

VAuLT: AUGMENTING THE VISION-AND-LANGUAGE TRANSFORMER FOR SENTIMENT CLASSIFICATION ON SOCIAL MEDIA

Georgios Chochlakis¹ Tejas Srinivasan² Jesse Thomason² Shrikanth Narayanan¹

¹ Signal Analysis and Interpretation Lab, University of Southern California, Los Angeles, CA 90089, USA

² GLAMOR Lab, University of Southern California, Los Angeles, CA 90089, USA

ABSTRACT

We propose the **Vision-and-Augmented-Language Transformer (VAuLT)**. VAuLT is an extension of the popular Vision-and-Language Transformer (ViLT), and improves performance on vision-and-language (VL) tasks that involve more complex text inputs than image captions while having minimal impact on training and inference efficiency. ViLT, importantly, enables efficient training and inference in VL tasks, achieved by encoding images using a linear projection of patches instead of an object detector. However, it is pretrained on captioning datasets, where the language input is simple, literal, and descriptive, therefore lacking linguistic diversity. So, when working with multimedia data in the wild, such as multimodal social media data, there is a notable shift from captioning language data, as well as diversity of tasks. We indeed find evidence that the language capacity of ViLT is lacking. The key insight and novelty of VAuLT is to propagate the output representations of a large language model (LM) like BERT to the language input of ViLT. We show that joint training of the LM and ViLT can yield relative improvements up to 20% over ViLT and achieve state-of-the-art or comparable performance on VL tasks involving richer language inputs and affective constructs, such as for Target-Oriented Sentiment Classification in TWITTER-2015 and TWITTER-2017, and Sentiment Classification in MVSA-Single and MVSA-Multiple. Our code is available at <https://github.com/gchochla/VAuLT>.

Index Terms— Vision-and-Language, Transformers, Social media, Sentiment Classification

1. INTRODUCTION

The study of social media can have positive societal outcomes. Interest in the importance of emotion has been growing beyond science [8] to law [27], and politics [26]. Social media users readily express their emotions about their experiences and world events through multimedia content such as videos, images, memes, and raw text. This allows researchers to use machine learning (ML) methods to study, for instance, the dynamics of public perception and opinion and to provide feedback to policy makers and inform future actions [12]. Social media are also ripe with opportunities for social manipulation [4]. This manipulation can be personalized based on user data [1], or use network effects [11]. Mitigating these risks and identifying widespread manipulation attempts hence becomes pivotal.

Due to the massive volume of data and the dynamic nature of social media, constant streams of data are required, thus efficiency is of the essence. In addition, multimedia content necessitates multimodal modeling, as users can use these extra degrees of freedom to convey complex and conflicting meanings, for instance in irony. Moreover, a multimodal model should be able to handle a subset of its input modalities if a post is missing the rest [24].

Many state-of-the-art vision-and-language models (VLM), e.g., VL-BERT [25] and ViLBERT [17], contain an object detector, such

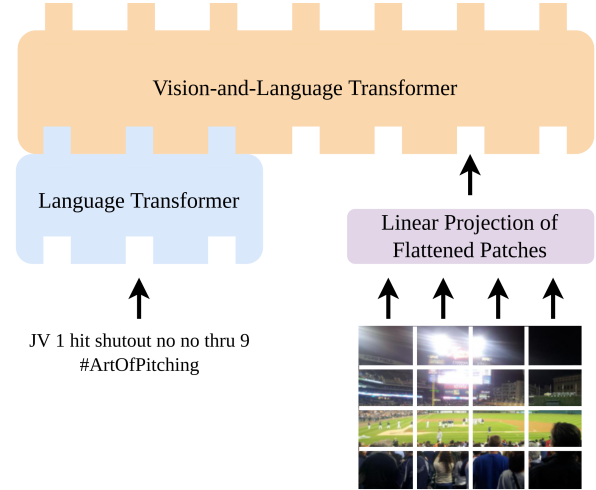


Fig. 1. VAuLT propagates representations from a LM to a VLM in a modular fashion. We show that even with little tuning data, language representations can be effectively used to significantly improve the performance of ViLT on tasks with greater linguistic and construct diversity than captioning datasets.

as Faster-RCNN [22], to compute input visual features. Object detectors are costly to run relative to the rest of the pipeline [15], which can be prohibitive for social media applications. Hence, we turn to the Vision-and-Language Transformer (ViLT) [15], whose visual input consists of linearly projected image patches.

Regardless, these models, designed to understand and extract information from multimedia content such as vision and language, are frequently trained *only* on images and their captions crawled from the web. Image captions are a weak form of language supervision, often describing literal content with simple syntax and little subtlety, deviating substantially from the text of social media posts.

In this paper, we propose the Vision-and-Augmented-Language Transformer (VAuLT), which directly addresses the impoverished language representations of ViLT by processing the language input through a pretrained large language model (LM) (Figure 1). We find that VAuLT is able to outperform ViLT on a variety of tasks involving affective constructs common in social media, including TWITTER-2015, TWITTER-2017 [28], and MVSA [21].

Our key contributions can be summarized as follows:

- In VAuLT, we stack Transformer architectures, a pretrained LM and VLM, and jointly train them end-to-end, yielding substantial performance gains from little tuning data.
- We show that VAuLT competes with or outperforms state-of-

the-art methods for the social media analysis tasks examined.

- We provide evidence that the performance trade-offs of ViLT and VAuLT depend on the complexity of the image and language components commensurate with the relative richness of, and variety within, the modalities in the source and target domains, as well as the tasks themselves.

2. RELATED WORK

2.1. Language Transformers

Transformers pretrained on general-purpose text grants them the ability to be used as backbones for many downstream tasks [6]. Such models can also be pretrained on more specialized corpora to improve their transferability to specific domains, like BERTweet [20], a model trained on tweets.

2.2. Vision-and-Language (VL) Transformers

Similar to language transformers, VL transformers [25, 17] are trained with multiple input modalities on web-crawled captions, and yield state-of-the-art results on many multimodal tasks [2]. These models frequently use region-of-interest (ROI) features extracted from a pretrained object detector [22] as their Transformer’s visual inputs. However, this visual embedding step is computationally expensive during training and inference [15], requiring caching of features. ViLT [15], based on the Vision Transformer (ViT) [7], proposed using image patches as visual input to the Transformer model instead. This results in slightly worse model performance, but is much more computationally efficient. However, ViLT also has limited language understanding capacity due to being initialized from ViT and pretrained solely on image caption data. VAuLT overcomes this issue by replacing ViLT’s language embedding inputs with features extracted from a large LM pretrained on greater linguistic diversity, which can also be selected to more closely suit the needs of the downstream task.

2.3. Affective Analysis on Social Media

Researchers tend to focus on single modalities when analyzing emotion and sentiment in social media posts, such as text-only [19], and image-only [23] datasets.

Emotion recognition from text has long relied on word-counting techniques [9], which can be deployed at scale but disregard context. LSTMs [3] and Transformers [5] capture context and provide superior performance with less interpretability.

On the other hand, multimodal affect modeling tasks with social media have predominantly focused on sentiment [28, 21, 13]. [28] have annotated tweets with their corresponding images for sentiment expressed towards entities in the text, resulting in the *Target-oriented Sentiment Classification* (TMSC) task and the TWITTER-15 and TWITTER-17 datasets. Similarly, MVSA [21] has annotations for sentiment expressed in the text and the image of a tweet. User-provided tags have been used as “self-reports” of emotion in scraped multimodal data from Tumblr [13].

The proposed approaches for the aforementioned tasks typically involve dedicated architectures or targeted modifications of VL Transformers. An example of the former is TomBERT [28] for TMSC. TomBERT consists of multiple components: first, a ResNet [10] extracts visual features. Then, a cross-attention layer is used to query these visual features using target embeddings extracted from BERT [6]. Parallel to the visual feature extraction, the text sequence

is passed through another BERT, and is concatenated with the refined visual features and fed to a multimodal encoder, also a BERT.

3. METHODOLOGY

3.1. Task Definitions

3.1.1. Sentiment Classification

For our purposes, “Sentiment Classification” corresponds to classifying an image and text pair as *positive*, *neutral*, or *negative*.

3.1.2. Target-oriented Sentiment Classification (TMSC)

TMSC extends the basic sentiment classification formulation to opinion expressed toward *given* targets within the text of multimodal tweets. In our setting, sentiment targets were repurposed from types of named entities, and an input tweet can contain multiple targets.

3.2. Vision-and-Language Transformer (ViLT)

ViLT is a multimodal transformer. The language component of ViLT follows conventional practices and embeds input tokens with a lookup table. On the other hand, the visual component of ViLT follows ViT [7], segmenting each image into non-overlapping square patches, flattening the resulting values, and mapping them to the input space using a linear projection.

ViLT is pretrained on image captioning datasets like MSCOCO [16]. Its pretraining objectives include: image-text matching and word-patch alignment, and masked language modeling and whole word masking. The former group of objectives focuses on enforcing joint modeling of the two modalities, while the latter is centered on the language understanding of the model. Explicit language training is necessary since the model is initialized from ViT weights, itself trained on images, in contrast to previous works that initialize weights from LMs [25, 17], thus *ViLT’s language understanding comes solely from image-captioning datasets*.

3.3. Vision-and-Augmented Language Transformer (VAuLT)

The proposed VAuLT can use any LM that produces token embeddings for tokenized text, irrespective of the specific tokenizer. We have this requirement because we use the sequence of embeddings produced by the LM to replace the canonical language input of ViLT. Therefore, we are essentially substituting the context-free representations from the lookup table of ViLT, with context-aware representations [6] from the output layer of a LM. Images are embedded in the canonical fashion with ViLT’s trained linear projection of patches.

Concretely, given an input image-text pair, and a specific LM and its tokenizer, we tokenize the input text sequence and pass this through the LM to produce its output contextual embeddings. Along with these embeddings, we pass these two input streams through ViLT. Note that position embeddings, etc., are added in both ViLT’s and the LM’s forward pass.

We use LMs like BERT and BERTweet that have not been trained on similar datasets with ViLT or tokenize text in the same manner as ViLT does. That is to say, the language tokens might not correspond to the tokens ViLT has been trained with.

4. EXPERIMENTS

Overall, our goal is to show that VAuLT improves upon ViLT on multimodal, small-scale, affective datasets derived from social media, in our case Twitter. More than that, we also show that ViLT

Model	TWITTER-15		TWITTER-17		MVSA-M		MVSA-S	
	Acc	mac-F1	Acc	mac-F1	Acc	w-F1	Acc	w-F1
ViLT	70.5 \pm 1.3	62.6 \pm 2.5	62.6 \pm 0.1	58.1 \pm 0.7	69.1 \pm 0.8	67.0 \pm 0.3	74.4 \pm 1.5	73.7 \pm 1.6
VAuLT [†]	75.6 \pm 0.8	70.0 \pm 1.7	70.2 \pm 0.4	67.8 \pm 0.1	70.0 \pm 0.7	68.9 \pm 0.7	78.0\pm1.2	77.4 \pm 1.4
VAuLT*	77.5\pm0.4	72.9\pm0.5	71.0 \pm 0.5	69.5\pm0.7	70.3 \pm 1.6	67.0 \pm 2.9	72.8 \pm 5.6	71.8 \pm 5.8
EF-CaTr-BERT [14]	77.9\pm0.8	73.9\pm0.8	72.3\pm0.3	70.2\pm0.2	-	-	-	-
ITIN [29]	-	-	-	-	73.5	-	75.2	-

Table 1. Direct comparison between ViLT, VAuLT ([†]: w/ BERT, bert-base-uncased, or *: w/ BERTweet, vinai/bertweet-base), and the state-of-the-art on each benchmark. **Bold** indicates best performance based on ranges, where available.

can achieve competitive performance over dedicated architectures for such datasets, and VAuLT can even exceed their performance.

4.1. Datasets

We use TWITTER-15 and TWITTER-17 [28] annotated for the task of TMSF. The former contains 1548 positive, 630 negative and 3169 neutral examples. The latter contains 2516 positive, 728 negative and 2728 neutral examples. We use the splits provided by the authors. Accuracy and macro F1 are typical evaluation criteria. We predict for one target at a time, even if the tweet contains multiple. We follow [28] in replacing the target of interest with the placeholder \$T\$, and appending the target as a second sequence.

We also evaluate our model on MVSA-Multiple and MVSA-Single [21], pre-processed for Sentiment Classification. Initially, both datasets contain annotations for the sentiment of both the image and the text of a tweet. MVSA-Single has annotations from solely one annotator, while MVSA-Multiple utilizes three separate annotators, plus more examples. We merge the annotations into a single annotation and aggregate across annotators, as is customary [29]. The filtered MVSA-Single contains 2683 positive, 460 negative and 1358 neutral examples. The filtered MVSA-Multiple contains 11318 positive, 4408 negative and 1298 neutral examples. We also had to remove 3 additional samples from MVSA-Multiple because of corrupted images (example IDs: 3151, 3910 and 5995). Since canonical splits are not provided and evaluation criteria are not well-defined, we follow previous work in randomly splitting the data 8:1:1 [29] and use accuracy to compare. We also include the weighted F1 score of our methods as our best guess as to what other methods are reporting as merely “F1 score”.

4.2. Implementation details

We use Python v3.7.4, PyTorch v1.11.0, and *transformers* v4.19.2.

We keep our learning rate at $2 \cdot 10^{-5}$ unless specified. We use linear warm-up for 10% of the training steps and then linear decay to 0. In re-implementations of published models, we use the specified hyperparameter configuration.

Test performance (mean and standard deviation) is reported after 3 different training runs with fixed hyperparameters. The main tuning of all of the hyperparameters but the number of training epochs was performed on TWITTER-15 and TWITTER-17. We searched for whether to use bias correction in AdamW, selecting no bias correction, the number of epochs (ran for 15 epochs and picked the best performing one based on average accuracy and macro F1 for MVSA-Multiple and MVSA-Single, and {8, 15, 25} for TWITTER-15 and TWITTER-17), and whether to integrate the placeholder \$T\$ to the tokenizer as a standalone token in TWITTER-15 and TWITTER-17 (yes for VAuLT; no for ViLT, existing baselines, and variants

thereof).

We find that divergence in the training of VAuLT can rarely occur, but those runs can be filtered by observing degradation in the training metrics. While we had to correct for such artifacts during development, no such instances occurred on the test sets.

We use simple random cropping instead of the augmentation utilized in ViLT for simplicity. We use augmentation in all cases where ViLT’s linear projection is solely used to extract image embeddings. We use *ekphrasis*¹ to pre-process the text of the tweets, but we remove the special characters added by the library around the names of the tags (e.g., when “<user>” would have been the original substitute for the actual Twitter handle appearing in a tweet, we simply use “user”). We also substitute emojis for their plain-text descriptions in parentheses, when they would be dropped by a tokenizer. In particular, in our work, only BERTweet’s tokenization preserves emojis. Because the text of TWITTER-15 and TWITTER-17 has been independently pre-processed, we do not use *ekphrasis* on them.

We use 2 NVIDIA GeForce GTX 1080 Ti GPUs (12GB of memory). We only use one GPU per model, limiting our training batch sizes to 16 or 32, always choosing the maximum possible of the two.

4.3. Comparison with ViLT

First and foremost, let us focus on the comparison of the evaluation metrics between ViLT and VAuLT, where we indeed find that VAuLT exceeds ViLT’s performance across all examined affective social media data. Results can be seen in Table 1. We can see that VAuLT with BERTweet improves upon ViLT in TWITTER-15 and TWITTER-17, with relative improvements between 9.9% and 19.6%. VAuLT with BERT also provides substantial improvements but still lags behind compared to BERTweet, and only we find overlapping ranges across the two versions of VAuLT on TWITTER-17’s accuracy.

Smaller yet still substantial improvements are observed in MVSA-Multiple and MVSA-Single. In particular, for this set of benchmarks, VAuLT with BERT performs favorably compared to ViLT and VAuLT with BERTweet. In fact, for the latter, we observe large fluctuations in performance, emphasized by the large deviation presented in Table 1. Otherwise, performance is analogous to VAuLT’s with BERT. Quantitatively, improvements over ViLT are up to approximately 5%. It has to be noted that ranges are overlapping for VAuLT and ViLT in MVSA-Multiple accuracy.

4.4. Comparisons with State of the Art

For completeness, we also compare VAuLT and ViLT with state-of-the-art models for the examined benchmarks. We do not elaborate on the state-of-the-art models because of the variety of tasks, as well as them being not central to the main goals of our work. We find

¹<https://github.com/cbaziotis/ekphrasis>

Model	TWITTER-15		TWITTER-17	
	Acc	mac-F1	Acc	mac-F1
ViLT	70.5 \pm 1.3	62.6 \pm 2.5	62.6 \pm 0.1	58.1 \pm 0.7
VAuLT [†]	75.6 \pm 0.8	70.0 \pm 1.7	70.2 \pm 0.4	67.8 \pm 0.1
VAuLT ^{*,f}	68.8 \pm 0.6	58.8 \pm 1.5	64.0 \pm 0.9	61.0 \pm 0.6
VAuLT [*]	77.5 \pm 0.4	72.9 \pm 0.5	71.0 \pm 0.5	69.5 \pm 0.7

Table 2. Comparing ViLT and VAuLT (^f: using a frozen LM, and [†]: w/ BERT, bert-base-uncased, or ^{*}: w/ BERTweet, vinai/bertweet-base) on TWITTER-15 and TWITTER-17. **Bold** indicates best based on ranges.

that VAuLT can outperform or be competitive with dedicated state-of-the-art models in terms of our evaluation criteria.

We observe, in Table 1, that VAuLT with BERTweet has overlapping performance ranges with the state-of-the-art model, EF-CaTriBERT, on TWITTER-15 and TWITTER-17, except for accuracy in TWITTER-17.

For MVSA-Multiple and MVSA-Single, we have mixed results. We achieve state-of-the-art results for MVSA-Single, where we improve upon the previous best accuracy by 3.8 absolute percentage points (note that our best-guess F1 score is also the state of the art if our assumption is correct). In MVSA-Multiple, our models are still outperformed by the competition. Caveats in the evaluation in MVSA include the random splits that are used in the literature, often with different cardinalities in the different splits, while evaluation criteria remain opaque, and pre-processing, aggregation and normalization techniques could differ between different works.

4.5. Ablation Studies

4.5.1. Frozen LM

We first examine whether ViLT can “read” the outputs of a LM without any adjustments on the part of the latter. That is to say, we keep the LMs frozen. [18] showed that early layers of transformers change significantly less than later layers during fine-tuning, implying that fine-tuning of the LM in VAuLT may be unnecessary. However, we find joint training to be essential.

Results can be seen in Table 2. Overall, fine-tuning of the language model appears essential, since we see such notable degradation in performance, as VAuLT with a frozen LM drops to the level or even below ViLT. This demonstrates the necessity of fine-tuning the LMs in conjunction with ViLT.

4.5.2. Deep Vision vs. Deep Language

In this section, we present an experiment supporting the hypothesis that the complexity of the modality-specific components and their effect on the final performance depends on the domain addressed, and that the linear projections used by ViLT and ViT do not inherently restrict the model’s performance. We do so by introducing and studying two variants of TomBERT we call TomVAuLT and TomViLT, solely for the purposes of examining our hypothesis in this section.

TomVAuLT extends TomBERT by replacing the multimodal encoder with ViLT. The multimodal encoder in TomBERT is responsible for the final classification, and its input consists of the queried features from ResNet and the contextual embeddings of the tweet from BERT. In this way, we replace the visual embeddings of ViLT with ResNet features, while we also retain the BERT that pre-processes the tweet before the multimodal encoder, giving us

Model	TWITTER-15		TWITTER-17	
	Acc	mac-F1	Acc	mac-F1
TomViLT	73.2 \pm 0.7	66.1 \pm 1.3	67.3 \pm 0.5	64.4 \pm 0.7
ViLT	69.6 \pm 1.0	60.3 \pm 1.5	64.0 \pm 1.2	58.0 \pm 1.5
TomVAuLT	73.2 \pm 0.5	67.1 \pm 0.4	67.3 \pm 0.9	61.1 \pm 0.9
VAuLT	75.0 \pm 0.6	69.0 \pm 0.9	67.8 \pm 1.0	64.9 \pm 0.9

Table 3. Comparing ViLT, VAuLT, TomViLT and TomVAuLT on TWITTER-15 and TWITTER-17. **Bold** indicates best based on means only.

VAuLT in that regard. By disposing of that BERT from TomVAuLT, we arrive at the other variant, TomViLT, where only ViLT is used to process the text inputs. We compare these with ViLT and VAuLT using BERT on the dev sets of TWITTER-15 and TWITTER-17. In this manner, we present all possible configurations: i) ViLT: no deep encoder, ii) VAuLT: deep language encoder, iii) TomViLT: deep visual encoder, iv) TomVAuLT: both deep encoders.

Our results, shown in Table 3, demonstrate that having a deep visual encoder actually hurts performance in the presence of a deep language encoder, contrasting previous assumptions that transferring a “strong” visual vocabulary is essential, which to the best of our knowledge is a novel analytical result.

5. CONCLUSION

In this work, we introduce VAuLT. It utilizes a large pretrained LM, such as BERT, to propagate enhanced language representations to ViLT in order to perform multimodal tasks. We find that this pre-processing to be integral, as ViLT’s impoverished language representations, owing to its limited language exposure during pretraining, cannot be easily fine-tuned to perform out-of-distribution tasks, such as affective analysis on multimodal social media data. Our approach obviates the requirements for extensive pretraining of ViLT on the desired domains, as the appropriate LMs can bridge the distribution shift between source and target. Importantly, we demonstrate state-of-the-art or competitive performance of VAuLT.

Our work opens up several research directions. Anecdotal evidence from experiments not presented in this work indicate VAuLT struggles with pure reasoning tasks compared to ViLT, a direction left for future research. Second, we observe rare divergences during training. This hints at possible misalignment of the output space of the LM and the input space of ViLT. Third, we offer some evidence that shallow visual encoders but deep language encoders can achieve better classification performance to alternatives, a paradigm that could be investigated in greater detail. Finally, VAuLT can be used in multilingual settings with the proper LM choice.

Despite the utility of multimodal emotion recognition on social media, as discussed earlier, multiple ethical concerns can be raised. These include utilization of these models for the identification and suppression of dissent and free speech beyond text, such as in the form of images and memes, as well as their general integration in the aforementioned social manipulation tools, rendering them more powerful with the inclusion of visual capabilities.

6. ACKNOWLEDGEMENTS

We thank Kristina Lerman and Keith Burghardt for their valuable guidance. This project was funded in part by DARPA under contract HR001121C0168.

7. REFERENCES

- [1] ALFANO, M., SULLIVAN, E., AND FARD, A. E. Ethical pitfalls for natural language processing in psychology. *The Atlas of Language Analysis in Psychology* (2022).
- [2] ANTOL, S., AGRAWAL, A., LU, J., MITCHELL, M., BATRA, D., LAWRENCE ZITNICK, C., AND PARIKH, D. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision* (2015), pp. 2425–2433.
- [3] BAZIOTIS, C., NIKOLAOS, A., CHRONOPOULOU, A., KOLOVOU, A., PARASKEVOPOULOS, G., ELLINAS, N., NARAYANAN, S., AND POTAMIANOS, A. Ntua-slp at semeval-2018 task 1: Predicting affective content in tweets with deep attentive rnns and transfer learning. In *Proceedings of The 12th International Workshop on Semantic Evaluation* (2018), pp. 245–255.
- [4] BRADSHAW, S., AND HOWARD, P. N. The global organization of social media disinformation campaigns. *Journal of International Affairs* 71, 1.5 (2018), 23–32.
- [5] CHOCHLAKIS, G., MAHAJAN, G., BARUAH, S., BURGHARDT, K., LERMAN, K., AND NARAYANAN, S. Leveraging label correlations in a multi-label setting: A case study in emotion. *arXiv preprint arXiv:2210.15842* (2022).
- [6] DEVLIN, J., CHANG, M.-W., LEE, K., AND TOUTANOVA, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [7] DOSOVITSKIY, A., BEYER, L., KOLESNIKOV, A., WEISSENBERN, D., ZHAI, X., UNTERTHINER, T., DEGHANI, M., MINDERER, M., HEIGOLD, G., GELLY, S., ET AL. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [8] DUKES, D., ABRAMS, K., ADOLPHS, R., AHMED, M. E., BEATTY, A., BERRIDGE, K. C., BROOMHALL, S., BROSCHE, T., CAMPOS, J. J., CLAY, Z., ET AL. The rise of affectivism. *Nature human behaviour* 5, 7 (2021), 816–820.
- [9] GARTEN, J., HOOVER, J., JOHNSON, K. M., BOGHRATI, R., ISKITCH, C., AND DEGHANI, M. Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior research methods* 50, 1 (2018), 344–361.
- [10] HE, K., ZHANG, X., REN, S., AND SUN, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778.
- [11] HOOVER, J., ATARI, M., MOSTAFAZADEH DAVANI, A., KENNEDY, B., PORTILLO-WIGHTMAN, G., YEH, L., AND DEGHANI, M. Investigating the role of group-based morality in extreme behavioral expressions of prejudice. *Nature Communications* 12, 1 (2021), 1–13.
- [12] HOOVER, J., JOHNSON, K., BOGHRATI, R., GRAHAM, J., AND DEGHANI, M. Moral framing and charitable donation: Integrating exploratory social media analyses and confirmatory experimentation. *Collabra: Psychology* 4, 1 (2018).
- [13] HU, A., AND FLAXMAN, S. Multimodal sentiment analysis to explore the structure of emotions. In *proceedings of the 24th ACM SIGKDD international conference on Knowledge Discovery & Data Mining* (2018), pp. 350–358.
- [14] KHAN, Z., AND FU, Y. Exploiting bert for multimodal target sentiment classification through input space translation. In *Proceedings of the 29th ACM International Conference on Multimedia* (2021), pp. 3034–3042.
- [15] KIM, W., SON, B., AND KIM, I. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning* (2021), PMLR, pp. 5583–5594.
- [16] LIN, T.-Y., MAIRE, M., BELONGIE, S., HAYS, J., PERONA, P., RAMANAN, D., DOLLÁR, P., AND ZITNICK, C. L. Microsoft coco: Common objects in context. In *European conference on computer vision* (2014), Springer, pp. 740–755.
- [17] LU, J., BATRA, D., PARIKH, D., AND LEE, S. Vilt: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems* 32 (2019).
- [18] MERCHANT, A., RAHIMTOROGHI, E., PAVLICK, E., AND TENNEY, I. What happens to bert embeddings during fine-tuning? In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP* (2020), pp. 33–44.
- [19] MOHAMMAD, S., BRAVO-MARQUEZ, F., SALAMEH, M., AND KIRITCHENKO, S. Semeval-2018 task 1: Affect in tweets. In *Proceedings of the 12th international workshop on semantic evaluation* (2018), pp. 1–17.
- [20] NGUYEN, D. Q., VU, T., AND NGUYEN, A. T. Bertweet: A pre-trained language model for english tweets. *arXiv preprint arXiv:2005.10200* (2020).
- [21] NIU, T., ZHU, S., PANG, L., AND EL-SADDIK, A. Sentiment analysis on multi-view social data. In *MultiMedia Modeling* (2016), p. 15–27.
- [22] REN, S., HE, K., GIRSHICK, R., AND SUN, J. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* 28 (2015).
- [23] SHARMA, C., BHAGERIA, D., SCOTT, W., PYKL, S., DAS, A., CHAKRABORTY, T., PULABAIGARI, V., AND GAMBACK, B. Semeval-2020 task 8: Memotion analysis—the visuo-lingual metaphor! *arXiv preprint arXiv:2008.03781* (2020).
- [24] SRINIVASAN, T., CHANG, T.-Y., ALVA, L. L. P., CHOCHLAKIS, G., ROSTAMI, M., AND THOMASON, J. Climb: A continual learning benchmark for vision-and-language tasks. *arXiv preprint arXiv:2206.09059* (2022).
- [25] SU, W., ZHU, X., CAO, Y., LI, B., LU, L., WEI, F., AND DAI, J. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530* (2019).
- [26] WAHL-JORGENSEN, K. *Emotions, media and politics*. John Wiley & Sons, 2019.
- [27] WISTRICH, A. J., RACHLINSKI, J. J., AND GUTHRIE, C. Heart versus head: Do judges follow the law of follow their feelings. *Tex. L. Rev.* 93 (2014), 855.
- [28] YU, J., AND JIANG, J. Adapting bert for target-oriented multimodal sentiment classification. In *IJCAI* (2019).
- [29] ZHU, T., LI, L., YANG, J., ZHAO, S., LIU, H., AND QIAN, J. Multimodal sentiment analysis with image-text interaction network. *IEEE Transactions on Multimedia* (2022).