’ feeling for the value of the Weasel analysis, other than (1) to note that in beer marketing, the average summer temperature is historically the best predictor for beer sales, and (2) people watching a beer commercial are not in a store actually buying beer – here, any real-time association between emotion and action is difficult to define. Measuring audience emotional response to a commercial or to a concert does not necessarily explain audience purchasing behavior or help differentiate audience preference for an abstract genre rather than a genre instance such as a particular concert on a particular evening with a particular program.

CONCLUSION

The measuring of audience reaction to an event is interesting and important. Yet, it is not clear that naïve approaches yield results significantly beyond the production of impressive (if uninterrupted) graphs.

During my working life, I have served customers at a fast food restaurant that gave the illusion of personalized control over the edible content being supplied. My experience was that people loved to place adjective-rich orders, but equally empowered whatever the content of the food bag contained. I have also seen how the presentation of data was often more important than the data itself – who has time to look at all of that data?

In this article, I have summarized some informal experiences that I have had with understanding how others value content that they receive. We have used this information to define personalized presentations within the scope of concert summaries and person-focused movies. It has never been possible to determine if our users appreciated the particular content streams that we were able to present, or whether they simply appreciated the potential of having some personal influence in the content delivered.

I follow work on audience emotional evaluation with both interest and skepticism. Our community needs to determine a strict set of measures that can attempt to properly profile audience participants, properly profile the performances they engage with and properly characterize the multiple levels of influences that are in play on the production and consumption side of the emotion chain between audience and performer. I feel we have a long way to go.

ACKNOWLEDGMENTS

I would like to thank the workshop organizers for their encouragement and patience.

REFERENCES

1. AAAA/ARF, New Thoughts on Measuring Emotional Response to Advertising (2006).
http://www.mrcouncil.org/uploadedarchives/MRC%20-%20Jan_20_06%20Speaker%20presentation-recd%20-10-06.pdf.
2. Christakis, Nicholas A., and Fowler, James H. *Connected: The surprising power of our social networks and how they shape our lives*. 2011.
3. Daniel, Jessica A., *The Participatory potential of early cinema: A reexamination of early projected films*. Master’s Thesis, UIUC, 2010.
4. Fleming, George S., *What Happened on Twenty-third Street*, New York City.
(<http://www.imdb.com/title/tt0312017/>)
5. Gunning, Thomas. *An Aesthetic of Astonishment: Early Film and the (In)Credulous Spectator*,” in Linda Williams, ed. *Viewing Positions*, (New Brunswick: Rutgers, 1995), pp. 114-133.
6. Hess, Amanda. *Break Out the Chewbacca Mask and You Won’t Laugh Alone*, *New York Times*, 2016.
<http://nyti.ms/1WqVn0WC>.
7. Karasec, Hellmuth, *Ein Kurzfilm wirkte besonders nachhaltig, ja er erzeugte Furcht, Schrecken, sogar Panik*. <http://www.spiegel.de/spiegel/print/d-13687466.html>
8. LaBarbera, Priscilla A., and Joel D. Tucciarone. *GSR reconsidered: a behavior-based approach to evaluating and improving the sales potency of advertising*. *Journal of Advertising Research*. Sept.-Oct. 1995: 33+.
9. Lumière, Auguste and Lumière, Louis. *L’Arrivée d’un train en gare de La Ciotat*
(<http://www.imdb.com/title/tt0000012/>)
10. Micu, Anca C. and Plummer, Joseph T. *Measurable Emotions: How Television ads really work*. *J. Advertising Research* (50) 2, 2010.
11. NYT, *Edison’s Vitascope Cheered: ‘Projecting Kinetoscope’ Exhibited for the First Time at Koster & Bial’s*, *New York Times*, April 24, 1896, p. 5.
12. Patnaik, Sandeep and Purvis, Scott. *Comparative analysis of emotion measures: A case study of the Heineken “Weasel” commercial*. *Quirk’s Marketing Research Review* (www.quirks.com). March, 2011.
13. Radbourne, Jennifer, Johanson, Katya, Glow, Hilary, White, Tabitha. *The Audience Experience: Measuring Quality in the Performing Arts*. *Int. J. of Art Management*, (11) 3, 2009.