

Advertising Positions: Data Portraiture as Aesthetic Critique

Daniel C. Howe, Qianxun Chen and Zong Chen

ABSTRACT

Advertising Positions integrates 3D scanning, motion capture, novel image mapping algorithms and custom animation to create data portraits from the advertisements served by online trackers. Project volunteers use bespoke software to harvest the ads they receive over months of browsing. When enough ads have been collected, the volunteer is interviewed, 3D scanned and motion captured. Each ad is then mapped to a single polygon on the textured skin of their virtual avatar. Outcomes have been displayed as 2D/3D images, animations and interactive installations.

Advertising Positions synthesizes several areas of emergent technology: 3D scanning, motion capture and novel image/texture mapping algorithms to explore the problematic dynamics of online tracking by marketers, then renders this data contextually, as portraits and animations in 2D and 3D (see Fig. 1). For each user, our software collects the ads they are served over months of daily browsing. Once enough ads have been collected (several thousand in the examples discussed), the volunteer is 3D scanned and motion captured, then skinned with their unique set of advertisements, according to a mapping process that minimize the distortion of ads and maximizes correspondence with skin tones. The result includes still and moving images that synthesize technical, conceptual, and aesthetic elements in a cohesive fashion.

Background and Research

There was a brief time when the Web was a space largely free of commerce and capitalism. Today, however, each online action is tracked by sophisticated automated systems that mine our personal data to analyze, predict, and influence our behavior. While this vast surveillance architecture has been used by governments as a means of control [1], its unchecked growth has been largely due to corporations. Users are rarely asked to consent, but must instead accept such tracking as the necessary cost of access to the “free” information we can no longer live without.

The practice of repackaging consumers of media as products for advertisers is by no means unique to the Internet age—in fact it is as old as the newspaper ad itself. Ad technology has, however, evolved tremendously since then. Marketers, publishers and consumers are now dynamically connected by hundreds of intermediaries—tracking companies, data brokers, advertising networks—all in an effort to deliver billions of ads from thousands of marketers to millions of publishers [2]. In 2017, revenue from online advertising in the United States alone reached \$83 billion USD. Of that, the vast majority was spent on behavioral targeting; using an individual’s prior behavior (e.g. searches, pages previously visited, OS and browser information, as well as demographic data including age, location, gender, ethnicity, sexual orientation, interests, political affinities, etc.) to select ads to display [3].

Daniel C. Howe
School of Creative Media
City University Hong Kong
<daniel@rednoise.org>

Qianxun Chen
Brown University, U.S.A.
<qianxun_chen@brown.edu>

Zong Chen
School of Creative Media
City University Hong Kong
<dorothychen0303@gmail.com>

See <mitpressjournals.org/toc/leon/51/4> for supplemental files associated with this issue.




Fig. 1. Still image outputs created from three user profiles.
(© Daniel C. Howe)

To understand this system, we can consider a typical webpage visit. Code stored in various page locations creates bid requests that include the user's profile and the page being loaded. The request goes from the page to the advertising exchange where the page location and user's data are made available to advertisers on the network. Bids are instantaneously and automatically submitted for the spot, using rules determined in advance. The impression goes to the highest bidder and their ad is then displayed [4]. This process is repeated for each ad location on the page, with transactions occurring at lightning speed while the page is loading [5]. The criteria for bids can be very complex, taking into account detailed behavioral profiles and conversion probabilities.

All behavioral advertising has in common a dependency on the collection of user data for targeting. Unfortunately, once collected, such data may be used for a wide range of purposes. It can be traded and sold, aggregated across networks, and even used for questionable practices like discriminatory pricing, in which different users are given different prices for an item [6]. Because the data collected is so personal, the set of ads served to a user represents a portrait of what has been called her "data double" [7]. As such, there is significant potential for harm, a fact that has long been recognized by the privacy community. It is only recently however that such risks have been addressed in mainstream media, and thus become more widely appreciated by average users. And while some user-level tools have been created to address these risks (Privacy Badger, uBlock, Lightbeam, etc.), tracking has rarely been the focus of visually oriented artistic practice. *Advertising Positions* represents a clear counter example here; one that also raises important questions regarding socially engaged art that leverages strong aesthetic elements (questions that we return to below).

Technical

Prototyping and Collection

Targeted ad content served to a particular online user is determined by a number of factors: page, time of access, user location and, as mentioned above, a wide range of demographic data. To explore the relationship between these factors and the ads served, we initially created a prototype system using software bots. As each crawled the Web, using a modified version of Google's Chrome browser, visiting pages and clicking links, their advertisements were stored. So that these bots might collect ads similar to those a human user would see, we gave each bot a fictional profile; e.g. Wilson: a 39-year-old single male in Hong Kong, with interests in travel, tech, cooking, dating and adult entertainment. For efficient crawling, we used Selenium, an automated browser framework, in which the AdNauseam extension [8] was installed.

To ensure that the ads served to each bot corresponded to its profile, a Gmail account was created for each. In the ad settings for the account, we input a range of information, including gender, age and interests. Before each browsing session, bots would begin by logging into their Gmail account before starting to search keywords and then selecting links to follow from the results. If the process stalled—for example, on a page without links—the bot could look up further sites related to its interests via Web indices listing pages by genre.

3D Modeling

Unlike the 2D images shown in Fig. 2, in later phases of the project we used 3D models for animation and rendering. Constructing a base model from human users involved profile studies and in situ 3D scanning via white-light technology, resulting in models composed of ~200,000 triangles (Fig. 3). Although these models are highly accurate, they were difficult to animate due to the high polygon count, the triangle-based topology and inconsistencies between their structure and those of the human face. To address these issues, we recomposed the models into low-poly count equivalents, with topology aligned to facial musculature. Each vertex in the new model is drawn along the surface of the scanned model and strategically positioned to maximize the difference between the slope values it separates, maintaining a delicate balance between highlighting important detail, following muscle patterns, minimizing polygons




Fig. 2. Ad collage of the user profile—Wilson. (© Daniel C. Howe)

and equally distributing vertices. This process yielded animation-ready models with fewer than 5,000 quads (Fig. 3), enabling more natural animation and increased rendering efficiency.

Ad Texture Mapping

To layout advertisements on the model's skin, we mapped each advertisement to a single quad on the 3D surface. Our mapping algorithm involved initially extracting vertex positions for each quad, then crafting a cost-minimization function for computing optimal ad/quad pairs. To map advertisements onto the 3D surface, a 2D projection of the 3D grid was created with each pixel translated from 2D to 3D space. For optimal results the algorithm needed to consider the distortion that occurs when a rectangular image is mapped to an arbitrary 3D quad, after aligning to the viewer's natural perception. To do so, we ordered the quad vertices in perspective space, then extracted the 3D positions of each vertex and paired it with 2D positions on the UV map. This allowed us to generate a UV-based 2D image textured with the advertisements, which could then be translated directly onto the 3D model.

Once we can map correctly ordered quad vertices, we must design a function to produce optimal pairing of images and quads. Here we note that the problem reduces to a classical CS problem called the Stable Marriage Problem [9], having a polynomial solution. Thus we need only define a cost




Fig. 3. Shaded (top) and wireframe (bottom) view of scanned model (left) and recomposed model (right). (© Daniel C. Howe)

function to input into a standard solver. Pairings that minimize cost minimize a) the distortion of each positioned image and b) the difference between image brightness and quad brightness from the UV map. The following cost function $f(q,i)$ computes the cost of each quad-image pair:

$$f(q,i) = (A * distWarp(q,i) + B * maxWarp(q,i)) + |i.brightness - q.brightness| < C ? 0 : \infty$$

Here the two image distortion metrics, controlled by tunable constants A and B, can be weighted to minimize warp, and pairings where the absolute difference is above threshold C are given infinite cost and excluded. Figure 4 shows an example output from this algorithm.

Animation

The animation process involves taking reference images from the subjects' motion capture sequences, decomposing them into movements of individual body parts (this increases facial flexibility and helps make more natural expression transitions), reconstructing them on the model, and exaggerating them for artistic purposes. The resulting animations explore subjects' initial reactions to their aggregate ad sets, with responses that included skepticism, surprise, pleasure, distress and others.

Outcomes

For each user profile, three outcomes have been generated thus far: animations of the 3D models, rendered collages of aggregate ad sets, and high-resolution still images. The project debuted in 2017 at the Oi! Street gallery in Hong Kong (Fig. 5), and was later featured at Art Center Nabi in Seoul, South Korea. In current iterations of the project, we import models into Unity3D to enable programmatic control and simple interaction (models gazing at audience members, etc.).

Aesthetic Considerations

The focus of our critique lies not in personalized marketing, but in the opaque system through which consumers are targeted without their consent. To address this issue, *Advertising Positions* presents the human body as a canvas on which to visualize aspects of the human experience that we are forced to sell each day (our attention, our privacy, culture itself). Unlike forms of manipulated recording (for instance, photography or filming), 3D modeling is a medium that not only depicts, but embeds advertisements—




Fig. 4. The 2D UV map with quads only (left) and with matching advertisements (right). (© Daniel C. Howe)



Fig. 5. Installation shot at Oil gallery in Hong Kong. (© Daniel C. Howe)

the tangible evidence of our profiles being sold—within the aesthetic object itself. Our subjects here are deliberately objectified, positioned in a frame, staged under a spotlight, in an attempt to expose the nature of the system that places us, without our consent or knowledge, on the virtual bidding table. The presentation of subjects is deliberately minimal—clothing removed, torso removed, representations of distinct physical identity removed—to reveal the emotional vulnerability residing in our new identities as online targets.

In the animated sequences, subjects are seen to undergo combinations of surprise, frustration and distress, emotions that reflect research on the dilemma consumers experience in relation to ad targeting. When told they must trade privacy for the somehow intangible benefits of customized ads, most experience a sense of violation, yet also feel unable to act upon the situation in any viable manner [10]. *Advertising Positions* attempts to explore exactly this tension via aesthetic means, highlighting the powerless position many consumers experience.

With its compelling visual elements, the project also raises important questions for art that critically addresses problematic social dynamics. How does the aestheticization of destructive practices affect users' understanding and emotional response to them? As a canonical example, we might consider Edward Burtynsky's hauntingly beautiful aerial photos [11] of environmental destruction or Trevor Paglen's often sublime long-distance photos [12] of secret military bases. Although we cannot speak for other artists, our goal here is to create a visceral tension between the aesthetic appeal of the work, on the one hand, and the disturbing social dynamics from which it emerges, on the other. Thus it is crucially important that users experience not only the more directly engaging visual aspects of the work, but also perceive its conceptual and theoretical foundations. This is the work of narrative—how best to make the work's process and material substrate an organic element of the viewing experience, rather than an afterthought to be found only on a caption card. Thus we also aim, through interviews with gallery visitors following visits, to gain insight into the crucial mechanics of user reception, and the ways in which aesthetic, conceptual and technical elements of the project integrate within their experience.

Conclusion and Future Directions

Over the course of development, we have identified several improvements and new directions for *Advertising Positions*. We are currently exploring real-time programmatic control of the models, to enable subtle interactions with users and variations in behavior over longer durations. We are also experimenting with new motion capture and 3D-scanning techniques that better map to our process. Perhaps most interesting of all will be the post-interviews from visitors (now being transcribed and coded), as to their subjective experience negotiating aesthetic and narrative elements of the project.

Acknowledgments

Special thanks to Leoson Cheon, Miu Ling Lam, Diane To, Tobias Klein, Hector Rodriguez and Tamas Waliczky for their generous feedback and assistance.

References and Notes

1. A. Soltani, A. Peterson and B. Gellman, “NSA Uses Google Cookies to Pinpoint Targets for Hacking,” *The Washington Post*, 10 December 2013.
2. LUMA Partners, “Display LUMAScape”: <www.lumapartners.com/lumascapes/display-ad-tech-lumascapes/> (accessed January 2018).
3. L. Johnson, “U.S. Digital Advertising Will Make \$83 Billion This Year, Says EMarketer”: <www.adweek.com/digital/u-s-digital-advertising-will-make-83-billion-this-year-says-emarketer> (accessed January 2018).
4. S. Yuan, J. Wang and X. Zhao, “Real-time Bidding for Online Advertising: Measurement and Analysis,” in *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising* (New York: ACM, 2013) p. 3.
5. S. Horowitz, “What Is Real Time Bidding and Why Is It More Effective than Direct Bidding Methods?": <<http://executive-digital.com/blog/real-time-bidding-vs-direct-bidding/>> (accessed January 2018).
6. M.R. Wigan and R. Clarke, “Big Data’s Big Unintended Consequences,” *Computer* **46**, No. 6, 46–53 (2013).
7. K.D. Haggerty and R.V. Ericson, “The Surveillant Assemblage,” *British Journal of Sociology* **51** No. 4, 605–622 (2000).
8. D.C. Howe and H. Nissenbaum, “Engineering Privacy and Protest: A Case Study of AdNauseam,” in *Proceedings of the IEEE International Workshop on Privacy Engineering* (San Jose, CA, 2017).
9. D.G. McVitie and L.B. Wilson. “Stable Marriage Assignment for Unequal Sets,” *BIT Numerical Mathematics* **10**, No. 3, 295–309 (1970).
10. J. Turov, M. Hennessy and N. Draper, “The Trade-off Fallacy: How Marketers Are Misrepresenting American Consumers and Opening Them Up to Exploitation”: <www.ftc.gov/system/files/documents/public_comments/2015/09/00012-97594.pdf> (accessed March 2018).
11. M. Schubert, *Edward Burtynsky: Oil* (Göttingen: Steidl, 2016).
12. J. Bryan-Wilson et al., *Trevor Paglen* (London and New York: Phaidon, 2018).