

Ad Empathy: A Design Fiction

Michael Skirpan

University of Colorado Boulder
Boulder, CO USA
Michael.Skirpan@colorado.edu

Casey Fiesler

University of Colorado Boulder
Boulder, CO USA
Casey.Fiesler@colorado.edu

ABSTRACT

Industry demand for novel forms of personalization and audience targeting paired with research trends in affective computing and emotion detection puts us on a clear path toward emotion-sensitive technologies. Written as API documentation for an AI marketing solution that provides “emotion-sensitive marketing decisions,” this design fiction presents one possible future application of today’s research. Offering a demonstrable grey area in technology ethics, Ad Empathy should help to ground debates around fair use of data, and the boundaries of ethical design.

Author Keywords

advertising; API; design fiction; emotion; ethics; social computing; speculative fiction; neural networks; machine learning; target marketing

ACM Classification Keywords

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

PRODUCT INTRODUCTION

Today’s competitive attention economy requires brands to reach customers in personal and affective ways. We have long known from research that personalization is both effective and an important factor in determining consumer attitudes towards advertising [20]. However, personalization is also now a saturated approach as it has become easier for online advertisers to know *what* a customer wants. Companies wanting the competitive edge now need to know *when* a product is best advertised and *how* it should be framed. Knowing this demand, we are happy to launch Ad Empathy, an AI marketing solution helping brands make emotion-sensitive marketing decisions.

Our API Resources are designed to help our clients generate content ad impressions that cater to the dynamic needs of their audience. We work with most major social

Please do not modify this text block until you receive explicit instructions.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CONF '22, Jan 1 - Dec 31 2022, Authorberg.

Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-xxxx-yyyy-z/zz/zz...\$zz.00.

unique doi string will go here

media platforms and search engines to create connected profiles of customers that can be accessed from any ad client via the Ad Empathy API. For each advertising platform you would like to integrate with Ad Empathy, simply add your company’s registered tokens using the Ad Empathy Dashboard and within 48 hours we will have trained models for each of your customers. From that point onward, you can use the Ad Empathy API to design your ad impressions on any connected platform. To use Ad Empathy as a full-cycle marketing platform, you may also register your product inventory with our platform to track emotional responses to product-specific brand interactions and improve our models.

GETTING STARTED

Before making any requests using our models, you should contact a member of our Sales Team to discuss pricing options or obtain a free trial. All API Resource requests must contain a valid token pair <client-token> and <client-secret>, a <cookie-id> for the user, and optionally a <platform-id> to specify the ad client platform. Developers building platform-agnostic services can use our Accounts API to obtain valid <cookie-id>’s for building cross-platform ad campaigns and event triggers.

API RESOURCES

Once you have obtained valid token pairs, integrated your external ad platform’s tokens, and see the green check mark at the top corner of your Ad Sense Dashboard, you can begin using any Ad Empathy API Resource

Mood

Mood

get - *GET* /mood/now/<cookie-id>

Returns current emotional state (mood) of user as a list of top ten moods by confidence

list - *GET* /mood/list/<cookie-id>

Returns a list of frequencies for all moods categories that Ad Empathy has related to the specified user.

Mood.product

list - *GET* /mood/product/<product-id>/<cookie-id>

Returns a list of product IDs and the mood that is most positively associated with a customer interaction.

Mood.topic

list - *GET* /mood/topic/<cookie-id>

Returns a list of content topics and our highest confidence mood association for that topic.

Trend

Trend

get - GET /trend/now/<cookie-id>

Returns the predicted emotional states, ordered by confidence, for upcoming 30-minute time interval.

list - GET /trend/daily/<cookie-id>

Returns a list of 30-minute time intervals over 24-hours with the most common emotional state associated to each interval.

Response

(API Resource available only to customers using Ad Empathy Trackers for their product inventory)

Response

get - GET /response/<product-id>/<cookie-id>

Returns the user's last cached online emotional response to an interaction with <product-id>.

Expression

Expression.text

get - GET /expression/single/<emotion>/<cookie-id>

Returns the syntax tokens most commonly associated with the user's online expression of the emotion.

list - GET /expression/all/<cookie-id>

Returns a paginated list of emotional states, sorted by their frequency, and the most common syntax tokens associated to that state.

HOW DOES IT WORK?

Ad Empathy is a state-of-the-art multi-model AI ecosystem that leverages the volume and velocity of online behavioral data by training user-specific machine learning models. The core of the system is a Long Short Term Memory (LSTM) neural network trained specifically to predict the evolution of moods using temporally-structured data coming from online activities (eg., text from posts, click content, reactions to others' posts). Our company began training this model nearly five years ago when researchers first found Gated Recurrent Units as a solution to cutting through the noise of online data [15]. After years of fine-tuning and learning how to transfer models between different users and incorporate multi-modal data, we found we had sown the seeds of something much bigger than a mood prediction model. In short, this core model became the heart of a system of interacting models. Developing our expertise in model transfer allowed our team to take layers of our novel LSTM model and combine them with convolutional layers, other Recurrent Neural Network (RNN) language-processing layers, and add them into Generative

Adversarial Networks to blossom the wide functionality you see today.

When your company opens an account with Ad Empathy, our system begins by data mining all social media content and tracked brand interactions available for your customer base. After mining all historical data about your customers, we place their user accounts into our reactive event loop that keeps tabs on new activities across any connected platform. Prior to training, we run all the data through a noise reduction network trained specifically to identify relevant emotional content. Using the filtered data set, we fork fresh versions of our base model and begin training a unique mood model for each of your customers. This training continues until the confidence of our predictions meets a certain threshold. Testing is done using a data set we capture and separate during the data-mining phase. Our central model (the one underneath the Mood API) takes in time-structured online activity for a user and outputs a likely current mood given the most recent observation. This model is then transferred into our second network, which chunks your users' history into 24-hour segments and trains a model that predicts the upcoming 24-hour emotional cycle (and provides the backbone of our Trends API!).

Once we have accurate models for our Moods and Trends API, we begin fine-grain analysis on specific data such as text and photos. This process starts by performing a topic-modeling analysis on all user text and browsing history to break up each user's history into topic-specific data sets. Further, each user photo is analyzed for facial expression, object detection, and captioning to develop visual insights into the personal aesthetics of your customer's emotions. A core value that Ad Empathy offers is recognizing that each product a customer purchases is embedded in a different context and thus requires a different cognitive model to understand underlying emotional relationships. We develop those models along many dimensions that account for complex relationships between emotions and brand sentiments.

Important to understanding how Ad Empathy works is that each API your team uses is operating with different custom models and parsing techniques that branch out from our central mood-recognition network. Our Expression API, for instance, uses sentiment analysis in tandem with a generative adversarial network to parse user text and then learn how to generate novel text that expresses the same sentiments while staying within the known vernacular of your customer. The adversarial network is trained against the core mood model, which allows rapid exploration of the syntax space observed and parsed from your customers' online platforms.

If your company would like to learn even more about the inner-workings of Ad Empathy, feel free to make an appointment with our Machine Intelligence Team to discuss specifics or let us know how you think we could improve our process.

EXAMPLE USE

Working with customers, we have found solutions that mix and match our APIs to help you generate the relevant content and design marketing campaigns most appropriate to your products. We explain some of our most successful applications below:

Time Cycling

Our research has shown that many customers have predictable emotional response patterns based on time of day. It is often reliable that a customer will elicit more positive emotions to food around 11AM; however, this response will diminish leading up to around 2PM as it becomes more likely they already ate lunch. For this reason we recommend *time cycling* campaigns for products with emotions that are highly correlated to temporal patterns.

For this, we recommend analysis of your products with our *Trends* Resource to discover your most temporally stable products and to make inferences about how they are associated across time. Then using our *Expressions* Resource, you can design context-sensitive Content Ads that can portray your product regularly at the times associated to the emotion best suited for your product.

A/B Emotional Testing

Not sure whether your product is better fit to when your customer feels happy or angry? Try A/B Testing emotions instead of features. Combining our Impressions and Response APIs, your team can try your ad impressions against different emotional conditions to see what elicits the most positive response. This can improve how you understand how your product is being perceived and better inform our models.

For well-modeled user profiles, your team may try running simulations using our Impressions and Expressions APIs. You can pilot your A/B tests, discovering correlations between ad impression and emotional responses and designing ad impressions with the right emotional language.

APPROPRIATE USE OF AD EMPATHY

The purpose of Ad Empathy is to support businesses in employing emotional insights as they create online advertisements. We love seeing our customers doing rapid prototyping of new ad campaigns and trying out new combinations of our models to maximize the utility emotions and timing play in your ad impressions. Ad Empathy, however, is *not* meant to be used as a research platform nor should it be used to target specific customers and invade their privacy. We do not approve of customer-specific analysis that exposes potentially sensitive vulnerabilities related to private dimensions of a customer's mental state.

Ad Empathy should also never be used in relation to medical data or to support mental health inference relative to emotional trends. Similarly, our insights should remain in the realm of marketing and should not be used in decision-making algorithms related to employment, education, housing, or health. Though we are proud of the accuracy of

our system, it is not appropriate to use such predictions to make firm decisions that could negatively impact your customers. If your company is focused on biomedical or employment-related inference, please contact our Customer Relations Team to discuss fair uses of data and how to access our models for purposes outside of our available products. Projects that are funded by a government agency should speak to an Ad Empathy representative before using our products. If your use of Ad Empathy goes beyond marketing, we offer consulting services to help your company develop an ethical and accurate system that incorporates emotional insights.

Thank you again for using Ad Empathy!

AUTHOR'S STATEMENT

The goal of this design fiction is to structure discussion around a technology that is at the cusp of creation, regardless if it emerges in this exact form. Industry demand for novel forms of personalization and audience targeting paired with research trends in affective computing and emotion detection puts us on a clear path toward emotion-sensitive technologies. With both the capability and economic incentives in place, we must, as a community, carefully define lines between what we consider fair marketing applications of technology versus unwelcome and unfair intervention or even exploitation.

Design fiction is one way to consider these possible futures. As a conflation of design, science fact, and science fiction, the medium is a method for exploring ideas, implementation, and consequences [6]. Importantly, as Baumer points out in an introduction to a set of fictional conference abstracts, these visions of tomorrow can help shape the research directions of today [3]. Lindley further proposes design fiction as a methodology for considering the ethics of radical digital interventions [12]. Therefore, we ask: how could a vision of tomorrow inform the ethical considerations of the research we are conducting today?

Written as an API, the piece situates itself both in technical and social literatures of computing. Questions have already been raised about the ethics of corporate experimentation and the fine line between product testing and harmful intervention [13]. Research has shown that users may not really understand what they are consenting to when agreeing to a terms of service [2,10]. They may also find certain uses of their data to be "creepy" or invasive when it comes to behavioral advertising [19]. When asked about the process of data merging and aggregation, users tend to feel they are not the ones receiving a true benefit [5].

Though these user attitudes may raise red flags, research and industry continue expanding our capabilities in this area. In computer vision, deep neural nets have been a boon for new models that aid in extracting emotion from facial images posted online [4,11]. Text is no different as research continues to improve our ability extract emotional insights from syntax tokens [1,14]. Separately, researchers have proven capabilities to make mental health inferences using

social media data [7,8]. Typically, future directions for this kind of work involve technology design for helping people. However, there are other potential uses for this technology, including online marketing tactics.

If we consider the bleeding edge of marketing and artificial intelligent, we see very similar forms of emotional targeting being brandished as the next wave [16]. Yet, when users actually find out that they are being classified on psychological and emotional terms, it foments anger and is seen as “overstepping boundaries” [17]. In academic circles, researchers such as Zeynep Tufekci and Kate Crawford have begun structuring a debate around new kinds of privacy harms caused by advancements in AI and algorithmic methods [9,18]. Their concern is based on the fact that predictive inference is now able to go beyond what users openly disclose about themselves.

Ad Empathy and its API Resource offer a demonstrable grey area in technology ethics. The product very clearly meets the path we are trending toward, yet it should provoke some sense of caution or discomfort in its ability to find users at their most vulnerable moments. Without a doubt, this kind of system will become possible and machines will continue pushing the limits of our cognitive capacity to recognize manipulation, presenting ethical issues that are worthy of close consideration and skepticism. As a discussion piece, the Ad Empathy fiction should work to ground debates around fair use of data, and the boundaries of ethical design.

REFERENCES

1. Ameeta Agrawal and Aijun An. 2012. Unsupervised emotion detection from text using semantic and syntactic relations. In *Proceedings of the IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology*.
2. Yannis Bakos, Florencia Marotta-wurlger, and David R Trossen. 2009. Does Anyone Read the Fine Print? Testing a Law and Economics Approach to Standard Form Contracts. *New York University Law and Economics Working Papers Paper* 195. Retrieved from http://lsr.nellco.org/nyu_lewp/195
3. Eric P S Baumer, et al. 2014. CHI 2039: Speculative Research Visions. ACM CHI 2014.
4. C. Fabian Benitez-Quiroz, Ramprakash Srinivasan, and Aleix M. Martinez. 2016. EmotioNet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In *Proceedings of the IEEE Conference on Computer vision and Pattern Recognition (CVPR)*.
5. Igor Bilogrevic and Martin Ortlieb. 2016. "If You Put All The Pieces Together..." Attitudes Towards Data Combination and Sharing Across Services and Companies. ACM CHI 2016.
6. Julian Bleecker. 2009. Design Fiction: A Short Essay on Design, Science, Fact and Fiction. *Near Future Laboratory*: 4–97.
7. Stevie Chancellor, Zhiyuan Lin, Eric L. Goodman, Stephanie Zerwas, and Munmun De Choudhury. 2016. Quantifying and predicting mental illness severity in online pro-eating disorder communities. ACM CSCW 2016.
8. Munmun De Choudhury, Scott Counts, and Eric Horvitz. 2013. Predicting postpartum changes in emotion and behavior via social media. ACM CHI 2013.
9. Kate Crawford and Jason Schultz. 2013. Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms. *Boston College Law Review* 55, 1.
10. Casey Fiesler, Cliff Lampe, and Amy S. Bruckman. 2016. Reality and Perception of Copyright Terms of Service for Online Content Creation. ACM CSCW 2016.
11. Youngsung Kim, ByungIn Yoo, Youngjun Kwak, Changkyu Choi, and Junmo Kim. 2017. Deep generative-contrastive networks for facial recognition. *arXiv (working paper)*. Retrieved from <https://arxiv.org/abs/1703.07140>
12. Joseph Lindley. 2015. Operationalising Design Fiction for Ethical Computing. *ACM SIGCAS Computers and Society* 45, 3: 79–83.
13. Michelle N. Meyer. 2015. Two Cheers for Corporate Experimentation: The A/B Illusion and the Virtues of Data-Driven Innovation. *Colorado Technology Law Journal* 13: 273–332.
14. Myriam Munezero, Calkin Suero Montero, Maxim Mozgovoy, and Erkki Sutinen. 2013. Exploiting sentiment analysis to track emotions in students’ learning diaries. In *Proceedings of the Koli Calling International Conference on Computing Education*.
15. Rajib Rana. 2016. Gated Recurrent Unit (GRU) for Emotion Classification from Noisy Speech. *arXiv (working paper)*. Retrieved from <https://arxiv.org/abs/1612.07778>
16. Gargi Sharma. 2017. How emotion detection technology can make marketing more effective. *ParallelDots*. Retrieved from <http://blog.paralldots.com/technology/changing-marketing-with-emotion-detection-technology/>
17. Olivia Solon. 2017. “This oversteps a boundary”: Teenagers perturbed by Facebook surveillance. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2017/may/02/facebook-surveillance-tech-ethics>

18. Zeynep Tufekci. 2015. Algorithmic Harms Beyond Facebook and Google: Emergent Challenges of Computational Agency. *Colorado Technology Law Journal* 13: 203–218.
19. Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Symposium on Usable Privacy and Security (SOUPS)*.
20. David Jingjun Xu. 2006. The influence of personalization in affecting consumer attitudes towards mobile advertising in China. *Journal of Computer Information Systems* 47, 2: 9–19.