

Deciphering Implicit Hate: Evaluating Automated Detection Algorithms for Multimodal Hate

Austin Botelho*

Anti-Defamation League
abotelho@adl.org

Scott A. Hale

University of Oxford and Meedan
scott.hale@oii.ox.ac.uk

Bertie Vidgen

The Alan Turing Institute
bvidgen@turing.ac.uk

Abstract

Accurate detection and classification of online hate is a difficult task. Implicit hate is particularly challenging as such content tends to have unusual syntax, polysemic words, and fewer markers of prejudice (e.g., slurs). This problem is heightened with multimodal content, such as memes (combinations of text and images), as they are often harder to decipher than unimodal content (e.g., text alone). This paper evaluates the role of semantic and multimodal context for detecting implicit and explicit hate. We show that both text- and visual- enrichment improves model performance, with the multimodal model (0.771) outperforming other models' F1 scores (0.544, 0.737, and 0.754). While the unimodal-text context-aware (transformer) model was the most accurate on the subtask of implicit hate detection, the multimodal model outperformed it overall because of a lower propensity towards false positives. We find that all models perform better on content with full annotator agreement and that multimodal models are best at classifying the content where annotators disagree. To conduct these investigations, we undertook high-quality annotation of a sample of 5,000 multimodal entries. Tweets were annotated for primary category, modality, and strategy. We make this corpus, along with the codebook, code, and final model, freely available.

1 Introduction

Although its prevalence is fairly low (Vidgen et al., 2019), the effects of online hate can be deeply pernicious, risking real harm on targeted victims and their communities (Müller et al., 2019; Guadagno et al., 2013). A 2021 survey by Anti-Defamation League found that 81% of Americans agree social media companies should do more to counter online hate (Anti-Defamation League, 2021).

Research into automated hate detection has primarily focused on explicit varieties. However, many purveyors of hate have adopted more complex and nuanced strategies, such as dog whistling: the use of intentionally ambiguous rhetorical techniques to express hateful messages which only some audiences will recognize. For instance, calls by right-wing American political figures to “protect the suburbs” cloak racial grievances and concerns about whiteness in more prosaic concerns about community protection.

Performance in online hate classification has improved substantially from static methods like GloVe and fastText through the use of context-aware word embeddings, in particular those computed by transformer-models with self-attention (Badjatiya et al., 2019; Mozafari et al., 2020; Polignano et al., 2019; Sabat et al., 2019; Zampieri et al., 2019; Yang et al., 2019; Sohn and Lee, 2019; Kennedy et al., 2020; Vidgen et al., 2020). However, most hate detection models are text-only and cannot be applied to non-textual content (such as images and audio) and do not account for non-textual information contained in multimodal content (such as memes). This is a problem of both task definition and modeling; most hateful content training datasets do not take into account non-textual features when annotations are made, which means that non-textual systems cannot be trained and evaluated on them. The lack of detailed- and expertly-annotated datasets means that many key aspects of multimodal content classification have not yet been explored.

We address these gaps in research, making three primary contributions. First, we present a newly annotated dataset of 5,000 multimodal tweets, with labels for primary category, modality, and strategy. We make the annotation guidelines, code, and best performing models publicly available. Second, we show that as models better take into ac-

*Work performed while at the University of Oxford

count contextualization, from context-invariant to context-aware and unimodal to multimodal, how accurately they detect hateful content significantly improves. Though the unimodal-text context-aware model performs the best on the implicit hate sub-task, the multimodal model is better overall due to a lower propensity towards false positives. Third, we show that all models perform considerably worse on ambiguous content (as determined by annotator disagreement).

2 Related Work

The networked structure of online platforms means that hate is often able to spread far beyond the author's original intended audience (Walther et al., 2011). These "masspersonal" (O'Sullivan and Carr, 2018) networks blur the divide between public and private discourse, resulting in "context collapse" as multiple audiences converge towards a singular unbounded one (Marwick and boyd, 2010; boyd, 2017). This can increase the social costs of spreading hateful messages as wider audiences and platform moderators may disapprove of this content. Consequentially, hateful actors are incentivized to employ implicit rhetorical strategies to circumvent these costs. Whereas explicit forms of hate (e.g., slurs or calls to violence) are likely to draw attention, subtle forms of hate, such as dog whistles, can be effective in avoiding detection.

Dog whistles comprise a range of strategies anchored in polysemy including pseudo-factual claims (Meddaugh and Kay, 2009), normative statements (Pettigrew and Meertens, 1995), coded hate terms (Magu et al., 2017), and artistic license and humor (Milner, 2013) to create implied meanings. This enables hateful actors to target their messages at different audiences such that the hateful elements are only recognized by people who are predisposed to respond favorably (Albertson, 2015). This gives their speakers plausible deniability, allowing them to avoid any social, legal, or platform-based punishment for the content they produce.

Multimodal communication, in particular, memes, are susceptible to co-optation by hateful actors for use as dog whistles because of its ability to convey incongruent ideas through each modality ('modal dissonance'). Hateful content can be passed under the façade of a shared, seemingly benign, meme macro (or 'template') (Zannettou et al., 2018). For instance, Vidgen et al. (2019) describe how a non-hateful image (e.g., a group of Mus-

lims in prayer) can be combined with a non-hateful text (e.g., the words 'Woken up yet?') to express prejudice. If the words or images were changed to something benign then the meme would no longer be hateful). This has led to a culture of 'shit posting' and trolling (Pelletier-Gagnon and Pérez Trujillo Diniz, 2018; Phillips, 2012).

Previous research on automated detection of hate has primarily relied on unimodal approaches with text-based features. These features have been passed through a variety of classification models (Fortuna and Nunes, 2018; Schmidt and Wiegand, 2017). Neural network architectures harnessing advances in convolutions (Gambäck and Kumar Sikdar, 2017; Ribeiro and Da Silva, 2019; Zhang et al., 2018), recurrence (Pitsilis et al., 2018), long-term memory (Badjatiya et al., 2019; Pitsilis et al., 2018), and bidirectionality (Caselli et al., 2018; Qian et al., 2018) have been applied to improved accuracy. However, a shift towards fine-tuning large, pre-trained models has yielded the best results with BERT and its varieties being the models of choice (Sohn and Lee, 2019; Zampieri et al., 2019; Mozafari et al., 2020; Polignano et al., 2019; Vidgen et al., 2020, 2021).

Despite these strides, many challenges persist as real-world interactions are noisy, varied, and multimodal. Most applications of multimodal hate speech detection have combined text with meta-information like user characteristics, comment thread information, and network connections (Chandrasekharan et al., 2017; Fehn Unsvåg and Gambäck, 2018; Gao and Huang, 2017; Qian et al., 2018; Vijayaraghavan et al., 2019). Early examples of combining text and image data yielded mixed results (Gomez et al., 2020; Sabat et al., 2019; Yang et al., 2019) leading companies like Facebook to initiate financial awards for improved performance (Kiela et al., 2020). Success in other domains like identifying pro-eating disorder content (Chancellor et al., 2017), gang activity on social media (Blandfort et al., 2019), demographic inference (Wang et al., 2019), and cyberbullying (Zhong et al., 2016) highlight the potential positive effects of multimodal approaches.

3 Schema

Our taxonomy comprises three main categories: (1) Primary attribute, (2) Modality, and (3) Strategy. The taxonomy and definitions were developed by reviewing existing theoretical frameworks for on-

line hate and multimodal content (Vidgen et al., 2019; Waseem and Hovy, 2016; Kiela et al., 2020; Citron and Norton, 2011) and by iteratively investigating samples of tweets from the dataset.

Primary For the *Primary* attribute annotators selected one of four options: Hate, Counterspeech, Reclaimed, and Neutral. Similar to Davidson et al. (2017, p. 512), “Hate” is defined as “language that is used to express hatred towards a targeted group or is intended to be derogatory, to humiliate, or to insult the members of the group.” This definition’s grounding in group identity differentiates hate from other forms of abusive content (such as interpersonal abuse) and corresponds with the definitions enforced by digital platforms like Facebook, Google, and Twitter (YouTube, 2020; Twitter, 2020; Facebook, 2020).

“Counterspeech” is defined as any response to hateful speech that undermines it or expresses support to a group that it targets. This category is needed as models trained on datasets with only a ‘Neutral’ category may struggle to differentiate between pro-social (Galinsky et al., 2013) tweets and hateful ones if they have similar lexical content. “Reclaimed” is defined as the use of slurs self-referentially, whereby oppressive language is reappropriated for in-group use. This category is particularly important given well-established biases in classification models, whereby they disproportionately classify the vernacular of African Americans (and other groups) as hate (Sap et al., 2019; Davidson et al., 2019). “Neutral” is defined as content which did not fall into the other three categories.

Modality For the *Modality* attribute annotators labeled the modality (image, text, or both) that was informative when making the *Primary* annotation. This was needed because although all entries contained both a text and image, both modes were not always used to express hate. In some cases the hate was expressed solely by the text or by the image, and in others both were used together.

Strategy The *Strategy* attribute captures the rhetorical devices used to convey hate. It expands upon the implicit/explicit distinction proposed by Waseem et al. (2017) and adopted by others, such as Caselli et al. (2020) and Zampieri et al. (2019). Strategy is hierarchical: if annotators identify Hate then they can select “Explicit”, “Psuedo-factual”, “Normative statements”, “Coded

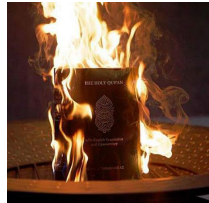




Strategy	Text	Image
Explicit	<pre><user> <user> suck a pig dick cunt <url></pre>	
Pseudo-factual	<pre><user> ille- gal criminals protected by liberals. #buildthewal <url></pre>	
Normative	<pre><user> <user> come on booker bring in the “race” card you always do. <url></pre>	
Coded	<pre>best #npc meme #npcmeme #sjw <url></pre>	
Creative	<pre>good morn- ing tweeps #friday #teamtrump #buildthewal 🔴🔵 <url></pre>	

Table 1: Examples of tweets for each implicit strategy.

language”, or “Creative expressions” (see Table 1). The latter four were collapsed into a single “Implicit” category. These strategies reflect previous research on the varieties of implicit hate (Meddaugh and Kay, 2009; Dvorak, 1999; Poynting and Noble, 2003; Hughey and Daniels, 2013).

4 Dataset and annotation

The dataset used to train and evaluate the models originates from the MMHS150K collected by Gomez et al. (2020). It comprises 150,000 English-language tweets which all contain text and an image, and have been annotated for hate (split into 5 subcategories: racist, sexist, homophobic, religion

based attacks or attacks to other communities).¹

We enrich a sample of 5,000 multimodal tweets from the MMHS150K dataset with re-annotation.² Half of the dataset (2,500 tweets) was sampled using a 17-term query to provide more coverage of hateful tweets, especially covert ones.³ The other half (2,500 tweets) was randomly sampled to ensure heterogeneity and to offset biases associated with focused sampling (Wiegand et al., 2019).

Annotation process

The annotators were given a detailed codebook to inform their decisions, with definitions, prototypical examples, and edge cases. Each tweet was labeled by two annotators. All annotators had prior experience annotating online hate, and each completed a minimum of four weeks of training. Given the frequency of mislabeled hate speech due to a lack of domain expertise (van Aken et al., 2018), we prioritized annotator experience over more scalable crowdsourcing options like Amazon Mechanical Turk. All entries with disagreement were sent for review by an expert annotator. The expert was a PhD student researching online and offline hate, who had previously worked on two annotation projects.

The Kappa score for the dataset is 0.40, indicating low to moderate agreement. However, it is a nearly three-fold increase from the MMHS150k’s Kappa of 0.15 despite the increased difficulty of the task. This is equivalent to other hate speech datasets. For instance, Wulczyn et al. (2017) report Krippendorff’s alpha of 0.45 and Sanguinetti et al. (2018) report category-wise Kappas of 0.37 for *offense* and 0.54 for *hate*.

In 51% of cases annotators disagreed on any of the *Primary*, *Modality*, or binarized *Strategies* – all of which were sent for review by the expert. The annotators agreed on labels far more frequently for tweets in the “None” category (68.8% of the time) than the others. Initial agreement (i.e. before expert adjudication) was 29% for Hateful, 15% for Counterspeech and 14% for Reclaimed. These relatively low agreement levels were primarily because broader situational context is often needed to make these judgments.

¹The dataset can be accessed at: <https://gombru.github.io/2019/10/09/MMHS/>

²The 5,000 tweet dataset can be downloaded from https://github.com/botelhoa/Dog_Whistle_Hate

³The terms are: wall, card, confederate, maga, islam, sjw, gender, crim, npc, normie,))) , muslim, illegal, caravan, obama, hillary, america.

Annotations	Label Breakdown
Primary	Hateful: 1850 (37.0%), Counterspeech: 113 (2.3%), Reclaimed: 366 (7.3%), None: 2,671 (53.4%)
Modality	Unimodal-text: 874 (37.5%), Unimodal-image: 25 (1.1%), Multimodal: 1,430 (61.4%)
Strategy	Explicit: 31.8%, Normative: 30.4%, Coded: 22.3%, Creative: 7.8%, Psuedo-factual: 7.7%

Table 2: Dataset label breakdown.

Annotators agreed less on “Multimodal” tweets than “Unimodal-Text” tweets. This suggests the richer semantic context from the different modalities helps annotators to clarify what is being expressed. Unexpectedly, agreement was higher when labeling implicit rather than explicit hate (note that this is only for whether implicit hate was expressed, rather than identifying the particular technique used). This may be because the codebook contained more explanation of implicit hate, given we anticipated difficulties in annotating for them or because it can be difficult to ascertain when explicit slurs are hateful versus reclaimed when few signals pointing to the author’s identity are available.

Dataset composition

The final labeled dataset is 37.0% Hateful, 2.3% Counterspeech, 7.3% Reclaimed, and 53.4% None (Table 2). Annotators relied on both modalities in a majority of cases (61.4%). In terms of strategy, of the tweets marked Hateful, 36.9% were explicitly hateful and 63.1% were implicitly hateful. The implicit strategies were Normative claims (30.4%), Coded language (22.3%), Creative (7.8%), and Psuedo-factual (7.7%).

5 Model implementation

All models were evaluated using the same 80/10/10 train, validation, and test split, stratified by class across the sets. Computation was completed using a single CUDA-enabled Nvidia Tesla K80 GPU in Google Colab.

5.1 Input features

The curators of the MMHS150k dataset represented all graphics (images, GIFs and video) as thumbnails. They were resized to be a pixel dimension of 500 in the smallest direction while maintaining the original aspect ratio. Textual features are derived from two sources: the tweet body and the image text, extracted using OCR. Both text sources underwent the same pre-processing procedure, using the Ekphrasis Python library.⁴ We de-noised the data by replacing hyperlinks, mentions, and dates with tags, decomposing hashtags into their constituent words, and normalizing elongated words and punctuation, in line with Mozafari et al. (2019) and Sohn and Lee (2019). To retain indicators of sentiment, variables were added for whether capital letters were used, elongated words and punctuation (Hutto and Gilbert, 2014). The remaining text was truncated to a length of 100 tokens. Longer tweets were abridged while shorter ones were padded.

5.2 Models

Four classes of models with varying levels of semantic contextualization were trained, including three classes of unimodal models—unimodal-image models, context-invariant unimodal-text models (LSTM), and context-aware unimodal-text models (Transformers)—and multimodal models. Only the best performing model in each class is reported and their tuning described. Full results are available in the supplemental materials.

For unimodal-image models, Xception, NASNet, and Inception-ResNet V2 (Chollet, 2017; Szegedy et al., 2016a; Zoph et al., 2018) were tested and Xception had the highest performance as measured by the weighted F1 score. For context-aware unimodal-text models, albert-xxlarge-v2, bert-large-uncased, electra-large-discriminator, and roberta-large (Lan et al., 2019; Clark et al., 2020; Liu et al., 2019) were tested with roberta-large performing best. For multimodals, the intermediate concatenations approaches of Gomez et al. (2020) and Sabat et al. (2019) and a joint representation approach (“MMBT”) (Kiela et al., 2019) were tested with the joint representation approach performing best.

5.2.1 Unimodal-Image

The first level of semantic contextualization consists of only image information. The extraction

⁴The documentation is available at <https://github.com/cbaziotis/ekphrasis>

of image features was conducted with Xception (Chollet, 2017). Xception decouples the mapping of cross-channel and spatial correlations by performing a depthwise convolution before a pointwise convolution. This improves Top-1 and Top-5 accuracy on ImageNet compared to Inception and a significant increase in performance on the larger JFT image corpus despite maintaining the same number of model parameters (Chollet, 2017). It has yet to be applied to the task of hateful image recognition but outperforms methods that have (Gomez et al., 2020; Yang et al., 2019; Sabat et al., 2019) in general image recognition tasks (Soo Ko, 2020).

The weights pre-trained on ImageNet were downloaded from the Keras library.⁵ Data augmentation was applied to the images in the train set prior to passing them through the network, including slight random rotations, height and width shifts, and horizontal flips. The images were passed in batch sizes of 32, following the approach of Szegedy et al. (2016b). The weights in the bottom layers were frozen while updates occurred only on the top 5%. A classifier was placed atop the CovNet which used two-dimensional Global Average Pooling followed by a fully connected layer of 1024 nodes with ReLU activation and a SoftMax output layer with dropout. During training, the same hyperparameters proposed by the original paper for the ImageNet task were applied with the addition of the early stopping regularization technique.

5.2.2 Unimodal-Text

In order to understand the impact of textual contextualization on the detection of implicit hate, two approaches to text-only classification were implemented: (1) a context-invariant LSTM model and (2) context-aware transformer-based models.

Context-Invariant LSTM The baseline context-invariant model consists of a bi-LSTM with pre-trained fastText embeddings (Schuster and Paliwal, 1997). This is the strongest alternative to a transformer model because it considers future context (Graves and Schmidhuber, 2005) and approximately represents OOV words via character n-grams (Bojanowski et al., 2016; Joulin et al., 2016).

The LSTM comprised an embedding layer of length 300, two bidirectional LSTM layers of 256 hidden nodes with a dropout of 0.2, and a fully con-

⁵These weights can be found here: <https://keras.io/api/applications/>

nected output layer. It was trained over 50 epochs with early stopping of a 10-epoch patience and 0.01 minimum average validation loss improvement using mini-batches of size 64 bucketed by length to reduce the need for padding. Parameters were optimized using the weighted ADAM algorithm (Loshchilov and Hutter, 2017) with a Cross Entropy Loss function. A hyperparameter search was conducted across the learning rates $\in \{0.0001, 0.001, 0.01, 0.1, 1\}$.

Context-Aware Transformer The context-aware model applies a transformer architecture, namely roberta-large (Liu et al., 2019). Within the transformer, each use of a word is treated independently from its other uses. This means that, in principle, it can distinguish between the phrase “race card” when used in horse racing versus in reference to the view that racial prejudice can be advantageous to its victims. The Transformer model was coded in PyTorch with pre-trained weights loaded through the HuggingFace library.⁶ It was trained for 10 epochs with early stopping of a two-epoch patience and 0.005 minimum average validation loss improvement. Parameters were optimized using the weighted ADAM algorithm (Loshchilov and Hutter, 2017) with a 0.1 weight decay and slanted triangular schedule (Howard and Ruder, 2018) with a warmup of 0.06. Backpropagation was conducted using Cross Entropy Loss. A hyperparameter search borrowed from (Liu et al., 2019) was implemented across the learning rates $\in \{1e-5, 2e-5, 3e-5\}$ and mini-batch sizes $\in \{16, 32, 64\}$.

5.2.3 Multimodal

Lastly, deeper semantic contextualization may be achieved through the inclusion of multimodal data. Such models should, in theory, more accurately identify implicit hate by drawing from information contained by both the image, text, and their interaction. Recent approaches use a transformer’s attention mechanism to generate joint representations of images and text (Kiela et al., 2019; Li et al., 2019; Lu et al., 2020).

The MMBT first encodes image data using ResNet-152 with a generalized final pooling layer pre-trained on the ImageNet dataset. The image embeddings are combined with the tokenized text and passed through a bidirectional transformer ar-

⁶These are available at <https://huggingface.co/transformers/>.

chitecture that was initialized using pre-trained BERT weights before an output layer with SoftMax activation makes the classification (Kiela et al., 2019). This was implemented using the Simple Transformers library.⁷ A batch size of eight was used, with a learning rate of 1^{-5} trained with early stopping, weighted ADAM optimization (Loshchilov and Hutter, 2017) with 0.1 weight decay, and a slanted triangular schedule (Howard and Ruder, 2018) with a warmup of 0.06.

6 Results

6.1 Inter-modal Performance Comparisons

Performance for the strongest models in each of the four modality types is displayed in Table 3. Metrics were computed with a weighted average to accommodate class imbalances. The multimodal model performed the best across the four metrics, albeit only marginally so compared to the second best model, the unimodal-text context-aware model. Both unimodal-text models noticeably outperformed the unimodal-image model.

6.2 Performance by Strategy

To assess the secondary effects of defining training objectives, accuracy based on the *Strategy* annotation was calculated. These labels were not shown to the model to demonstrate performance variations hidden by oversimplified label categories.

The unimodal-text context-aware model had the highest accuracy when identifying both “Explicit” and “Implicit” hate (Table 3). By contrast, the unimodal-image model most accurately identifies non-hateful tweets, but struggles with hateful ones.

6.2.1 Ambiguity

Ambiguity is the final characteristic for which performance was assessed (Table 4). A decision in the *Primary* annotation is considered ambiguous if the two annotators provided conflicting decisions.

All models had significantly higher F1 scores on data deemed unambiguous. The multimodal model dealt the best with ambiguity (0.692) followed closely behind by the unimodal-text context-aware model (0.682). The unimodal-image model was most affected (0.758 vs 0.398) by ambiguity.

7 Error Analysis

We conducted a qualitative analysis on the errors of the multimodal model similar to the one in Vidgen

⁷<https://github.com/ThilinaRajapakse/simpletransformers>

	Overall				Accuracy by Strategy		
	Accuracy	Precision	Recall	F1	Non-Hateful	Explicit	Implicit
Unimodal-Image	0.604	0.560	0.544	0.544	0.760	0.273	0.370
Unimodal-Text							
Context-Invariant	0.737	0.707	0.737	0.719	0.713	0.742	0.798
Context-Aware	0.765	0.759	0.765	0.754	0.678	0.864	0.941
Multimodal	0.785	0.763	0.785	0.771	0.732	0.833	0.899
Count	502	502	502	502	317	66	119

Table 3: Overall performance of models and performance split by strategy. The multimodal model performs best overall. The context-aware unimodal-text model performs best on both implicit and explicit forms of hate.

	Unambiguous	Ambiguous
Unimodal-Image	0.758	0.398
Unimodal-Text		
Context-Invariant	0.827	0.643
Context-Aware	0.875	0.682
Multimodal	0.845	0.692
Count	272	230

Table 4: Model F1 Score by ambiguity. Entries where annotators disagree are considered ambiguous and entries with full agreement are considered unambiguous. All models perform better on Unambiguous content. The multimodal model performs best on ambiguous content.

et al. (2020). Errors were inductively categorized into “mutually exclusive and collectively exhaustive” groups (Vidgen et al., 2020, p. 7). This is visualized in the Tree Diagram in Figure 1.

The first branch on the tree is the *annotator errors*. This occurred when the model’s classification better represented the tweet’s content, based on the annotation codebook (as determined by this paper’s authors). This represented 25% of the errors and suggests that model performance could be improved further if annotator errors were eradicated. Model errors were the most frequent, accounting for 66%. Lastly, in some cases we determined that the true label was a classification other than those provided by the annotators or the model, which account for 9% of errors.

Model errors were further subdivided, as shown on the tree diagram branches. *Hate when None* includes instances where the model classified “None” content as “Hateful”. This was likely caused by an over reliance on the use of slurs (i.e., a slur being used non-hatefully), misidentification of in-

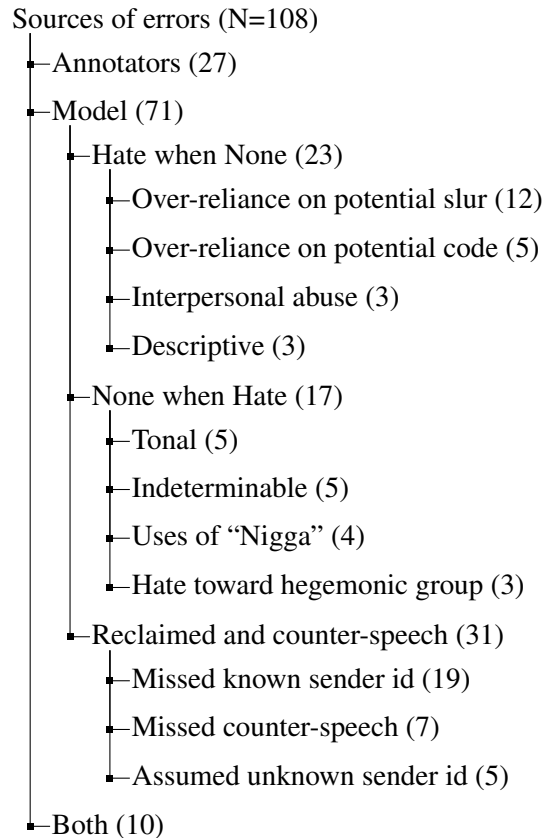


Figure 1: Sources of classification errors for the multimodal model.

terpersonal abuse or confusion caused by content which describes/reports on (but does not endorse) hateful activities. *None when Hate* is the inverse. It includes all instances where the model classified “Hateful” content as “None”. This was likely caused by a combination of tone misjudgements (e.g., hateful language which had positive sentiment), uses of the word *Nigga*⁸, hate towards a

⁸When the identity of the speaker was unknown and the use was not evidently “Hateful”, annotators were instructed to treat the tweet as “None” to avoid penalizing language most

hegemonic group (e.g., referring to White people as *Hillbillies*), and, in select cases, indeterminable reasons. The final category, *Reclaimed and Counterspeech*, comprises confusion beyond the “None”–“Hateful” distinction. These were caused by missing the sender’s identity when it was knowable (i.e., “Reclaimed” language was classified as “None”), assuming the sender’s identity when it was unknowable (i.e., “None” content was classified as “Reclaimed”, and failing to predict “Counterspeech”). This final type did not exhibit systematic misclassifications, likely because there were few instances overall.

8 Discussion

As hypothesized, the unimodal-text context-aware model outperformed the unimodal-text context-invariant model. The 3.5 percentage point increase in F1 scores when distinguishing between *Primary* labels (“Hate”, “Counterspeech”, “Reclaimed”, and “None”) is in line with other results which range from 0–8 percentage point improvements depending on the dataset (Vidgen et al., 2019; Sohn and Lee, 2019; Zampieri et al., 2019; Mozafari et al., 2019).

When isolating performance by hate strategy, this gap substantially widens to 12.2 (“Explicit”) and 14.3 (“Implicit”) percentage points. This demonstrates the value of semantic contextualization for accurately identifying hate speech grows as its overtness diminishes. This is not the case with non-hateful content. On those, the unimodal-text context-invariant model is 3.5 percentage points more accurate.

Content modalities interact in ways that can create more subtle forms of hate. For this reason, the multimodal model was expected to outperform unimodal ones. This is marginally supported by the 1.7 percentage point gap in F1 score on the *Primary* labels between the multimodal model and the best performing unimodal model. This improvement is larger than that achieved by Gomez et al. (2020); Sabat et al. (2019); Yang et al. (2019), but smaller than that by Kiela et al. (2020) in similar hate speech detection tasks. Improvements from the incorporation of image data are minor compared to the gains from the transformer model. This may either signal underlying patterns in how multimodality is used in online content or reflect the

commonly used by Black communicators. Ergo, the model may associate its use with “None”.

data collection methods which relied on text-based query methods through the Twitter API.

However, it is the context-aware unimodal-text model that performs the best when assessing by *Strategy*. The gaps between the unimodal-text model and multimodal model are 2.8 percentage points (“Non-Hateful”), 3.1 percentage points (“Explicit”), and 4.2 percentage points (“Implicit”). The models which only consider image data are more attuned to “Non-Hateful” tweets to the determinant of hate identification. This implies a dissonance between modalities which mitigates, sometimes rightly and other times wrongly, hateful signals from the text.

All models more accurately identify implicit than explicit hate. This is a surprising result which may reflect their higher representation within the training data rather than any inherent property that makes them more detectable.

9 Conclusion

This paper substantiates the need to consider varying forms of hate with different modalities and levels of overtness. It investigates the value of context-aware textual and multimodal features finding that both improve model F1 score with the multimodal model performing the best (0.771). Further, we find that model performance is directly contingent upon annotator agreement levels (referred to as ‘ambiguity’ in the main body of the paper). These findings are generated from a newly-annotated dataset of 5,000 tweets containing information on each entry’s primary attribute, modality, and strategy. This dataset along with the annotation codebook, model training code, and model weights are available to encourage future research on the topic.

Acknowledgments

This work was supported by Wave 1 of The UKRI Strategic Priorities Fund under the EPSRC Grant EP/T001569/1, particularly the “Criminal Justice System” theme within that grant, and by The Alan Turing Institute. We are grateful to all our annotators and appreciate the feedback from our reviewers as well as that from the 2020 cohort of the MSc in Social Data Science at the University of Oxford.

References

- Betty van Aken, Julian Risch, Ralf Krestel, and Alexander Löser. 2018. [Challenges for Toxic Comment Classification: An In-Depth Error Analysis](#). *arXiv*, pages 33–42.
- Bethany L Albertson. 2015. [Dog-Whistle Politics: Multivocal Communication and Religious Appeals](#). Technical Report 1.
- Anti-Defamation League. 2021. [Online Hate and Harassment: The American Experience 2021](#). Technical report.
- Pinkesh Badjatiya, Shashank Gupta, Manish Gupta, and Vasudeva Varma. 2019. [Deep learning for hate speech detection in tweets](#). In *26th International World Wide Web Conference 2017, WWW 2017 Companion*, pages 759–760. International World Wide Web Conferences Steering Committee.
- Philipp Blandfort, Desmond U. Patton, William R. Frey, Svebor Karaman, Surabhi Bhargava, Fei-Tzin Lee, Siddharth Varia, Chris Kedzie, Michael B. Gaskell, Rossano Schifanella, Kathleen McKeown, and Shih-Fu Chang. 2019. [Multimodal Social Media Analysis for Gang Violence Prevention](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 13:114–124.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. [Enriching Word Vectors with Subword Information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- danah boyd. 2017. [Social Network Sites: Public, Private, or What?](#) *Knowledge Tree*.
- Tommaso Caselli, Valerio Basile, Jelena Mitrović, Inga Kartoziya, and Michael Granitzer. 2020. [I Feel Offended, Don't Be Abusive! Implicit/Explicit Messages in Offensive and Abusive Language](#). In *Proceedings of LREC*.
- Tommaso Caselli, Nicole Novielli, Viviana Patti, and Paolo Rosso. 2018. [EVALITA Evaluation of NLP and Speech Tools for Italian](#). *Proceedings of the Final Workshop*, 12:235–238.
- Stevie Chancellor, Yannis Kalantidis, Jessica A. Pater, Munmun De Choudhury, and David A. Shamma. 2017. [Multimodal classification of moderated online pro-eating disorder content](#). In *Conference on Human Factors in Computing Systems - Proceedings*, volume 2017-May, pages 3213–3226. Association for Computing Machinery.
- Eshwar Chandrasekharan, Umashanthi Pavalanathan, Anirudh Srinivasan, Adam Glynn, Jacob Eisenstein, and Eric Gilbert. 2017. [You can't stay here: The efficacy of Reddit's 2015 ban examined through hate speech](#). *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):1–22.
- Francois Chollet. 2017. [Xception: Deep learning with depthwise separable convolutions](#). *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Danielle Keats Citron and Helen L. Norton. 2011. [Intermediaries and hate speech: Fostering digital citizenship for our information age](#).
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#).
- Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. 2019. [Racial Bias in Hate Speech and Abusive Language Detection Datasets](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 25–35, Florence. Association for Computational Linguistics.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. [Automated Hate Speech Detection and the Problem of Offensive Language](#). *Proceedings of the 11th International Conference on Web and Social Media, ICWSM 2017*, pages 512–515.
- Richard Dvorak. 1999. [Cracking the Code: De-Coding Colorblind Slurs during the Congressional Crack Cocaine Debates](#). *Michigan Journal of Race & Law*, 5.
- Facebook. 2020. [Community Standards](#).
- Elise Fehn Unsvåg and Björn Gambäck. 2018. [The Effects of User Features on Twitter Hate Speech Detection](#). In *Proceedings of the Second Workshop on Abusive Language Online*, pages 75–85, Brussels.
- Paula Fortuna and Sérgio Nunes. 2018. [A survey on automatic detection of hate speech in text](#).
- Adam D Galinsky, Cynthia S Wang, Jennifer A Whitson, Eric M Anicich, Kurt Hugenberg, and Galen V Bodenhausen. 2013. [The Reappropriation of Stigmatizing Labels: The Reciprocal Relationship Between Power and Self-Labeling](#). *Psychological Science*, 24(10):2020–2029.
- Björn Gambäck and Utpal Kumar Sikdar. 2017. [Using Convolutional Neural Networks to Classify Hate-Speech](#). In *Proceedings of the First Workshop on Abusive Language Online*, pages 85–90.
- Lei Gao and Ruihong Huang. 2017. [Detecting Online Hate Speech Using Context Aware Models](#). In *International Conference Recent Advances in Natural Language Processing, RANLP*, pages 260–266. Association for Computational Linguistics (ACL).
- Raul Gomez, Jaume Gibert, Lluís Gomez, and Dimosthenis Karatzas. 2020. [Exploring Hate Speech Detection in Multimodal Publications](#). Technical report.

- Alex Graves and Jürgen Schmidhuber. 2005. [Frame-wise phoneme classification with bidirectional LSTM and other neural network architectures](#). In *Neural Networks*, volume 18, pages 602–610. Pergamon.
- Rosanna E Guadagno, Daniel M Rempala, Shannon Murphy, and Bradley M Okdie. 2013. [What makes a video go viral? An analysis of emotional contagion and Internet memes](#).
- Jeremy Howard and Sebastian Ruder. 2018. [Fine-tuned language models for text classification](#). *CoRR*, abs/1801.06146.
- Matthew W Hughey and Jessie Daniels. 2013. [Racist comments at online news sites: a methodological dilemma for discourse analysis](#). *Culture & Society*, 35(3):332–347.
- C J Hutto and Eric Gilbert. 2014. [VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text](#). In *The International AAAI Conference on Web and Social Media (ICWSM)*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. [Bag of Tricks for Efficient Text Classification](#). *15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017 - Proceedings of Conference*, 2:427–431.
- Brendan Kennedy, Xisen Jin, Aida Mostafazadeh Davani, Morteza Dehghani, and Xiang Ren. 2020. [Contextualizing hate speech classifiers with post-hoc explanation](#).
- Douwe Kiela, Suvrat Bhooshan, Hamed Firooz, and Davide Testuggine. 2019. [Supervised multimodal bitransformers for classifying images and text](#).
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine Facebook. 2020. [The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes](#). *ArXiv*, pages 1–17.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. [Albert: A lite bert for self-supervised learning of language representations](#).
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. [Visualbert: A simple and performant baseline for vision and language](#).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#).
- Ilya Loshchilov and Frank Hutter. 2017. [Decoupled weight decay regularization](#).
- Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 2020. [12-in-1: Multi-task vision and language representation learning](#). In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Rijul Magu, Kshitij Joshi, and Jiebo Luo. 2017. [Detecting the Hate Code on Social Media](#). Technical report.
- Alice E Marwick and danah boyd. 2010. [Article I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience](#). *New Media & Society*, 13(1):111–133.
- Priscilla Marie Meddaugh and Jack Kay. 2009. [Hate Speech or “Reasonable Racism?” The Other in Stormfront](#). *Journal of Mass Media Ethics*, 24(4):251–268.
- Ryan M Milner. 2013. [FCJ-156 Hacking the Social: Internet Memes, Identity Antagonism, and the Logic of Lulz](#). *The Fibreculture Journal*, (22):62–92.
- Marzieh Mozafari, Reza Farahbakhsh, and Noel Crespi. 2019. [A BERT-Based Transfer Learning Approach for Hate Speech Detection in Online Social Media](#). *Studies in Computational Intelligence*, 881 SCI:928–940.
- Marzieh Mozafari, Reza Farahbakhsh, and Noël Crespi. 2020. [A BERT-Based Transfer Learning Approach for Hate Speech Detection in Online Social Media](#). In *Studies in Computational Intelligence*, volume 881 SCI, pages 928–940. Springer.
- Karsten Müller, Carlo Schwarz, Sascha Becker, Christopher Blattman, Leonardo Bursztyjn, Mirko Draca, Thimo Fetzer, Evan Fradkin, Matthew Gentzkow, Andy Guess, Vardges Levonyan, Atif Mian, Magne Mogstad, Sharun Mukand, Hans-Joachim Voth, Fabian Waldinger, and Noam Yuchtman. 2019. [Fanning the Flames of Hate: Social Media and Hate Crime *](#). Technical report.
- Patrick B O’Sullivan and Caleb T Carr. 2018. [Masspersonal communication: A model bridging the mass-interpersonal divide](#). *New Media & Society*, 20(3):1161–1180.
- Jérémie Pelletier-Gagnon and Axel Pérez Trujillo Diniz. 2018. [Colonizing Pepe: Internet Memes as Cyberplaces](#). *Space and Culture*, page 120633121877618.
- T. F. Pettigrew and R. W. Meertens. 1995. [Subtle and blatant prejudice in western Europe](#). *European Journal of Social Psychology*, 25(1):57–75.
- Whitney M. Phillips. 2012. [This is why we can’t have nice things: The origins, evolution and cultural embeddedness of online trolling](#). Ph.D. thesis, University of Oregon.

- Georgios K. Pitsilis, Heri Ramampiaro, and Helge Langseth. 2018. [Effective hate-speech detection in Twitter data using recurrent neural networks](#). *Applied Intelligence*, 48(12):4730–4742.
- Marco Polignano, Pierpaolo Basile, Marco De Gemmis, and Giovanni Semeraro. 2019. [Hate Speech Detection through AIBERTo Italian Language Understanding Model](#). Technical report.
- Scott Poynting and Greg Noble. 2003. [‘Dog-Whistle’ Journalism and Muslim Australians since 2001](#). *Media International Australia incorporating Culture and Policy*, 109(1):41–49.
- Jing Qian, Mai Elshierief, Elizabeth M Belding, and William Yang Wang. 2018. [Leveraging Intra-User and Inter-User Representation Learning for Automated Hate Speech Detection](#). *arXiv*, pages 1–6.
- Alison P Ribeiro and Nádia F F Da Silva. 2019. [INF-HatEval at SemEval-2019 Task 5: Convolutional Neural Networks for Hate Speech Detection Against Women and Immigrants on Twitter](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval-2019)*, pages 420–425.
- Benet Oriol Sabat, Cristian Canton Ferrer, and Xavier Giro-i Nieto. 2019. [Hate Speech in Pixels: Detection of Offensive Memes towards Automatic Moderation](#). In *AI for Social Good workshop at NeurIPS*, Vancouver, CA.
- Manuela Sanguinetti, Fabio Poletto, Cristina Bosco, Viviana Patti, and Marco Stranisci. 2018. [An Italian Twitter corpus of hate speech against immigrants](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, Noah A Smith, and Paul G Allen. 2019. [The Risk of Racial Bias in Hate Speech Detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence. Association for Computational Linguistics.
- Anna Schmidt and Michael Wiegand. 2017. [A Survey on Hate Speech Detection using Natural Language Processing](#). In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 1–10, Valencia.
- Mike Schuster and Kuldip K Paliwal. 1997. [Bidirectional Recurrent Neural Networks](#). Technical Report 11.
- Hajung Sohn and Hyunju Lee. 2019. [MC-BERT4HATE: Hate speech detection using multi-channel bert for different languages and translations](#). In *IEEE International Conference on Data Mining Workshops, ICDMW*, volume 2019-November, pages 551–559. IEEE Computer Society.
- Byung Soo Ko. 2020. [ImageNet Classification Leaderboard](#).
- Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alex Alemi. 2016a. [Inception-v4, inception-resnet and the impact of residual connections on learning](#).
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016b. [Rethinking the Inception Architecture for Computer Vision](#). In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2016-December, pages 2818–2826. IEEE Computer Society.
- Twitter. 2020. [The Twitter Rules](#).
- Bertie Vidgen, Austin Botelho, David Broniatowski, Ella Guest, Matthew Hall, Helen Margetts, Rebekah Tromble, Zeerak Waseem, and Scott Hale. 2020. [Detecting east asian prejudice on social media](#).
- Bertie Vidgen, Helen Margetts, and Alex Harris. 2019. [How much online abuse is there? A systematic review of evidence for the UK Policy Briefing-Full Report Public Policy Programme Hate Speech: Measures and Counter Measures](#). Technical report, The Alan Turing Insitute.
- Bertie Vidgen, Dong Nguyen, Helen Margetts, Patricia Rossini, and Rebekah Tromble. 2021. [Introducing CAD: the contextual abuse dataset](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2289–2303, Online. Association for Computational Linguistics.
- Prashanth Vijayaraghavan, Hugo Larochelle, and Deb Roy. 2019. [Interpretable Multi-Modal Hate Speech Detection](#). In *International Conference on Machine Learning AI for Social Good Workshop*, pages 1–5, Long Beach.
- Joseph B. Walther, Caleb T. Carr, Scott Seung W. Choi, David C. DeAndrea, Jinsuk Kim, Stephanie Tom Tong, and Brandon Van Der Heide. 2011. [Interaction of Interpersonal, Peer, and Media Influence Sources Online: A Research Agenda for Technology Convergence](#). In Zizi Papacharissi, editor, *A Networked Self: Identity, Community, and Culture on Social Network Sites*, chapter 1, pages 17–38. Routledge, New York.
- Zijian Wang, Scott Hale, David Ifeoluwa Adelani, Przemyslaw Grabowicz, Timo Hartman, Fabian Flöck, and David Jurgens. 2019. [Demographic inference and representative population estimates from multilingual social media data](#). In *The World Wide Web Conference, WWW ’19*, page 2056–2067, New York, NY, USA. Association for Computing Machinery.
- Zeerak Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. [Understanding abuse: A typology of abusive language detection subtasks](#). In

- Proceedings of the First Workshop on Abusive Language Online*, pages 78–84, Vancouver, BC, Canada. Association for Computational Linguistics.
- Zeerak Waseem and Dirk Hovy. 2016. [Hateful symbols or hateful people? predictive features for hate speech detection on twitter](#). In *Proceedings of the NAACL Student Research Workshop*, pages 88–93, San Diego, California. Association for Computational Linguistics.
- Michael Wiegand, Josef Ruppenhofer, and Thomas Kleinbauer. 2019. [Detection of abusive language: the problem of biased datasets](#). The 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Proceedings of the Conference Vol. 1. Minneapolis, Minnesota, June 2 - June 7, 2019, pages 602–608, Stroudsburg, PA, USA. The Association for Computational Linguistics.
- Ellery Wulczyn, Nithum Thain, and Lucas Dixon. 2017. [Ex machina: Personal attacks seen at scale](#). In *26th International World Wide Web Conference, WWW 2017*, pages 1391–1399. International World Wide Web Conferences Steering Committee.
- Fan Yang, Xiaochang Peng, Gargi Ghosh, Reshef Shilon, Hao Ma, Eider Moore, and Goran Predovic. 2019. [Exploring Deep Multimodal Fusion of Text and Photo for Hate Speech Classification](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 11–18, Florence. Association for Computational Linguistics.
- YouTube. 2020. [Hate speech policy - YouTube Help](#).
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. [SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media \(OffensEval\)](#). *arXiv*, pages 1–12.
- Savvas Zannettou, Tristan Cauleld, Jeremy Blackburn, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Guillermo Suarez-Tangil. 2018. [On the origins of memes by means of fringe web communities](#). In *Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC*, pages 188–202. Association for Computing Machinery.
- Ziqi Zhang, David Robinson, and Jonathan Tepper. 2018. [Detecting Hate Speech on Twitter Using a Convolution-GRU Based Deep Neural Network](#). In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 10843 LNCS, pages 745–760. Springer Verlag.
- Haoti Zhong, Hao Li, Anna Squicciarini, Sarah Rajtmajer, Christopher Griffin, David Miller, and Cornelia Caragea. 2016. [Content-Driven Detection of Cyberbullying on the Instagram Social Network](#). In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 3952–3958, New York. International Joint Conferences on Artificial Intelligence.
- Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. 2018. [Learning transferable architectures for scalable image recognition](#). *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*.