

Leveraging Customer Feedback for Multi-modal Insight Extraction

Sandeep Sricharan Mukku , Abinesh Kanagarajan , Pushpendu Ghosh , Chetan Aggarwal

Amazon

{smukku, abinesk, gpushpen, caggar}@amazon.com

Abstract

Businesses can benefit from customer feedback in different modalities, such as text and images, to enhance their products and services. However, it is difficult to extract actionable and relevant pairs of text segments and images from customer feedback in a single pass. In this paper, we propose a novel multi-modal method that fuses image and text information in a latent space and decodes it to extract the relevant feedback segments using an image-text grounded text decoder. We also introduce a weakly-supervised data generation technique that produces training data for this task. We evaluate our model on unseen data and demonstrate that it can effectively mine actionable insights from multi-modal customer feedback, outperforming the existing baselines by 14 points in F1 score.

1 Introduction

Customer feedback is essential for businesses to design and improve their products and services, according to customer expectations. (Luo et al., 2022) observe that the multi-modal feedback is growing rapidly. However, most existing solutions (Mukku et al., 2023; Sircar et al., 2022; Liu et al., 2022) do not take into account the rich information that image and video feedback contain, which can enhance the actionability and improve the customer experience. To address this challenge, we propose a novel multi-modal architecture that extracts pairs of text segments and corresponding images that are relevant and actionable for a given product from customer feedback. These pairs can help businesses to increase the actionability, improve the product catalogue quality, enhance the customer experience and thereby reduce returns and replacements.

2 Related work

In recent years, multi-modal tasks, combining various data types such as images and text, have garnered significant interest in artificial intelligence

and natural language processing (Goyal et al., 2017; Zhou et al., 2020; Lu et al., 2023). Among these tasks, Visual Question Answering (VQA) stands out as a prominent domain (Antol et al., 2015), wherein the aim is to generate textual answers to questions based on images. VQA has evolved significantly over the years, thanks to various advancing contributions. Early works on VQA (Antol et al., 2015) laid the groundwork, delineating the fundamental framework and challenges of the task. Subsequent research delved into innovative methodologies, such as leveraging transformer networks (Zhou et al., 2020; Lu et al., 2023), and devising techniques for seamlessly integrating visual and textual information (Li et al., 2019; Lu et al., 2019; Kim et al., 2021). The field progressed with more advanced models and datasets, such as VQAv2 (Goyal et al., 2017), which improved the robustness and performance benchmarking of VQA models. Moreover, the strides in pre-training on both vision and language have significantly influenced the landscape, as evidenced by groundbreaking models like OSCAR (Li et al., 2020), BEiT-3 (Wang et al., 2022), and VLMO (Bao et al., 2021, 2022), which have achieved remarkable results. Additionally, techniques such as counterfactual data augmentation (Chen et al., 2020) have further enhanced the performance of VQA models.

Meanwhile, in the NLP domain, verbatim extraction and text summarization tasks have attracted a lot of attention. Models like BERT (Devlin et al., 2018) and GPT-3 (Brown et al., 2020) have shown remarkable abilities in extracting and summarising textual feedback. Abstractive text summarization models, such as the T5 (Raffel et al., 2019), have emerged as the state-of-the-art (SOTA) solutions for generating short and coherent summaries from longer texts.

Our work falls at the intersection of these two domains, where we propose a novel problem at the intersection of VQA and text summarization, where

the input is a customer feedback text and corresponding images, and the objective is to extract the verbatim (exact segments of the textual feedback) that best highlights or talks about the feedback image. We design a multi-modal approach that builds on VQA models and extends them with text generation capabilities, introducing a new technique for multi-modal insight extraction. Our work bridges the gap between multi-modal understanding and verbatim extraction, and enhances the interpretability of images within textual context, especially for customer feedback and images.

3 Problem Statement

Given a customer feedback consisting of a set of images $I = \{I_1, I_2, \dots, I_n\}$ and a text T , we extract m verbatims from T , denoted by $V = \{v_i \mid 1 \leq i \leq m, v_i \in T\}$. The aim is to extract all the relevant and actionable verbatims $\{v_k \in V\}$ that corresponds to the given image I_k .

4 Proposed Approach

In this section, we first present our data generation method that creates training data from raw feedback. Then, we introduce our model architecture that extracts pairs of verbatims and images that are relevant and actionable for a given product.

4.1 Weakly supervised Training Data Generation

We segment raw feedback text to extract actionable verbatim¹ (refer appendix section A.1 for exact process). We then generate training data by obtaining $m \times n$ possible verbatim-image pairs for each feedback text, where m is the number of verbatim and n is the number of images. Next, we compute cosine similarity scores for every verbatim-image pair using pre-trained CLIP (Radford et al., 2021). To evaluate the relevance of the feedback image-verbatim pairs predicted by base CLIP², we had them manually annotated as positive or negative pairs with clearly defined annotation guidelines (refer section 4.1.1). We established a validation set, denoted as **StratSet-1k**, comprising 1,000 image-verbatim pairs, carefully stratified (Algorithm 1) across 27 distinct product categories sourced from

¹key text-phrases extracted from customer feedback, that can be utilised to take actions

²We used the ViT-B/32 version of CLIP

the raw dataset by clustering^{3,4} the actionable verbatims.

Algorithm 1 Stratified Sampling(K)

```

1: Let  $C$  be the number of clusters of verbatims spread across
    $L$  product categories.
2: Let  $K$  be the no. of verbatims to be sampled
3:  $S \leftarrow \emptyset$ 
4: for  $c \in C$  do
5:    $V_c \leftarrow$  verbatims in cluster  $c$ 
6:    $D_c \leftarrow$  {product categories in  $V_c$ }
7:   for  $d \in D_c$  do
8:      $V_{c,d} \leftarrow$  verbatims in category  $d$  of cluster  $c$ 
9:      $k_d \leftarrow \lfloor K \cdot \frac{1}{|D_c|} \rfloor$ 
10:     $S \leftarrow S \cup \text{RandomSample}(V_{c,d}, k_d)$ 
11:   end for
12: end for
13: return  $S$ 

```

To determine the optimal threshold, we systematically adjusted it within the range of 0.19 to 0.31 and recorded the resulting precision and product category coverage. Our analysis (refer Figure 2) revealed that a threshold of 0.27 strikes a favorable balance between product category coverage and precision, effectively reducing noise in the training data. We prioritized precision, aiming for a minimum of 90%, which sufficiently covers the majority of product categories while accepting a lower level of recall, as these matched pairs are primarily used for fine-tuning the CLIP.




Image-verbatim similarity score on **StratSet-1k**

Figure 2: Product category coverage and precision as a function of raw data threshold for fine-tuning CLIP. The graph shows the trade-off between coverage and precision for different values of the threshold

³We used Fast-clustering to cluster the verbatims

⁴Clustering verbatims helps group the relevant intents and thereby diversify the sample set

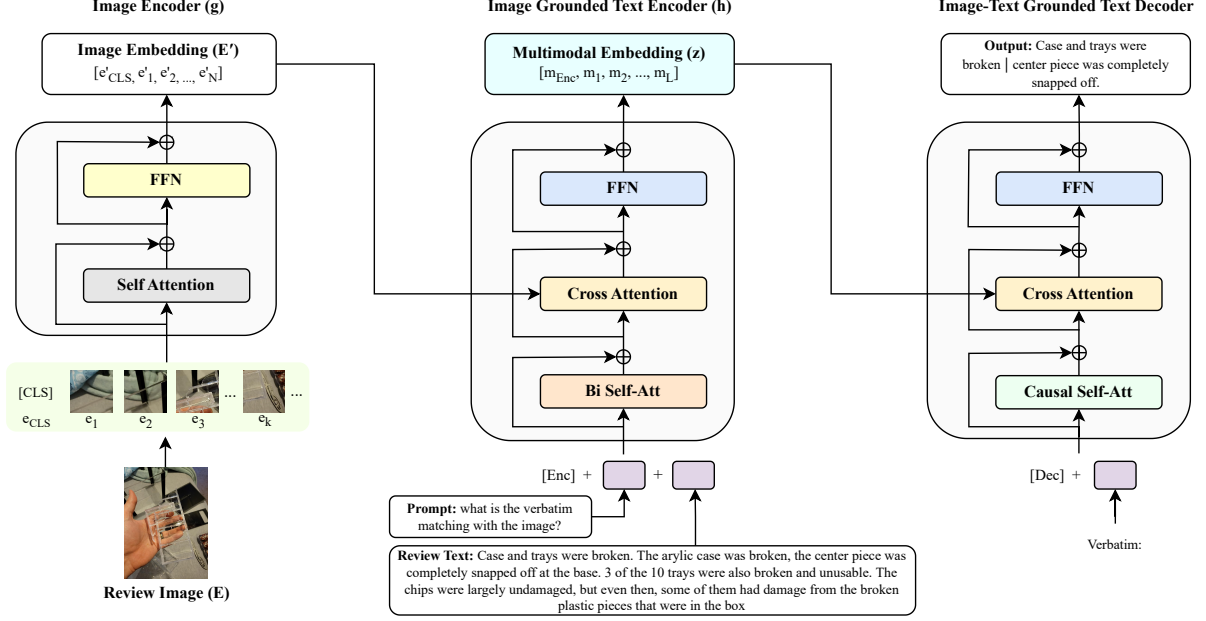


Figure 1: Multi-modal INSight Extraction (MINE) Architecture

After that, we fine-tune the CLIP using 210K positive pairs with a symmetric cross-entropy loss (refer Eqn. 2) that aligns image and text embeddings. The scaled pairwise cosine similarities between image and text embeddings are computed as follows:

$$\text{logits}_{I,T} = \frac{e_I^T e_T}{|e_I||e_T|} / \tau = \cos(\theta_{I,T}) / \tau \quad (1)$$

where e_I and e_T are the image and text embeddings, and $\theta_{I,T}$ is the angle between them. We compute the loss over both images and verbatim (text) by averaging the losses for each modality:

$$\begin{aligned} \text{loss} &= (\text{loss}_i + \text{loss}_t) / 2 \\ &= -\frac{1}{2} \sum_{i=1}^n \left[\log \frac{\exp(\cos(\theta_{i, \text{labels}_i}) / \tau)}{\sum_{t=1}^m \exp(\cos(\theta_{i,t}) / \tau)} \right. \\ &\quad \left. + \log \frac{\exp(\cos(\theta_{\text{labels}_i, i}) / \tau)}{\sum_{j=1}^n \exp(\cos(\theta_{j,i}) / \tau)} \right] \quad (2) \end{aligned}$$

where $\text{labels}_i \in V$ is the most relevant text of the i -th image, j iterates over the image classes, and τ is a temperature parameter that controls the smoothness of softmax distribution. This loss function encourages the model to learn embeddings that are close to each other for positive pairs and far apart for negative pairs. The loss is symmetric as it is

computed over both modalities. Finally, we use this trained model to obtain the matched (positive) and mismatched (negative) pairs from segments and images along with their similarity scores (sample training example shown in appendix section A.3).




Image-verbatim similarity score on **StratSet-1k**

Figure 3: Training data threshold selection

We employed the fine-tuned CLIP model to make inferences on the StratSet-1k dataset. Subsequently, we manually annotated the inferred pairs as either positive or negative, adhering to Inter-Annotator Agreement (IAA) process (Artstein, 2017). We then computed Precision, Recall,

and F1-score across a range of threshold values, as illustrated in Figure 2. Our analysis revealed that a threshold of 0.225 offers an optimal balance between precision and recall, thus serving as the foundation for creating training data for our problem. This refinement resulted in a notable enhancement of the fine-tuned CLIP model, achieving 75% F1-score, representing a 3% improvement compared to the off-the-shelf CLIP model.

4.1.1 Annotation Guidelines: Verbatim-Image relevancy

Following are the annotations guidelines given to annotators that define the relevancy of Verbatim-Image pair:

1. **Object relevance:** When certain object discussed in the text is found in the image in any forms which is may not be explicitly mentioned, it will be considered as relevant pair.
2. **OCR relevance:** When certain entity discussed in the text is represented in the image by some form of text, the pair is considered as relevant.
3. **Semantic relevance:** When the information discussed in the text is contextually represented in the image, the pair to be marked relevant.

Each pair is annotated (*relevant / not relevant*) by two experts separately and resolved by third in-case of conflicting annotation. We used IAA process and achieved Cohen’s kappa (Cohen, 1960) of 0.89.

4.2 Model Components

In this section, we describe the components of our proposed architecture for multi-modal insight extraction, which we refer to as MINE. Motivated by BLIP (Li et al., 2022), MINE consists of three main modules: an image encoder, an image-grounded text encoder, and an image-text grounded text decoder. The image encoder uses a visual transformer to extract visual features from the input image. The image-grounded text encoder fuses the input text with the visual features using cross-attention layers. The image-text grounded text decoder generates the output text using causal self-attention and cross-attention layers conditioned on the multi-modal representation from the encoder and the previous tokens.

4.2.1 Image Encoder

Image encoder is built using the visual transformer ViT-B/16 (Dosovitskiy et al., 2021), which consists of an image encoder and a transformer encoder. The image encoder splits the input image into patches of size 16×16 pixels and converts each patch into a vector of 768 dimension embedding. The transformer encoder (which has 12 layers and each layer performs different operations on the embeddings, such as attention, normalization, and feed-forward networks) takes the sequence of embeddings and outputs a new sequence of embeddings that contains more information about the image content and context. We formulate the transformer encoder as a function g that maps a sequence of embeddings E to another sequence of embeddings E' :

$$E' = g(E) = [e'_{CLS}, e'_1, e'_2, \dots, e'_N] \quad (3)$$

where e'_{CLS} is the updated [CLS] token (added to represent the embeddings as image features), e'_k is the updated embedding for the k -th patch (e_k), and N is the number of patches in the image.

4.2.2 Image Grounded Text Encoder

The text encoder is based on BERT base (Devlin et al., 2018), which has 12 transformer blocks with self-attention and feed-forward network (FFN) layers. To fuse the visual embedding from the image encoder, an extra cross-attention layer is inserted between the self-attention and FFN layers in each transformer block. This layer updates the text embeddings by attending to the image embeddings. The attention mechanism in the cross-attention layer computes attention scores for each token in the text sequence with respect to the image embeddings and determine how much importance each token in the text should place on the information in the image embeddings. A special token [ENC] is appended to the start of the input text (refer Section 5.2 for exact prompts used) to provide an identity for the encoder input. The text encoder outputs a multi-modal embedding of the image-text pair as follows:

$$z = h(T, E') = [m_{ENC}, m_1, m_2, \dots, m_L] \quad (4)$$

where h is the image-grounded text encoder function, T is the text input, E' is the embedding from image encoder (g), m_{ENC} is the multi-modal embedding for the [ENC] token, m_i is the multi-modal

embedding for the i -th text token, and L is the length of the text input.

4.2.3 Image-Text Grounded Text Decoder

The decoder follows the same architecture as the encoder, except that it uses causal self-attention instead of bidirectional self-attention. The multi-modal embedding is also fused as a cross-attention layer between FFN and attention layer. The decoder shares the parameters of FFN and cross-attention layers with the encoder, which improves training efficiency and enables multitask learning. The causal self-attention layer allows the decoder to generate text tokens (y_t), conditioned on the previous tokens ($y_{<t}$) and the multi-modal representation (z). For a given input token (x_t), the output of the decoder at each time step t is:

$$y_t = f(x_t, y_{<t}, z) \quad (5)$$

4.3 Overall Architecture

We encode the review text and image pair into a multi-modal representation and decode it into a sequence of verbatims (feedback segments) that are relevant to the image. Figure 1 illustrates the overall architecture of MINE. The encoding process consists of two steps: we apply the image encoder to the review image to obtain an image embedding; then, we pass the review text and the image embedding to the image-grounded text encoder to produce a multi-modal embedding that fuses both modalities. The decoding process uses the image-text grounded text decoder, which takes the multi-modal embedding as input and generates verbatim tokens conditioned on it. During training, we provide the ground truth verbatim as input to the decoder and optimize it to predict the next token. We fine-tune our model, with an objective of extracting the insightful segment by optimizing cross entropy loss:

$$L(\theta) = -\frac{1}{N} \sum_{n=1}^N \sum_{t=1}^{T_n} \log p(y_t^n | y_{<t}^n, x^n, z^n; \theta) \quad (6)$$

where N is the number of review-image pairs in the dataset, T_n is the length of the verbatim sequence for the n -th pair, y_t^n is the verbatim token at time step t for the n -th pair, $y_{<t}^n$ is the sequence of previous tokens, x^n is the review verbatim, z^n is the multi-modal embeddings from the image grounded text encoder, and θ are the model parameters. Dur-

ing inference, we only give the [DEC] token as input to the decoder and let it generate verbatims.

5 Experimental settings

5.1 Dataset & Analysis

We obtain raw reviews dataset⁵ provided by (Ni et al., 2019), which contains over 233 million reviews from 29 unique product categories. We used the sample subset (K-cores subset, released as part of the original dataset) which have 973k reviews with images of 27 product categories for our analysis.

5.2 Prompt Engineering

We use the similarity scores that are obtained from the training data generation, as the confidence scores for each verbatim (see appendix section A.3 for an example). We compare three different prompting approaches for extracting actionable verbatim from the review text and image pair, and are as follows:

1. **Comprehensive Segment Extraction with Confidence Scores (CSECS):** We generate all possible verbatim from the text and their scores based on how well they match the image. The ground truth includes all verbatim and their scores for each pair.
2. **Matching Segment Extraction with Confidence Scores (MSECS):** We generate only the verbatim that match image and their scores. The ground truth includes only matching verbatim and their scores for each pair.
3. **Matching Segment Extraction (MSE):** We generate only the verbatim that match the image without any scores. The ground truth includes only the matching verbatim predicted by the fine-tuned CLIP model for each pair.

Refer appendix section A.4 for exact prompts and targets used during training.

5.3 Approaches

As an additional multi-modal approaches, we tried ALBEF (Li et al., 2021) and VL-T5 (Cho et al., 2021) along with MINE. We use 80K verbatim-image pairs sampled from all product categories to train both approaches.

ALBEF: Given that the training data was generated using a weakly supervised methodology and

⁵dataset can be found [here](#)

Model	Prompting Approach	Precision (Correctness)	Recall	F1-score	Completeness
ALBEF	CSECS	0.61	0.52	0.56	0.71
	MSECS	0.63	0.59	0.61	0.76
	MSE	0.65	0.60	0.62	0.80
VL-T5	CSECS	0.65	0.53	0.58	0.76
	MSECS	0.65	0.57	0.61	0.82
	MSE	0.66	0.61	0.63	0.89
MINE	CSECS	0.69	0.70	0.69	0.73
	MSECS	0.71	0.76	0.73	0.91
	MSE	0.76	0.77	0.77	0.93

Table 1: Experimental results

may contain noise, we employ ALBEF, a state-of-the-art robust multi-modal vision and language representation learning model. ALBEF leverages momentum distillation, a self-training technique, to glean knowledge from pseudo-targets, enhancing its resilience to noise within the training data. We considered the VQA setting⁶ for our baselines. ALBEF consists of an image encoder and a text encoder, followed by a 6-layer auto-regressive answer decoder, where we used the pretrained encoder weights and only finetuned the decoder for all the three prompting approaches, as mentioned in section 5.2.

VL-T5: We fine-tune VL-T5, on our task of generating actionable verbatim from review text and image. We use Faster R-CNN (Ren et al., 2016) to obtain 36 object features from the review image, which is concatenated with the tokenized input text consisting of the prompt question and the review text. The resulting sequence is fed into the bidirectional multi-modal encoder of VL-T5, which learns to encode both textual and visual information. The decoder of VL-T5 then generates the actionable verbatim as the output.

5.4 Training Details

MINE was initialized using the pre-trained weights of BLIP base model and fine-tuned using AdamW optimizer (Loshchilov and Hutter, 2019) for 10 epochs with a cosine learning rate scheduler (Loshchilov and Hutter, 2017) and a minimum and initial learning rate of $1e^{-6}$ and $5e^{-5}$ respectively. In comparison, ALBEF was fine-tuned using 128 as batch size for 20 epochs with the following hyperparameters: AdamW optimizer, a learning rate of $1e^{-5}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$ and

⁶<https://github.com/salesforce/ALBEF/blob/main/VQA.py>

a weight decay of 0.05. Similarly, we fine-tuned VL-T5 for 5 epochs using AdamW optimiser with a learning rate of $5e^{-5}$ and linear warmup scheduler of 5%. At inference, we adjusted the max length to 512 to enable the model to generate longer and more relevant verbatim. Additionally, we employed an image resolution of 225px, consistent with the BLIP base model. We also experimented with different decoding methods for MINE, such as beam search (Hu et al., 2015) (*beam_size* = 10), consistent top-k sampling (Welleck et al., 2020) (*top_k* = 50) and nucleus sampling (Holtzman et al., 2019) (*top_k* = 50 and *top_p* = 0.95).

5.5 Results

We found that MSE approach with beam search decoding produced the most complete and relevant verbatim. We compared different prompt settings and approaches, and measured the quality of the generated verbatim using precision (correctness), recall and completeness metrics. We gave the exact definitions of these metrics in appendix section A.2. Table 1 summarized the results of our experiments. We see that our approach is effective than existing baselines like ALBEF and VL-T5 in both identifying the exact intent and extracting the complete verbatim, present in the image. In addition to beam search, we experimented with different decoding methods for MINE and reported the results in Table 2.

6 Conclusion

In this paper, we introduced a novel architecture MINE, for extracting insights from multi-modal customer feedback encompassing both text and images. We proposed a weakly-supervised data generation approach that leveraged raw feedback text and images to create relevant verbatim and image

Prompt Type	Decode method	Precision (Correctness)	Recall	F1-score	Completeness
CSECS	topK	0.57	0.48	0.52	0.66
	nucleus	0.62	0.6	0.61	0.71
MSECS	topK	0.6	0.59	0.59	0.69
	nucleus	0.65	0.69	0.67	0.83
MSE	topK	0.57	0.7	0.63	0.72
	nucleus	0.68	0.72	0.7	0.86

Table 2: Additional Experimental results

pairs. We also proposed an unsupervised insight extraction approach that used an image-grounded text encoder to learn latent representations of the feedback text and corresponding images, and an image-text grounded text decoder to extract actionable verbatim relevant to image. We evaluated our approach on real-world dataset and demonstrated its effectiveness in extracting actionable insights from multi-modal feedback with minimal supervision, achieving a 14-point improvement in F1-score over the existing baselines and uni-modal insight extraction.

Future work: Our research opens up several avenues for further exploration. As a future direction, we would like to extend our method to video feedback by incorporating the temporal dimension and extracting relevant video snippets, as mention in the feedback text. We also aim to structure the extracted insights in hierarchical way that facilitates decision making for users. Moreover, we plan to localise the regions of interest in the feedback images that correspond to the verbatim, using bounding boxes, to help identify the specific features or issues that customers refer to in their feedback. Finally, we intend to generate suggestions for product enhancements or fixes based on the extracted insights and the customer sentiment.

Limitations

Like any other ML architecture, even MINE architecture has some limitations. One of the limitations is that the MINE architecture is not pre-trained with language modelling as an objective (such as Prefix language modelling-style / Deshuffling-style / i.i.d noise, replace spans-style (Raffel et al., 2020), or BERT-style (Devlin et al., 2019), or MASS-style (Song et al., 2019)). This means that the model may not always generate responses that are identical to the feedback texts in terms of wording, even though they generate text the same intent. We found that around 19% of the cases had different

wording but with same intent. Another limitation is that the extraction task can be computationally demanding with a higher number of beams, which may affect the performance and scalability of the approach using low compute resources.

Ethics Statement

In this paper, we used a publicly available dataset in an almost unsupervised fashion, with minimal reliance on human annotations for labeling. Annotators are subject matter experts who specialize in annotating image-text samples, whether they are relevant or not. They are compensated accordingly, following industry standards set by the organisation. We carefully considered all aspects related to annotator agreement and fairness. However, the dataset is anonymized and does not contain any identifiable or traceable information. Thus, we respect the privacy and confidentiality of the consumers of the product and do not expose them to any potential harm or misuse. The dataset is widely used and well-known in the natural language processing community. It has been previously analyzed and evaluated by several researchers and practitioners. Our work does not introduce any bias or prejudice either, as we do not make any assumptions or judgments based on the feedback text or images. Our work maintains a purely empirical and objective approach, without favoring or disadvantaging any individual or product. Throughout our research process and reporting, we adhered to the ACL code of ethics and professional conduct.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. [Vqa: Visual question answering](#). In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Ron Artstein. 2017. [Inter-annotator agreement](#). *Handbook of linguistic annotation*, pages 297–313.

- Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. 2021. [Beit: Bert pre-training of image transformers](#). *arXiv preprint arXiv:2106.08254*.
- Hangbo Bao, Wenhui Wang, Li Dong, Qiang Liu, Owais Khan Mohammed, Kriti Aggarwal, Subhojit Som, Songhao Piao, and Furu Wei. 2022. [Vlmo: Unified vision-language pre-training with mixture-of-modality-experts](#). *Advances in Neural Information Processing Systems*, 35:32897–32912.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#).
- Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. 2020. [Counterfactual samples synthesizing for robust visual question answering](#). In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10800–10809.
- Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. [Unifying vision-and-language tasks via text generation](#).
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and psychological measurement*, 20(1):37–46.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *arXiv preprint arXiv:1810.04805*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). *ICLR*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. [Making the v in vqa matter: Elevating the role of image understanding in visual question answering](#).
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. [The curious case of neural text degeneration](#). *arXiv preprint arXiv:1904.09751*.
- Xiaoguang Hu, Wei Li, Xiang Lan, Hua Wu, and Haifeng Wang. 2015. [Improved beam search with constrained softmax for NMT](#). In *Proceedings of Machine Translation Summit XV: Papers*, Miami, USA.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. [Vilt: Vision-and-language transformer without convolution or region supervision](#).
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. [Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training](#). In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11336–11344.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. [Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation](#). In *ICML*.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. 2021. [Align before fuse: Vision and language representation learning with momentum distillation](#).
- 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](#).
- Yang Liu, Varnith Chordia, Hua Li, Siavash Fazeli Dehkordy, Yifei Sun, Vincent Gao, and Na Zhang. 2022. [Leveraging seq2seq language generation for multi-level product issue identification](#). In *Proceedings of The Fifth Workshop on e-Commerce and NLP (ECNLP 5)*, pages 20–28.
- 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. [Sgdr: Stochastic gradient descent with warm restarts](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#).
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. [Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks](#).
- Siyu Lu, Yueming Ding, Mingzhe Liu, Zhengtong Yin, Lirong Yin, and Wenfeng Zheng. 2023. [Multiscale feature extraction and fusion of image and text in vqa](#). *International Journal of Computational Intelligence Systems*, 16(1):54.
- Hanyang Luo, Wanhua Zhou, Wugang Song, and Xiaofu He. 2022. [An empirical study on the differences between online picture reviews and text reviews](#). (7).
- Sandeep Sricharan Mukku, Manan Soni, Chetan Aggarwal, Jitenkumar Rana, Promod Yenigalla, Rashmi Patange, and Shyam Mohan. 2023. [Insightnet: Structured insight mining from customer feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 552–566.
- Minh Van Nguyen, Viet Dac Lai, Amir Poursan Ben Veyseh, and Thien Huu Nguyen. 2021. [Trankit: A light-weight transformer-based toolkit for multilingual natural language processing](#). In *Proceedings of*

the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations.

- Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. [Justifying recommendations using distantly-labeled reviews and fine-grained aspects](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong, China.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Michael Narang, Adam Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#).
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2016. [Faster r-cnn: Towards real-time object detection with region proposal networks](#).
- Prateek Sircar, Aniket Chakrabarti, Deepak Gupta, and Anirban Majumdar. 2022. [Distantly supervised aspect clustering and naming for e-commerce reviews](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 94–102.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. [Mass: Masked sequence to sequence pre-training for language generation](#).
- Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, and Furu Wei. 2022. [Image as a foreign language: Beit pretraining for all vision and vision-language tasks](#).
- Sean Welleck, Ilya Kulikov, Jaedeok Kim, Richard Yuanzhe Pang, and Kyunghyun Cho. 2020. [Consistency of a recurrent language model with respect to incomplete decoding](#).
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. [Unified vision-language pre-training for image captioning and vqa](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):13041–13049.

A Appendix

A.1 Verbatim Extraction

We use the following steps to extract verbatims (actionable feedback text segments) from raw customer feedback source:

1. **Text Pre-processing:** We apply text pre-processing techniques to remove html tags, urls, accented characters, extra spaces and other noise from the feedback text.
2. **Segmentation:** We use Trankit (Nguyen et al., 2021), a fast and lightweight transformer-based toolkit for natural language processing, to segment each feedback into meaningful units. We also implement some heuristics to avoid having segments with only single word.
3. **Removing non-actionable (neutral) segments:** We use a RoBERTa (Liu et al., 2019) based 3-class sentiment classifier⁷ to evaluate the sentiment of each feedback segment. It has been rigorously evaluated and demonstrates an impressive 92% F1-score on our specific sentiment classification task. We make the decision to exclude segments classified as neutral, as these segments may not provide actionable feedback for brands and selling partners who seek feedback to improve their products and services.

A.2 Verbatim evaluation metrics

1. **Completeness:** This metric measures the proportion of segments that convey a complete meaning. A segment is considered complete if it expresses a coherent and relevant idea about the feedback.

$$\text{Completeness} = \frac{\text{Number of complete segments}}{\text{Total number of segments generated}}$$

2. **Correctness/Precision:** This metric measures the proportion of segments that match both the text segment and the image. A segment is considered correct (relevant) if it accurately reflects the information and sentiment from

both sources.

$$\text{Correctness} = \frac{\text{Number of correct segments}}{\text{Total number of segments generated}}$$

3. **Recall:** This metric measures how many relevant (to the image) verbatims are extracted among all the relevant verbatims present in the feedback.

$$\text{Recall} = \frac{\text{Number of relevant verbatims}}{\text{Total number of relevant verbatims}}$$

⁷<https://huggingface.co/siebert/sentiment-roberta-large-english>

A.3 Training Data








Input		Output		
Feedback Text	Feedback Image	Feedback Image	Verbatim	Similarity Score
<p>They don't look nice. It looked nice for a brief period of time; then the finish came off, that fadeout the pleated color and became brownish. The one I'd purchased elsewhere before lasted years, looked nice and was retired only because the "stone" was lost.</p>			They don't look nice	0.19
			It looked nice for a brief period of time	0.12
			then the finish came off	0.29
			fadeout the pleated color and became brownish	0.31
			The one I'd purchased elsewhere before lasted years	0.13
			looked nice and was retired only because the "stone" was lost	0.21

Table 3: Training Data

A.4 Experimented Prompts

Task	Input Prompt	Target
CSECS	<p>Extract all the verbatim and confidence score of each matching with image? Feedback: <feedback text></p>	<p><verbatim 1>; <confidence 1> <verbatim 2>; <confidence 2></p>
MSECS	<p>Extract all the verbatim and confidence score of each matching with image? Feedback: <feedback text></p>	<p><matching verbatim 1>; <confidence 1> <matching verbatim 2>; <confidence 2></p>
MSE	<p>What is the verbatim matching with the image? Feedback: <feedback text></p>	<p><matching verbatim 1> <matching verbatim 2></p>

Table 4: Experimented Prompts

A.5 MINE: Sample Predictions























Model Input (Multi-modal Feedback)		Model Output	
Text	Image	Image	Verbatim
<p>Paint is outside the lines and it really looks sloppy. So, i really dont like it. It looks so bad. I would not order this it is product. And it was rolled up like a newspaper. Complete and total waste. The paint is pasted on the nose as well.</p>			Looks like its used.
			water marks coming out of the steam holes.
<p>I was disappointed. The shoulder straps did not have anything to secure the loose strap so you essentially have a loose strap getting in the way. This makes the purse look really cheap. Inside, the zipper is misaligned so when it is zipped, you see a gap. When I try to rezip, it doesn't correct itself. Lastly, the inside pocket must have been oversewn, because when you open the pocket, you can see the bright green threading. Lastly, the gold-plated clasps that connect the purse and shoulder straps together were also not aligned.</p>			the zipper is misaligned
			shoulder straps together were also not aligned
			inside pocket must have been oversewn
<p>I wish the product to be good and expecting for it. Got the product just by today morning. It never met my expectations and I am totally disappointed!. This set is really small. The shakers don't fit in the holder and the holder is pathetic! Not made well at all!!!</p>			Not made well at all
			
<p>I've had these boots for a few months and only wear them when I ride (less than twice a week). For no reason at all the heel on the right boot fell off (in half). Looks like the heel is made of particle cardboard. I would have expected that boots in this price range would not fall apart under light use in just a few months! What a rip off!</p>			heel on the right boot fell off,
			fall apart under light use.
<p>I was very pleased with the print and very excited to hang these up. We had a room that was lonely and in need of some fun print, these looked like the perfect fit. I was very disappointed once they were hung. The packaging stated 95" length; what I received was one of 88" and the other of 92". Although the print was nice, what can one do with them? very disappointed in this purchase.</p>			received was one of 88" and the other of 92"
			received was one of 88" and the other of 92"
			

Table 5 continued from previous page

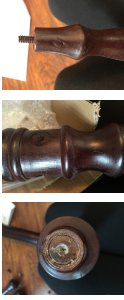





Model Input (Multi-modal Feedback)		Model Output	
Text	Image	Image	Verbatim
<p>The first one I ordered had a big chip on one of the small pieces and I found the loose chip in the package too. They sent a replacement pretty quick, but the second one was even worse - multiple cracks, uneven finish and one big problem - one of the main middle pieces did not fit the connecting piece. The hole was not big enough. This was a gift, and I felt so embarrassed.</p>			<p>even worse - multiple cracks</p>
			<p>hole was not big enough</p>
<p>Very disappointed in this ring. I have worn it most every day, taking it off for showers, for the past couple of months and it's already tarnishing, and it looks like the silver is coming off of the band. Looks more like it's silver plated instead of actual sterling silver. I have owned SS and Black Hills Gold Jewelry for years and I am careful not to expose my jewelry to harsh chemicals, etc.</p>			<p>it's already tarnishing</p>
			<p>it's already tarnishing</p>

Table 5: Sample Predictions