

Affective Signals in a Social Media Recommender System

Jane Dwivedi-Yu
Meta AI
London, UK
janeyu@fb.com

Yi-Chia Wang
Meta AI
Menlo Park, CA, USA
yichiauw@fb.com

Lijing Qin
Meta AI
Seattle, WA, USA
qinlijing@fb.com

Cristian Canton-Ferrer
Meta AI
Seattle, WA, USA
ccanton@fb.com

Alon Y. Halevy
Meta AI
Menlo Park, CA, USA
ayh@fb.com

ABSTRACT

People come to social media to satisfy a variety of needs, such as being informed, entertained and inspired, or connected to their friends and community. Hence, to design a ranking function that gives useful and personalized post recommendations, it would be helpful to be able to predict the affective response a user may have to a post (e.g., entertained, informed, angered). This paper describes the challenges and solutions we developed to apply Affective Computing to social media recommendation systems.

We address several types of challenges. First, we devise a taxonomy of affects that was small (for practical purposes) yet covers the important nuances needed for the application. Second, to collect training data for our models, we balance between signals that are already available to us (namely, different types of user engagement) and data we collected through a carefully crafted human annotation effort on 800k posts. We demonstrate that affective response information learned from this dataset improves a module in the recommendation system by more than 8%. Online experimentation also demonstrates statistically significant decreases in surfaced violating content and increases in surfaced content that users find valuable.

CCS CONCEPTS

• **Human-centered computing** → **Social networks**; • **Information systems** → **Sentiment analysis**; **Recommender systems**.

KEYWORDS

affective computing, recommendation systems, social media

ACM Reference Format:

Jane Dwivedi-Yu, Yi-Chia Wang, Lijing Qin, Cristian Canton-Ferrer, and Alon Y. Halevy. 2022. Affective Signals in a Social Media Recommender System. In *Proceedings of the 28th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '22)*, August 14–18, 2022, Washington, DC. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3485447.3512131>

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
KDD '22, August 14–18, 2022, Washington, DC
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-XXXX-X/18/06.
<https://doi.org/10.1145/3485447.3512131>

1 INTRODUCTION

Social media platforms have become a common means of interacting and connecting with others as well as finding interesting, informing, and entertaining content [5, 20, 51]. Users of those platforms depend on the ranking systems of the recommendation systems to show them information they will be most interested in, provide them with positive experiences, and safeguard them against offensive material. Users may also look for online content that will help them change or enhance their current affective state [31, 52], and social media is indeed rich in these signals [22, 23, 48]. Conceivably, Affective Computing, a field that develops methods to predict user's affects in a particular application context (e.g., customer support interaction, online learning), can contribute to the suite of methods used by recommender systems to ensure that users are having the best experiences possible [9, 12, 40, 50, 60]. This paper describes the challenges and solutions we developed to apply Affective Computing in the context of Facebook's ranking algorithm.

The first contribution of this paper is translating the vision of applying Affective Computing into a well-specified technical problem. Doing so involved several challenges. The first challenge was to understand where in the complex recommendation system an affective prediction can be useful for ranking and how to combine it with other ranking signals. The second challenge was to devise a set of criteria for determining which affects pertain to recommender systems (e.g., inspiration, entertainment, sadness, fear) and to apply these criteria to decide on a reasonably short list of affects to operationalize.

The second set of contributions concern the operationalization of Affective Computing where we developed models for predicting the potential affective response a user may have to a post. We began by developing methods for training affect classifiers based on users' engagement (reactions, shares, outbound clicks, negative user feedback, and comments) on the platform, which carries an important signal about their affective response towards content. However, engagement alone does not suffice, because a single response can sometimes indicate different intents given different contexts. For instance, consider the sad or sorry reaction used in response to a post. In some cases, the reaction may simply indicate the user's emotional support, but in other cases, it may indicate the user's sadness and a preference to never see that type of content again. Hence, we also supplemented the signal based on engagement with a carefully constructed human annotation pipeline to directly label affective response. Using this data, we developed a model for

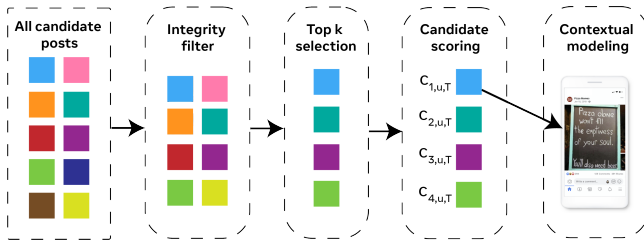


Figure 1: Example of a recommender system. The first step removes any violating content from the pool of candidate items. The top 500 or so of the remaining posts get scored on whether post p would produce value for the user u at time T , denoted $c_{p,u,T}$. Lastly, the selected posts undergo a re-ranking step to ensure the feed contains enough diversity.

predicting affective signals that was then injected into the ranking algorithm of the recommendation system, leading to substantial improvements of several metrics.

The paper is organized as follows. Section 2 describes where Affective Computing can contribute to a recommendation system and how we decided which affects to operationalize. Section 3 describes how we create training data for our classifiers and the modeling architecture and features for predicting affective signals. Section 4 provides an analysis of both the dataset and model when incorporated into the recommendation system, describing experiments that validate some of the choices we made in this work. Section 5 describes related work and Section 6 concludes and points to future avenues of research.

2 PROBLEM DEFINITION

This section begins by reviewing the basic aspects of recommendation systems that are relevant to our discussion (Section 2.1). We then define the concept of affective response and discuss how it can be incorporated into the recommendation system (Section 2.2). Finally, we discuss how we chose the affects to model in our work (Section 2.3).

2.1 Recommendation systems

The recommendation system is the component responsible for deciding which posts will be shown to a user and in what order. A typical architecture of a recommendation system inspired by [24] is shown in Figure 1.

The recommendation system starts with a pool of available posts. These posts may have been posted by the user’s friends or pages they follow, or come from sources that post about topics that the user seems to be interested in. The posts can also be curated by the platform, such as information about voting locations during an election cycle or about COVID vaccines.

Because much of the content is organic, the first step in the process is to remove items that violate the community policies set forth by the platform. Such items may include hate speech, bullying, nudity, calls for violence, or misinformation [13]. At the next step, the recommender uses a lightweight model to reduce the number of candidate items to about 500 per user.

In the main ranking step, the goal of the recommender system is to produce a score $c_{p,u,T}$ which determines whether post p is of value to user u at time T . The value of $c_{p,u,T}$ is computed as a weighted linear sum of a set of prediction models. Prediction models can include trying to predict engagement actions such as liking or resharing a post. They can also incorporate how recent the post is or whether other friends have engaged with it. Since not all types of user value can be captured through their engagement with posts, the platform also conducts surveys with small sets of users. These surveys can ask direct questions such as whether posts are worth their time, provide value, or contribute to their feeling of being informed. Based on these surveys, an additional set of predictors tries to predict the user’s answer to these survey questions.

The final step of the recommendation system is a re-ranking step, where the goal is to ensure that the feed that is produced contains enough diversity (e.g., not overwhelmed by a particular topic or set of users that post often).

2.2 Affective response

In this work, we use the term **affect** to refer to both an affective state (e.g., anger, joy) as well as a cognitive state (e.g., entertainment, inspiration). Our goal is to model the **affective response**, the affect that a user may have when exposed to a particular post, and use such a model to improve a scoring module of the recommendation system.

Predicting affective response is a form of content analysis. A common type of content analysis done for recommendation today is to determine the *topic* of a post [2, 37]. However, the topic alone does not capture its full nuance. For example, consider that a post on the topic of clowns can have an affective response varying from happy to scared. Our definition of affect as both affective and cognitive experiences is motivated by a need to narrow the gap between topic analysis and more nuanced analysis of a post.

We found that a prediction of affective signals can be useful in several of the scoring modules that are used in the recommender system, and therefore we create an embedding that can be used as a feature in any scoring module. The embedding is the penultimate layer of the network that predicts affective tasks.

An alternative option would have been to use the affective model as a scoring module in itself. We discarded this option because these affective signals alone do not necessarily indicate whether a post should be ranked higher or lower.

Finally, it is important to differentiate between affective response and the more prevalent work in Affective Computing on emotion detection. In the context of social media, emotion detection has been used to try to detect the emotion of the author of the post (referred to as the **publisher affect** in Chen et al. [3]), whereas our focus is on the anticipating the viewer’s affective response to the post. While the publisher affect may be relevant to the affective response, it is not always sufficient signal (see Figure 7 for an illustrative example). For example, a post with an excited publisher affect can induce an angry affective response. We decided to focus on affective response for two main reasons. First, detecting the emotional state of the author of a post is ethically questionable. Second, our goal here is to curate content that caters to the preferences of the user,

and therefore the possible affective responses should be taken into consideration.

2.3 Defining the affective taxonomy

In determining the set of affective responses to operationalize, we need to construct a taxonomy of labels with the proper granularity. There are two challenges in designing the taxonomy. First, the labels need to be discriminative enough to distinguish different use cases and serve a variety of user preferences. Take, for instance, the affective response of being *angered*. On social media, there are instances where collectively venting over a common issue can permit self-expression and community, and can be cathartic and useful to both the posters and viewers [17, 53]. On the other hand, there is clearly an unproductive type of anger that can arise, such as when viewing posts containing spam, toxic speech, or misinformation [4]. In the former case, we may still consider showing the user the candidate post, but in the latter case it's much more unlikely. Consequently, we realized that a single label for *angered* was too broad, and we constructed two types of *angered*: *constructively-angered* and *deconstructively-angered*. Another example is excited and relaxed, which in other works such as [43] are classified under a single category *joy*.

The second challenge we faced was to design a taxonomy that covers the important use cases but, for practical considerations, also minimizes the number of affective responses as much as possible. As noted in [19], the degree of difficulty in modeling affective responses would scale with an increase in the number of affects. In particular, the quality of human labels would decrease because it is impractical to ask annotators to distinguish between a large number of affects. In Section 4, we evaluate the interrater correlation and show average values much larger than that of state-of-the-art work in publisher affect detection.

Table 1: Affective responses (class) and their corresponding definitions of our taxonomy.

Class	Definition
Adoring	Response to finding something adorable.
Connected	Feeling more connected (either to the person making the post or something in the post).
Constructively-angered	Angered in response to content that is angering but informative, valuable, or promoting social good.
Destructively-angered	Angered or annoyed in response to unproductive, unhealthy, borderline violating content.
Entertained	Finding entertaining, amusing, or humorous.
Excited	Feeling joy, excitement, enthusiasm or eagerness.
Grateful	Feeling grateful or appreciative.
Informed	Feeling of being informed or having received new information regarding a subject, event, or topic.
Inspired	Feeling inspired, motivated, uplifted, or encouraged.
Neutral	Having a neutral feeling.
Relaxed	Feeling peaceful, calm, or relieved.
Saddened	Feeling grief, unhappy, sad.
Scared	Feeling of concern, anxiety, fear, or stress.
Surprised	Feeling shocked or astonished (either +/-).
Touched	Feeling moved, emotionally stirred.

3 METHODS

In this section, we describe methods for generating training data for our affective models and the architecture and features of the model that is trained on this data. Training data can be generated in two ways: engagement data that we have on the platform and human labeling of posts. Our goal is to leverage the engagement data as much as possible to reduce the costs of human labeling. Section 3.1 describes how we extract training data from these simple behaviors on the network (e.g., reactions to post). Note, we could have also relied on explicitly asking users their affective states regarding content, but this has the drawbacks of being intrusive and potentially unreliable [36]. In Section 3.2 we analyze the content of comments written in response to posts to generate training labels, and in Section 3.3 we discuss how we obtained annotations from human labelers. Together, the affective response labels and the engagement labels are used for training a two-tower architecture multi-class classifier described in Section 3.4. For this work, only de-identified Facebook posts were utilized.

3.1 Engagement signals

When using the platform, there are a number of ways that users can indirectly give feedback on how they feel about the content that they see. These include but are not necessarily limited to:

- User reactions (e.g. like, love, care, haha, wow, sad, angry as in Figure 8)
- User behaviors (e.g. share, outbound click)
- Negative user feedback (e.g. hide, snooze, unfollow, report as in Figure 9)

While these signals may not be directly indicative of an affective state, a subset of these categories could conceivably provide useful signals that are transferrable to learning the affective response. For instance, a user clicking the haha reaction to a post might indicate that they're feeling entertained by that content, while a user reporting a post might indicate that they are angered or offended. These engagement signals are also straightforward to incorporate into our model because they are already personalized (i.e., the engagement signal is both user and post dependent).

In the dataset we constructed, we used prediction of these engagement signals from the prior 90 days as a new training task (i.e., as prediction labels). Another possibility would be to use these engagement signals as features, rather than training labels, because they intuitively could help in predicting the affective response. There are two issues with this. Firstly, removing engagement as training labels, would greatly reduce our overall training set size and force us to learn a more complex feature space with less data. Secondly, we ideally would like to predict affective response at the time of the post's creation when there is little to no engagement yet. Consequently, assuming engagement signals are available to be used as features was not a suitable design.

3.2 Patterns from comments

In addition to engagement, comments that users write in response to posts contain valuable signal relevant to the affective response they had when viewing the post. For example, the expression "What a hilarious story" may indicate that a post is humorous, and "This is so cute" may indicate that a post is adorable. If comments of the

same flavor appear multiple times in a response to a post, we can use that as a label for the affective response.

We developed the CARE (Common Affective Response Expression) method [59], a means of obtaining labels for affective response in an unsupervised way from the comments written in response to online posts. Since these labels were going to be used as training data, we wanted to ensure that their precision is high, but we also wanted a flexible method that can be applied for new affects as they came up.

CARE uses patterns and a keyword-affect mapping to identify expressions in comments that provide high-precision evidence about the affective response of the readers to the post. We seed the system with a small number of high-precision patterns and mappings. We then iteratively and automatically expand on the initial set by considering frequent patterns and keywords in unlabeled comments on posts labeled by the previous iteration. The CARE method is illustrated in Figure 2. We stopped expansion of the system after reaching 23 distinct patterns and a lexicon of 163 keywords because this sufficed to generate enough labels for each class. Note, because these patterns are applied to the comments rather than the post, training any models on the content of the post will not be biased to these particular expressions.

Using the CARE method, we obtained 4 million labels for the affects *adoring*, *entertained*, *excited*, *saddened*, *scared*, *angered*, and *approving* (which is not in Table 1 but refers to expressing support, praise, or pride). To evaluate the quality of the labels generated by CARE, we randomly selected 6000 posts and asked human annotators to label them according to these labels. The comparison of the human annotations and the CARE labels are shown in Table 2 and indicate high agreement.

In implementing CARE, we noticed a few shortcomings of the method. Firstly, in analyzing the individual classes and mappings, there are some patterns which work for certain class-keyword mappings but not so for others. While this is a point for improvement in future work, we observe that in large numbers, the more error-prone combinations are infrequent compared to the highly-accurate ones. Consequently, the method as a whole is a reasonable cost-effective alternative to obtaining more human annotations.

# Agree	Any CARE	All CARE	Other
≥ 1	98	96	82
≥ 2	94	90	53
$= 3$	80	76	24

Table 2: The rate of agreement between the annotators and the labels proposed by CARE. The first column specifies the number of annotators to be used for consensus. The rest of the columns shows for all posts, the average rate of intersection of the human labels with at least one CARE label, all CARE labels, and any label that is not a CARE label.

3.3 Human labels

In addition to weakly-labeled data from engagement signals and comments, we also crowdsourced a ground-truth dataset, which allowed us to augment our original data with high-quality labels

that can be used for evaluation and analysis. In our annotation tasks, we restricted to posts with just text and image (no video) and to English only.

Labeling guidelines. Before deciding the details of the annotation procedure, it was important for us to first understand which frame of reference we wanted annotators to label from. Unlike in most labeling frameworks which ask objective tasks (e.g., is the post about baseball?), a labeler’s background and personality can greatly affect their answers in our context. More specifically, in asking individuals about the affective response to a post, they could either answer with their personal opinion or they could answer with what they perceive is a more universally accepted answer:

- (1) Personalized: The affective response of a post from a labeler’s personal perspective (i.e., how does this post make *you* feel?).
- (2) Unpersonalized: The affective response of a post from a common, universal perspective (i.e., how does this post make *most people* feel?). One could equate this framing to how genres on Netflix are labeled (e.g., feel-good, emotional, provocative).

It is clear that (1) lends itself best to personalized predictions and is more akin to the recommendation system setting. However human annotators are not our users and therefore we lack their user engagement history and other information that is normally critical for personalized prediction. Moreover, for each post, we have five distinct labelers annotate, which is not sufficiently large to leverage annotator agreement if operating under (1). For these reasons, (2) was the more effective solution, and so we asked the following question: How might someone feel after seeing the following post? Select up to 3 of the top options.

Personalizing annotated posts. As discussed previously, the labels we obtained from human annotation were not personalized, but our recommendation system is personalized. In order to incorporate the non-personalized human labels, we construct a personalized dataset with the following heuristic: if a user liked (or loved) a post that was annotated with an affect A by the annotators, we assume that the user also had the same affect towards the post. Specifically, for each post p labeled as a positive affect A by the human labelers, we collect all users \mathcal{U} who liked or loved the post. We then add a row (p, A, u) to the personalized dataset for each $u \in \mathcal{U}$. We note that this heuristic does not apply to the negative affects (*angered*, *saddened*, and *scared*), since ‘liking’ an angering post does not naturally imply that the user felt angered. Personalizing these negative affects remains as future work and they were withheld from the modeling stage. This personalization process also applied to the CARE labels discussed in Section 3.2 but was not necessary for the engagement-sourced labels since those are already personalized.

3.4 Modeling

Now that we have discussed each component of the training data, we combine all engagement signals (all reactions, unfollow, report, hide, share, and outbound click) and all affective response classes except *angered*, *saddened*, *neutral*, *other*, and *scared* (excluded due to the reasons discussed in the previous paragraph), resulting in

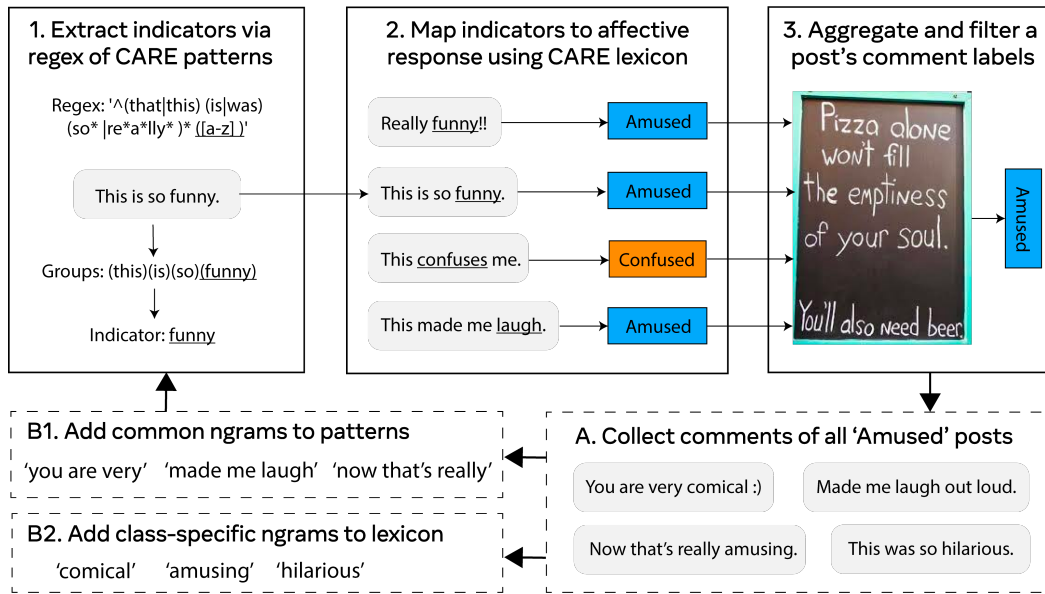


Figure 2: Overview of the CARE Method. The top part of the figure shows the process of labeling a post, while the bottom features how we expand the set of patterns and lexicon. In step (1), we apply CARE patterns to the comment and extract the indicator or keyword. In step (2), we map each comment to the corresponding affective response using the CARE lexicon. Step (3) aggregates the comment-level labels for a post-level label. In step (A) we collect all comments of all posts corresponding to a particular affective response and analyze its most frequent n-grams. N-grams common to multiple classes are added to the CARE patterns (B1) while frequent n-grams specific to a class are added to the lexicon (B2). This process can then be repeated until a satisfactory number of labels are generated.

a total of 23 classes. The training process involves taking a subset of 1 million samples from each prediction class, where half are positive samples and half are negative samples. We create a train/validation/test split of 80/10/10, respectively, on a dataset totaling 23 million samples.

Conceivably, our multi-label classification model trained on affective signals can be used for a number of applications. In some uses cases, personalization will be necessary, and in others, only content information be necessary. In order to have a model which is flexible to the needs of the downstream use cases, we used a two-tower architecture model (see Figure 3) where the left tower is used to model information about the content, and the right for the user. Each tower features a linformer transformer [55] followed by a multi-layer perceptron (MLP) module. The outputs of each tower are then fused in a secondary MLP module, which is then used for multi-label classification.

More concretely, the input to the model consists of the features for both the content and user tower. The content features, for instance, consist of properties of the post like the text of the title, body, optical character recognition, and video transcript, if available. The user features, on the other hand, consist of statistics from the user’s network, interests, and profile properties such as text from a user’s biography. The output of the model is a multi-label prediction for the 23 prediction classes (i.e., a vector of length 23 with binary values). After conducting a sweep over epochs and learning rate, we found training with 3 epochs and a learning rate of 0.0007 to be optimal.

4 RESULTS

In this section, we provide results and analysis of the dataset and model. Specifically, we first discuss statistics pertaining to the human-labeled dataset (e.g., annotator agreement) and second, we will discuss correlation with alternative labels, such as those discussed in Section 3.1. Lastly, we describe results for incorporating the trained model into the overarching recommendation system.

4.1 Analysis of human annotations

We collected human labels for nearly 820k posts with five annotators each, resulting in a total of 7.3 million annotations. In total, there were 348 unique human annotators and on average, each annotator selected 2.57 options per post. In analyzing the number of annotations per post, we compute these statistics using two sets of labels: labels selected by at least one out of the five annotators (1x) and labels selected by at least three of the five annotators (3x). Figure 4 shows the breakdown by class under both settings and indicates that the labels *informed*, *excited*, and *connected* are among the top most prevalent affects. Note, we did not source posts randomly. Instead, we iteratively trained simple binary models to predict for each affect, applied inference to a set of randomly selected posts, and then sent posts with high prediction scores (particularly for low-volume and important classes such as *entertained* and *inspired*). This was done in an effort to screen out neutral posts and reduce labeling costs. Thus, the distribution in Figure 4 is not reflective of the sampling distribution.

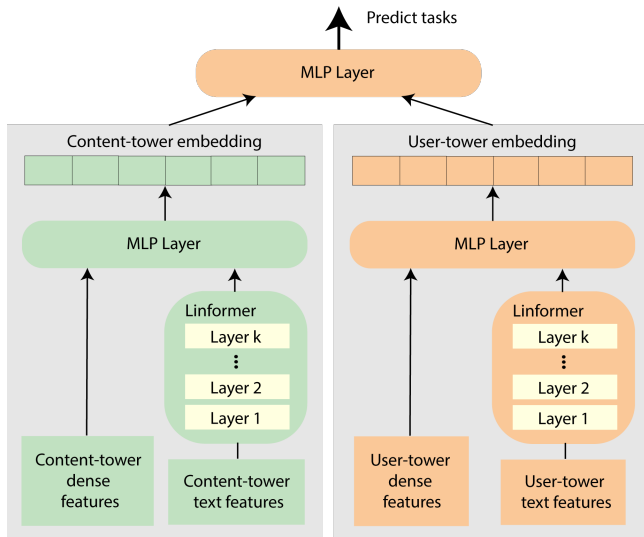


Figure 3: Architecture of the two-tower model. The left and right tower encode content and user information, respectively. The model is trained using the engagement signals and the affective response labels described in Section 3. Each tower encodes text features and dense features, where dense features are typically embeddings from other models which may capture high-dimensional information. One such embedding would be one that describes the visual components of an image in the post.

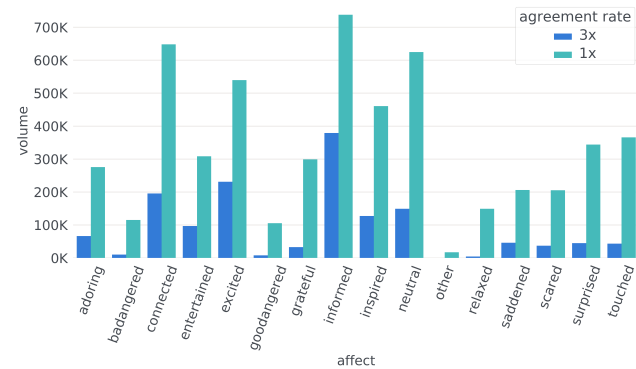


Figure 4: Number of annotations per affect where at least 3 out of 5 annotators agree (3x) or where any annotator selects the label (1x).

Naturally, the number of labels agreed upon by 3 annotators (3x) is smaller than labels with no agreement restriction (1x). The average degree of annotator support for each affect (given that at least one annotator suggests the affect) is shown in Figure 5. The numbers suggest that the classes *other*, *relaxed*, the two types of *angered*, and *grateful* have the lowest degree of agreement given 5 annotators. The *other* category here was an option in the labeling process to submit alternative affects not listed in our taxonomy. Some of the

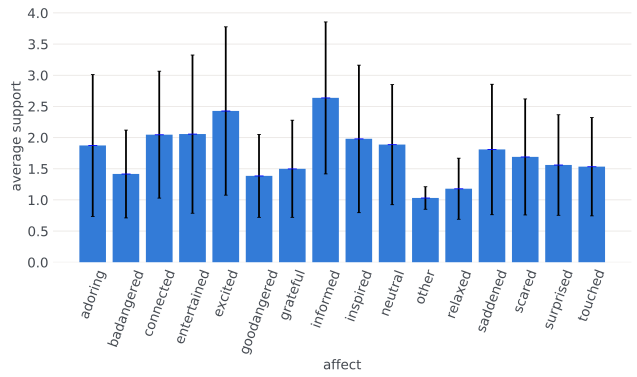


Figure 5: Average degree of annotator support for each class. Error bars indicate standard deviation across all posts.

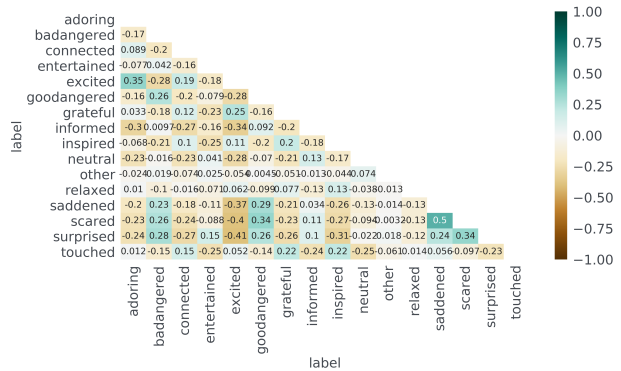


Figure 6: Pearson correlation between the different affective responses in the taxonomy.

most frequent suggestions were *confused*, *curious*, *yummy*, *beautiful*, *disgusted*, and *annoyed*, and these alternative suggestions are valuable for improving future iterations. We did consider distinguishing *disgusted* and *annoyed* from *destructively-angered*, but concluded that the use cases are too similar to justify additional labels.

Following the work of Demszky et al. [7] on publisher affect, we estimate rater agreement by interrater correlation [6], which is computed by taking the average correlation between each rater’s judgement and the mean of other rater judgements. We find that the average inter-rater correlation in our context is 0.52, which is much higher than the inter-rater agreement of 0.28 in [7], where they had 28 classes in their taxonomy (and only 3 labels).

To investigate whether the low agreement of certain classes is a consequence of conceptual overlap amongst the labels, we study the correlation between the affects in Figure 6. Here we use the interpretation of moderate (0.40–0.69) and weak (0.10–0.39) correlation as described in Schober et al. [41]. Figure 6 shows that weak correlations are indeed present. For instance, *constructively-angered* and *deconstructively-angered* are weakly correlated with each other (0.26) as well as several other affects like *scared*, *saddened*, and *surprised* (0.23 to 0.34). Other correlations that are weak to moderate include *adoring* and *excited* (0.35) as well as *saddened* and

scared (0.5). Interestingly, the affective response *connected*, which is not typically in traditional emotion detection taxonomies, seems to have weak correlations with feelings of excitement (0.19), gratitude (0.12), inspiration (0.1), and feeling touched (0.15).

4.2 Publisher affect vs. affective response

In Section 2.3 we distinguished between publisher affect of the poster and the affective response of the viewer, noting that the two are not always interchangeable. In what follows, we experimentally validate and quantify this hypothesis.

Like other social networking platforms, there exists a feature allowing posters to set their status updates, particularly with feelings such as *feeling blessed* or *feeling sad*, which we refer to as the poster-annotated feelings. Because the poster-annotated feelings come from the users themselves and denote their affective state, we can consider these labels as the publisher affect of the post. Hence, we can test our hypothesis given these two sets of labels: publisher affect from the user and affective response from human annotation.

For this analysis, we identify posts in our human-annotated dataset which contain poster-annotated feelings and compute their correlation (after filtering out feelings with frequency less than 1000). As shown in Figure 11, we find moderate correlations between *feeling sad* and *saddened* (0.5) and weak correlation for *feeling worried* and *scared* (0.36), which are intuitive relationships. Overall, however, these correlations between equivalent publisher affects and affective response such as for *excited* (0.12) and *grateful* (0.11) are weak, suggesting that making a distinction between publisher affect and affective response is valid. Interestingly, *touched* is weakly correlated with *feeling sad* (0.25) and *deconstructively-angered* and *constructively-angered* are weakly correlated with *feeling angry* and *feeling annoyed* (0.19 to 0.25), though this is more so the case for *deconstructively-angered*. While there is certainly noise in making these comparisons, the low degree to which these two types of affects correlate indicate that they are not completely synonymous.

4.3 Affective response vs. engagement

While the poster-annotated feelings give insight into the poster’s perspective, we also want to understand how affective response aligns with the viewer’s actual behavior and feedback. Figure 12 shows the correlation between several prevalent engagement signals and affective response labels. Firstly, the behavior and negative user feedback signals (*nufs*) don’t seem to have significant correlation with any of the affects (< 0.042), but this is perhaps due to their low overall prevalence (less than 1000). For reactions, where data is more abundant, we find that the anger, haha, and sorry reactions are weakly correlated with the affective responses *angered* (0.14 to 0.25) and *entertained* (0.41), and *saddened* (0.36), respectively, as one might anticipate. Intuitively, like and love reactions are associated with positive affects and inversely so to negative affects. The support reaction seems to be correlated most with *saddened* (0.14), which empirically seems to be because it is often used in response to expressing concern or sympathy to sad news. As seen earlier in Figure 6, the anger reaction is also weakly correlated with *saddened* (0.13), *scared* (0.16), and *surprised* (0.11). These results together suggest that some affective signals can be gleaned from engagement, as discussed in Section 3.4.

4.4 The affective model in the recommender

This section provides an analysis of the end-to-end system that incorporates the affective predictors. To evaluate our affective model, we experimented with using our model embedding in one of the scoring models of the recommendation system. Given a post and a user, the scoring model we chose tries to predict a user’s answer to a survey concerning their preferences regarding the post. We chose this predictor for two reasons: (a) understanding whether a user wants to see more of a particular content necessitates understanding the user’s affective response to a given post and (b) since this scoring model tries to predict the result of a survey, it by nature has much less training data, and hence can potentially benefit from the data-rich affect embeddings.

We conducted several ablation experiments, which involved a two-step process. The first step involved ablating features and parameters such as the number of encoding layers. The second step involved exporting the 32-length embedding from the content tower, and using this as a feature in the scoring model. Here we prefer to use the embedding from the content tower for internal infrastructure efficiency reasons, but we experimented with both the content and user tower embeddings (including the concatenation of both) and found the results for the latter to be only marginally higher than the content embedding alone.

To evaluate, we first created a static dataset of around 1.5 million survey responses. After running ablation experiments using this offline dataset that was split into train, validation, and test partitions, we identified an embedding with the highest statistical improvement. This model achieved an AUC-ROC loss reduction (the observed improvement normalized by the possible amount of improvement) of more than 8%, as computed by $\frac{S_{new} - S_{base}}{1 - S_{base}} * 100$ where S_{new} and S_{base} refer to the AUC-ROC of the model with the embedding and model without the embedding, respectively. We then conducted an online experiment involving more than 20 million users using the scoring model which uses the new affective embedding (the test group) and using the original scoring model running in production, which does not use the affective embedding (the control group). During the 14 days of online experimentation, we measured a number of metrics relevant to user satisfaction with the platform and benchmarked these values against those of the control group.

The results showed statistically significant decreases in visibility of integrity-violating content like misinformation and engagement bait (more than 0.6% decreases), and also demonstrated meaningful gains in engagement, like the number of like reactions (around 0.4% improvement), without causing detriment to other key important metrics. Additionally, the affective embedding ranked as the most important feature in the scoring model, reaffirming our approach to ranking from an affective response perspective. It also suggests that users, when given more control over their content, will overall choose higher quality content that encourages greater engagement. After launching to an even larger population, these overall trends still generally hold.

5 RELATED WORK

In this section, we situate our work with respect to previous research on related tasks.

5.1 Affective Recommender Systems

Our work is the first to demonstrate that affective signals can benefit recommender systems at large scale in the context of social networks. The following are related work in the field. A significant portion of these study implicit physiological signals like facial and audio tracking [21, 45, 50, 54] while others rely on explicit surveys [11, 25, 49], both of which are not feasible for social media at scale. Some of these works also focus on clean and curated multi-media datasets that are impractical for real-world settings [10, 33, 34, 60]. Orellana-Rodriguez et al. [35] applies affective recommendation to social media, particularly Youtube videos, by acquiring affective annotations from 80 human annotators along Plutchik’s eight basic emotions [38]. Much larger in scale is the work done by Mizgajski and Morzy [29] and Leung et al. [26] using feedback to online news and tweets, respectively, but these works do not leverage multiple types of affective sources. Qian et al. [39] combines user rating data, user social network data, and sentiment from user reviews as affective information, but again is small in scale. Additionally, many of these works utilize traditional machine learning techniques like similarity-based clustering or regression trees as the basis for their recommendation system [29, 35]. We note that this work also differs from others in the space of affective models in that it utilizes a two-tower architecture to jointly model the user and content features, particularly for live online prediction serving billions of users.

5.2 Methods for unsupervised labeling

A major bottleneck in developing models for emotion and affective response detection is the need for large amounts of training data. As an alternative to manually-labeled data, many works utilize metadata such as hashtags, emoticons, and Facebook reactions as pseudo-labels [14, 30, 46, 56]. The work we present here extracts labels from both engagement like Facebook reactions as well as free-form text in comments rather than metadata. The work done in Sintsova and Pu [44] is similar to our work on comments in that it pseudo-labels tweets and extends its lexicon, but the classifier itself is a keyword, rule-based approach and is heavily reliant on the capacity of these lexicons. In contrast, our work leverages the high precision of CARE on the comments and uses the post content to train a model, which is not constrained by the lexicon size in its predictions. Our method also employs bootstrapping to expand the set of patterns and lexicon, similar to Agichtein and Gravano [1] and Jones et al. [18] but focuses on extracting affect rather than relation tuples. Many works utilize engagement and social network structure as features instead of labels in their model [15, 27, 39, 57], but as explained in Section 3.1, our application needs to perform inference prior to engagement signals being available.

5.3 Affective taxonomies

Perhaps two of the most well known categorical organizations for emotion are Paul Ekman’s six basic emotions (happiness, sadness, disgust, fear, surprise, and anger) [8] and Robert Plutchik’s Wheel of Emotions (anger, anticipation, joy, trust, fear, surprise, sadness, and disgust) [38]. Arguably all of Ekman’s six basic emotions exist in our taxonomy, with the exception of disgust which is assumed by deconstructively-angered. These basic emotions are hardly sufficient,

which is in line with Plutchik’s theory that suggests few experiences are basic ones—they are often combination results, which necessitates the need for a more comprehensive taxonomy in practice. The Flickr LDL dataset [58], for example, contains images labeled according to a taxonomy that uses Ekman’s six but also includes amusement and contentment, akin to entertained and relaxed in our current work. We also know from prior work that adequately detecting inspirational [16] and informative [28, 32] content as well as content expressing gratitude [42] is beneficial for users. Craig et al. [5] also found that the primary reasons adolescents use social media is because they want to be entertained, be informed, and feel connected to others. Our taxonomy builds upon prior taxonomies, but includes affects with the intention to satisfy these user needs.

6 CONCLUSIONS

We described the challenges involved with incorporating affective signals in a large-scale recommendation system and the solutions we developed at Facebook. In particular, we designed an affective taxonomy customized to user needs on social media, and created training data for our models by combining engagement data and data from a human-labeling task. Our two-tower model learns from both engagement signals and affective response labels. Our results also provide new insights into the correlations among the affects in the taxonomy and correlations between publisher affect and viewer affective response, thereby justifying some of the design choices we made. We demonstrated that exporting the embedding of this model and using it as feature in one of the scoring models of the recommendation system greatly improves performance, both online and offline.

There are several avenues for improvement and additional research. Our taxonomy can be extended with affective responses that the human annotators frequently noted as missing and our techniques for personalizing human labels need to be extended to negative affects. We believe that more advanced analysis of images and videos can improve our models considerably. More broadly, our work considered the affective response the user may have to a *single* post. However, it is not clear how these individual affective responses combine to an affective response for a session that includes a *sequence* of posts, which is closer to the overall experience the user has on the platform.

REFERENCES

- [1] Eugene Agichtein and Luis Gravano. 2000. Snowball: Extracting relations from large plain-text collections. In *Proceedings of the fifth ACM conference on Digital libraries*, pages 85–94.
- [2] Sonia Bergamaschi and Laura Po. 2014. Comparing lda and lsa topic models for content-based movie recommendation systems. In *International conference on web information systems and technologies*, pages 247–263. Springer.
- [3] Yan-Ying Chen, Tao Chen, Winston H Hsu, Hong-Yuan Mark Liao, and Shih-Fu Chang. 2014. Predicting viewer affective comments based on image content in social media. In *Proceedings of international conference on multimedia retrieval*, pages 233–240.
- [4] Matteo Cinelli, Andraž Pelicon, Igor Mozetič, Walter Quattrociocchi, Petra Kralj Novak, and Fabiana Zollo. 2021. Dynamics of online hate and misinformation. *Scientific reports*, 11(1):1–12.
- [5] Shelley L Craig, Andrew D Eaton, Lauren B McInroy, Vivian WY Leung, and Sreedevi Krishnan. 2021. Can social media participation enhance lgbtq+ youth well-being? Development of the social media benefits scale. *Social Media+ Society*, 7(1):2056305121988931.
- [6] Rosario Delgado and Xavier-Andoni Tibau. 2019. Why cohen’s kappa should be avoided as performance measure in classification. *PLoS one*, 14(9):e0222916.

- [7] Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. GoEmotions: A dataset of fine-grained emotions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4040–4054. Online. Association for Computational Linguistics.
- [8] Paul Ekman. 1999. Basic emotions. *Handbook of cognition and emotion*, 98(45-60):16.
- [9] Belén Saldías F. and Rosalind W. Picard. 2019. Tweet moodifier: Towards giving emotional awareness to twitter users. In *8th International Conference on Affective Computing and Intelligent Interaction, ACII 2019, Cambridge, United Kingdom, September 3-6, 2019*, pages 1–7. IEEE.
- [10] Miguel Á García-Cumbreras, Arturo Montejo-Ráez, and Manuel C Díaz-Galiano. 2013. Pessimists and optimists: Improving collaborative filtering through sentiment analysis. *Expert Systems with Applications*, 40(17):6758–6765.
- [11] Gustavo Gonzalez, Josep Lluís De La Rosa, Miquel Montaner, and Sonia Delfin. 2007. Embedding emotional context in recommender systems. In *2007 IEEE 23rd international conference on data engineering workshop*, pages 845–852. IEEE.
- [12] Chuan Guo, Juan Cao, Xueyao Zhang, Kai Shu, and Miao Yu. 2019. Exploiting emotions for fake news detection on social media. *CoRR*, abs/1903.01728.
- [13] Alon Y. Halevy, Cristian Canton-Ferrer, Hao Ma, Umut Ozertem, Patrick Pantel, Marzieh Saeidi, Fabrizio Silvestri, and Ves Stoyanov. 2020. Preserving integrity in online social networks. *CoRR*, abs/2009.10311.
- [14] Maryam Hasan, Emmanuel Agu, and Elke Rundensteiner. 2014. Using hashtags as labels for supervised learning of emotions in twitter messages. In *ACM SIGKDD workshop on health informatics, New York, USA*.
- [15] Xiangnan He, Ming Gao, Min-Yen Kan, Yiqun Liu, and Kazunari Sugiyama. 2014. Predicting the popularity of web 2.0 items based on user comments. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, pages 233–242.
- [16] Oana Ignat, Y-Lan Boureau, Jane A. Yu, and Alon Halevy. 2021. Detecting inspiring content on social media. In *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 1–8. IEEE.
- [17] Harri Jalonen. 2014. Social media—an arena for venting negative emotions. *Online Journal of Communication and Media Technologies*, 4(October 2014-Special Issue):53–70.
- [18] Rosie Jones, Andrew McCallum, Kamal Nigam, and Ellen Riloff. 1999. Bootstrapping for text learning tasks. In *IJCAI-99 Workshop on Text Mining: Foundations, Techniques and Applications*, volume 1. Citeseer.
- [19] David R Karger, Sewoong Oh, and Devavrat Shah. 2013. Efficient crowdsourcing for multi-class labeling. In *Proceedings of the ACM SIGMETRICS/international conference on Measurement and modeling of computer systems*, pages 81–92.
- [20] Gohar Feroz Khan, Bobby Swar, and Sang Kon Lee. 2014. Social media risks and benefits: A public sector perspective. *Social science computer review*, 32(5):606–627.
- [21] Joep JM Kierkels, Mohammad Soleymani, and Thierry Pun. 2009. Queries and tags in affect-based multimedia retrieval. In *2009 IEEE International Conference on Multimedia and Expo*, pages 1436–1439. IEEE.
- [22] Funda Kivran-Swaine and Mor Naaman. 2011. Network properties and social sharing of emotions in social awareness streams. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pages 379–382.
- [23] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790.
- [24] Akos Lada, Meihong Wang, and Tak Yan. 2021. Facebook news feed ranking. <https://engineering.fb.com/2021/01/26/ml-applications/news-feed-ranking/>.
- [25] Peter J Lang. 2005. International affective picture system (iaps): Affective ratings of pictures and instruction manual. *Technical report*.
- [26] John Kalung Leung, Igor Griva, and William G Kennedy. 2020. Text-based emotion aware recommender. *arXiv preprint arXiv:2007.01455*.
- [27] Marek Lipczak, Borkur Sigurbjornsson, and Alejandro Jaimes. 2012. Understanding and leveraging tag-based relations in on-line social networks. In *Proceedings of the 23rd ACM conference on Hypertext and social media*, pages 229–238.
- [28] Sreejagadeesh Malla and PJA Alphonse. 2021. Covid-19 outbreak: An ensemble pre-trained deep learning model for detecting informative tweets. *Applied Soft Computing*, 107:107495.
- [29] Jan Mizgajski and Mikolaj Morzy. 2019. Affective recommender systems in online news industry: how emotions influence reading choices. *User Modeling and User-Adapted Interaction*, 29(2):345–379.
- [30] Saif M Mohammad and Svetlana Kiritchenko. 2015. Using hashtags to capture fine emotion categories from tweets. *Computational Intelligence*, 31(2):301–326.
- [31] Jessica Gall Myrick. 2015. Emotion regulation, procrastination, and watching cat videos online: Who watches internet cats, why, and to what effect? *Computers in human behavior*, 52:168–176.
- [32] Xiaochuan Ni, Gui-Rong Xue, Xiao Ling, Yong Yu, and Qiang Yang. 2007. Exploring in the weblog space by detecting informative and affective articles. In *Proceedings of the 16th international conference on World Wide Web*, pages 281–290.
- [33] Ante Odic, Marko Tkalcić, Jurij F Tasic, and Andrej Košir. 2012. Relevant context in a movie recommender system: Users' opinion vs. statistical detection. *ACM RecSys*, 12.
- [34] Ante Odic, Marko Tkalcić, Jurij F Tasic, and Andrej Košir. 2013. Predicting and detecting the relevant contextual information in a movie-recommender system. *Interacting with Computers*, 25(1):74–90.
- [35] Claudia Orellana-Rodriguez, Ernesto Diaz-Aviles, and Wolfgang Nejdl. 2015. Mining affective context in short films for emotion-aware recommendation. In *Proceedings of the 26th ACM Conference on Hypertext & Social Media*, pages 185–194.
- [36] Maja Pantic and Alessandro Vinciarelli. 2009. Implicit human-centered tagging [social sciences]. *IEEE Signal Processing Magazine*, 26(6):173–180.
- [37] Marco Pennacchiotti and Siva Gurumurthy. 2011. Investigating topic models for social media user recommendation. In *Proceedings of the 20th international conference companion on World wide web*, pages 101–102.
- [38] Robert Plutchik. 1980. Chapter 1 - a general psychoevolutionary theory of emotion. In Robert Plutchik and Henry Kellerman, editors, *Theories of Emotion*, pages 3 – 33. Academic Press.
- [39] Yongfeng Qian, Yin Zhang, Xiao Ma, Han Yu, and Limei Peng. 2019. Ears: Emotion-aware recommender system based on hybrid information fusion. *Information Fusion*, 46:141–146.
- [40] Santosh Rajamanickam, Pushka Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2020. Joint modelling of emotion and abusive language detection. In *Proceedings of the annual meeting of the association for computational linguistics*. Association for Computational Linguistics.
- [41] Patrick Schober, Christa Boer, and Lothar A Schwarte. 2018. Correlation coefficients: appropriate use and interpretation. *Anesthesia & Analgesia*, 126(5):1763–1768.
- [42] Simona Sciarra, Daniela Villani, Anna Flavia Di Natale, and Camillo Regalia. 2021. Gratitude and social media: A pilot experiment on the benefits of exposure to others' grateful interactions on facebook. *Frontiers in Psychology*, 12.
- [43] Phillip Shaver, Judith Schwartz, Donald Kirson, and Cary O'connor. 1987. Emotion knowledge: further exploration of a prototype approach. *Journal of personality and social psychology*, 52(6):1061.
- [44] Valentina Sintsova and Pearl Pu. 2016. Dystemo: Distant supervision method for multi-category emotion recognition in tweets. *ACM Trans. Intell. Syst. Technol.*, 8(1).
- [45] Mohammad Soleymani, Maja Pantic, and Thierry Pun. 2011. Multimodal emotion recognition in response to videos. *IEEE transactions on affective computing*, 3(2):211–223.
- [46] Jared Suttles and Nancy Ide. 2013. Distant supervision for emotion classification with discrete binary values. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 121–136. Springer.
- [47] Luis Tapia. *Clown on black background*. Shutterstock.
- [48] Mike Thelwall, David Wilkinson, and Sukhvinder Uppal. 2010. Data mining emotion in social network communication: Gender differences in myspace. *Journal of the American Society for Information Science and Technology*, 61(1):190–199.
- [49] Marko Tkalcić, Urban Burnik, and Andrej Košir. 2010. Using affective parameters in a content-based recommender system for images. *User Modeling and User-Adapted Interaction*, 20(4):279–311.
- [50] Marko Tkalcić, Ante Odic, Andrej Kosir, and Jurij Tasic. 2012. Affective labeling in a content-based recommender system for images. *IEEE transactions on multimedia*, 15(2):391–400.
- [51] Yalda T Uhls, Nicole B Ellison, and Kaveri Subrahmanyam. 2017. Benefits and costs of social media in adolescence. *Pediatrics*, 140(Supplement 2):S67–S70.
- [52] Erik van Ingen, Sonja Utz, and Vera Toepoel. 2016. Online coping after negative life events: Measurement, prevalence, and relation with internet activities and well-being. *Social Science Computer Review*, 34(5):511–529.
- [53] Anne Vermeulen, Heidi Vandebosch, and Wannes Heirman. 2018. # smiling, # venting, or both? adolescents' social sharing of emotions on social media. *Computers in Human Behavior*, 84:211–219.
- [54] Stefanos Vrochidis, Ioannis Kompatsiaris, and Ioannis Patras. 2011. Utilizing implicit user feedback to improve interactive video retrieval. *Advances in Multimedia*, 2011.
- [55] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*.
- [56] Wenbo Wang, Lu Chen, Krishnaprasad Thirunakaran, and Amit P Sheth. 2012. Harnessing twitter "big data" for automatic emotion identification. In *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*, pages 587–592. IEEE.
- [57] Dingqi Yang, Daqing Zhang, Zhiyong Yu, and Zhu Wang. 2013. A sentiment-enhanced personalized location recommendation system. In *Proceedings of the 24th ACM conference on hypertext and social media*, pages 119–128.
- [58] Jufeng Yang, Ming Sun, and Xiaoxiao Sun. 2017. Learning visual sentiment distributions via augmented conditional probability neural network. In *Thirty-first AAAI conference on artificial intelligence*.
- [59] Jane A. Yu and Alon Y. Halevy. 2022. The CARE dataset for affective response detection.
- [60] Yong Zheng, Bamshad Mobasher, and Robin D Burke. 2013. The role of emotions in context-aware recommendation. *Decisions@ RecSys*, 2013:21–28.

A BROADER IMPACT

Any work that touches upon recognizing affective response needs to ensure that it is sensitive to its application. Our work in detecting affective response is intended for anticipating the affective response of users to content, in order to better safeguard them against offensive material and provide them with content that better aligns with their user preferences. This work should not be used for ill-intended purposes, such as purposefully recommending particular content to manipulate a user’s perception or preferences. Additionally, any work that utilizes user information or content created by users must be careful in respecting the privacy preferences of its users. Before this research was conducted, it went through an extensive internal review process with a diverse team to delineate these bounds. Regarding our crowdsourcing process, the human annotators were paid a competitive and fair rate. The raters were also selected by diversifying the pool amongst several categories along five attributes: age, ideology, gender, ethnicity, and location.

B AFFECTIVE RESPONSE VS. PUBLISHER AFFECT



Figure 7: An example case of differing publisher affect and affective response. This work focuses on affective response through signals such as comments and reactions. Post image sourced from Shutterstock [47].

C FACEBOOK INTERFACE

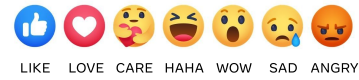


Figure 8: Facebook reactions.

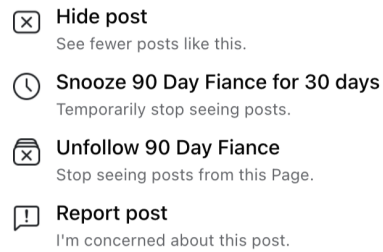


Figure 9: Facebook negative user feedback controls.

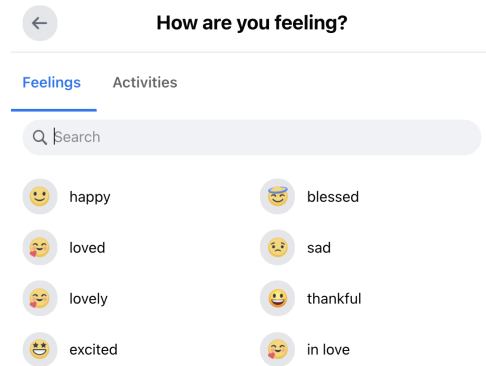


Figure 10: Facebook feature allowing posters to annotate their own posts with feelings.

D CORRELATION BETWEEN AFFECTIVE RESPONSE AND OTHER FACEBOOK SIGNALS

In this section, we show the correlation between affective response and poster-annotated feelings in Figure 11 and between affective response and engagement signals in Figure 12.

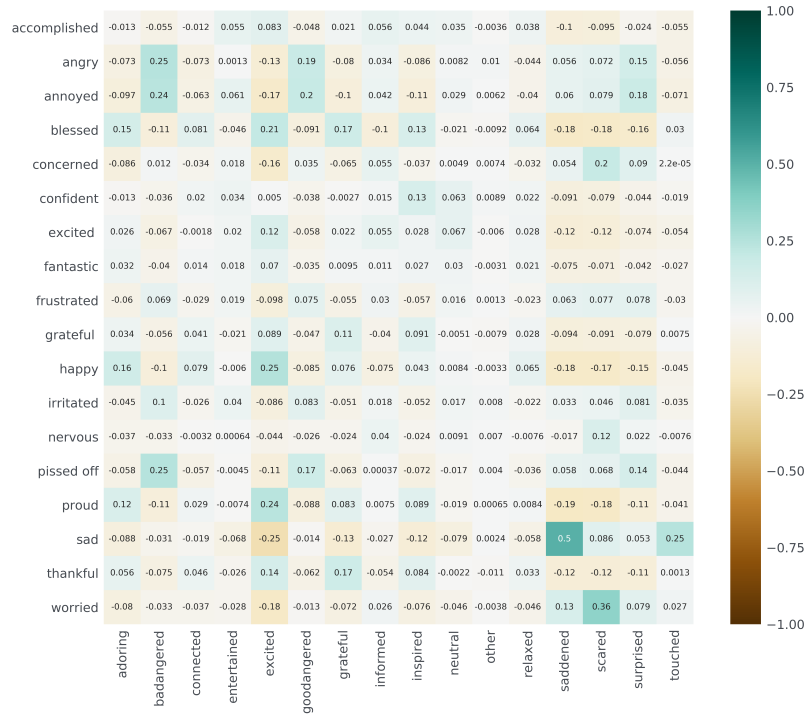


Figure 11: Pearson correlation between human annotations for affective response and poster-annotated feelings.

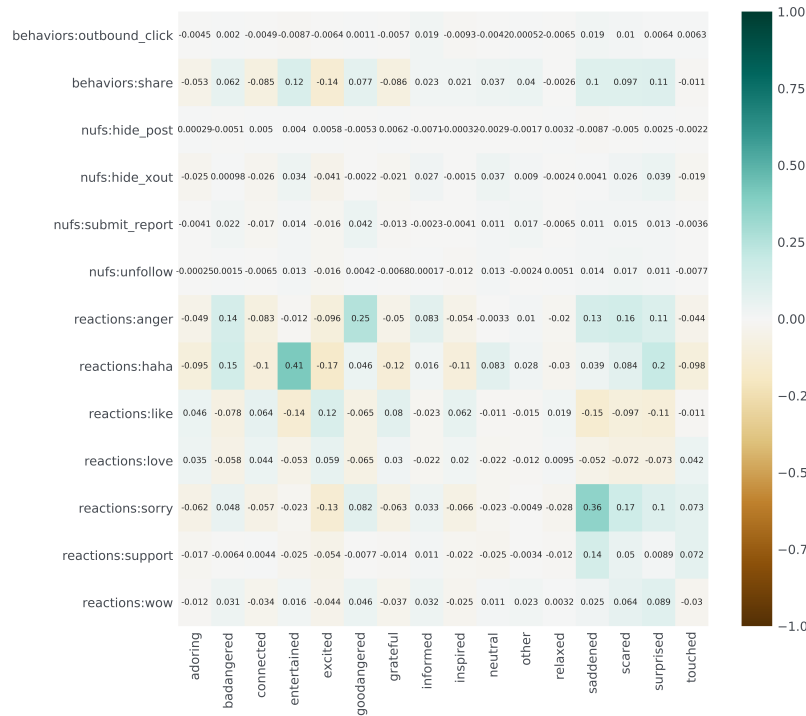


Figure 12: Pearson correlation between human-annotations for affective response and user engagement signals. Values are first normalized by engagement type (i.e., behaviors, negative user feedback, and reactions).