

InsightNet : Structured Insight Mining from Customer Feedback

Sandeep Sricharan Mukku , Manan Soni , Jitenkumar Rana , Chetan Aggarwal ,
Promod Yenigalla , Rashmi Patange and Shyam Mohan

Amazon

{smukku, sonmanav, jitenkra, caggar, promy, rpatang, shyaamm}@amazon.com

Abstract

We propose InsightNet, a novel approach for the automated extraction of structured insights from customer reviews. Our end-to-end machine learning framework is designed to overcome the limitations of current solutions, including the absence of structure for identified topics, non-standard aspect names, and lack of abundant training data. The proposed solution builds a semi-supervised multi-level taxonomy from raw reviews, a semantic similarity heuristic approach to generate labelled data and employs a multi-task insight extraction architecture by fine-tuning an LLM. InsightNet identifies granular actionable topics with customer sentiments and verbatim for each topic. Evaluations on real-world customer review data show that InsightNet performs better than existing solutions in terms of structure, hierarchy and completeness. We empirically demonstrate that InsightNet outperforms the current state-of-the-art methods in multi-label topic classification, achieving an F1 score of 0.85, which is an improvement of 11% F1-score over the previous best results. Additionally, InsightNet generalises well for unseen aspects and suggests new topics to be added to the taxonomy.

1 Introduction

Customer reviews provide rich insights for various stakeholders, such as businesses, brands, and customers. They can inform product development, enhance customer experience, track reputation, and guide purchase decisions. However, customer reviews pose several challenges for analysis, such as subjectivity, variation, noise, domain-specificity, volume, and dynamism. Existing solutions for extracting structured insights from reviews, such as topic classification (Zheng, 2021; Sánchez-Franco et al., 2019), polarity identification (Bilal and Almazroi, 2022; Gopi et al., 2023), and verbatim extraction (Majumder et al., 2022), suffer from several drawbacks that limit their effectiveness and

applicability. These drawbacks include: (1) low accuracy and reliability in generating and extracting insights from reviews, (2) lack of coherence and clarity in the output, which makes it hard to act upon, (3) high dependency on large amounts of annotated data, which are scarce and expensive to obtain, (4) task-specificity, (5) inability to handle multiple tasks simultaneously, (5) skewed data distribution towards a few dominant aspects, which biases the models' performance (81% of reviews are covered by just 12% of topics, see Appendix C.1 for detailed analysis), (6) reliance on predefined aspects and limitations to discover new topics.

In this paper, we present three key modules to address the challenges of the existing approaches as follows: (1) AutoTaxonomy: A method to generate a hierarchical taxonomy of aspects with minimal supervision (section 4.2). This helps to organise the output in a structured and hierarchical form (2) SegmentNet: An unsupervised data creation technique using semantic similarity based heuristics to produce labelled data (section 4.3) that contains topic, polarity and verbatim for each review. Here, verbatim is the exact segment of the review that describes the topic identified. (3) InsightNet: A generative model for insights extraction. We model aspect identification as a multi-task hierarchical classification problem and then leverage the generative model (section 4.1) to classify topic (granular aspect), identify sentiment, extract verbatim and also discover new topics that are not in the current taxonomy. We use T5-base (Raffel et al., 2020) as a pre-trained Large Language Model (LLM) and fine-tune it with the data obtained from SegmentNet. Thus, we do not require any manually annotated data to train InsightNet.

2 Related Work

Insight extraction from customer reviews is a well researched problem. Researchers have posed this problem in various frameworks such as heuristic based insight extraction, aspect based senti-

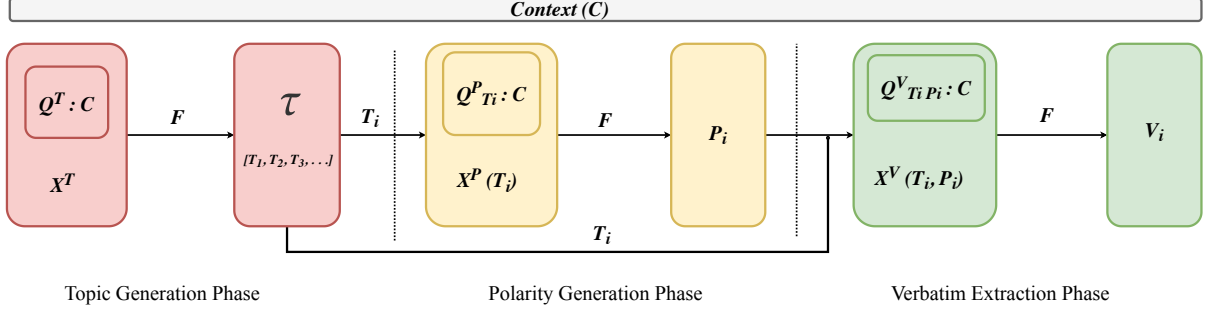


Figure 1: Decomposed Sequential Prompting - InsightNet

ment analysis, text summarisation, topic modeling, generative modeling. Rana and Cheah (2015); Kang and Zhou (2017) proposed a rule-based approach to extract insights from reviews. However, these approaches require huge manual efforts and domain expertise to discover patterns, update them frequently as new products launch and create rules. Hu and Liu (2004); Baccianella et al. (2009) propose aspect based sentiment analysis methods which first extract aspects and then rate reviews on each aspect. However, aspects obtained using these methods are not granular enough for actionability. Titov and McDonald (2008); Brody and Elhadad (2010); Sircar et al. (2022) proposed unsupervised approaches for aspect and sentiment analysis from reviews, but these methods suffer from two main limitations : a) redundancy of clusters and b) low interpretability, as the clusters produced are not actionable, structured or intuitive. Recently, generative approaches (Raffel et al., 2020; Brown et al., 2020) demonstrated promising performance on wide range of Natural Language Processing (NLP) tasks. Liu et al. (2022) used a seq-to-seq model to generate product defects and issues from customer reviews, but they lack structure.

3 Problem Statement

Given a customer review C , we aim to extract a set of all the relevant and actionable insights I_1, I_2, \dots, I_k , where each insight I_i is composed of a granular topic T_i , a corresponding polarity P_i , and a set of verbatims V_i associated with it.

4 Methodology

In this section, we present our generative approach, InsightNet, for mining insights (topics, polarities, verbatims) from raw reviews (obtained from Amazon US marketplace¹). Next, we describe how we construct a multi-level hierarchical taxonomy from

the reviews to help organize topics in a meaningful way. Then, we introduce how the labelled data is created using SegmentNet. Later, we explain how we apply post-processing techniques to eliminate redundant topics and surface new topics that are not covered by the taxonomy. Finally, we discuss the experiments that led us to the design of InsightNet.

4.1 InsightNet: Generative Multi-task model for Insights Extraction

The InsightNet architecture (Figure 1) is based on decomposed prompting (Khot et al., 2023), allowing to solve the complex task of extracting actionable verbatims and assigning a topic and a polarity to each verbatim. It consists of three phases of prompting, one for topic generation, one for polarity generation, and one for verbatim extraction.

4.1.1 Topic Generation Phase

We construct prompt X^T by appending question Q^T to context C (raw review), where Q^T is question prompt to generate list of granular actionable topics $\tau : [T_1, T_2, T_3, \dots]$. We feed InsightNet model (F), with X^T to generate actionable topics list, τ

$$X^T = Q^T : C \quad ; \quad \tau = F(X^T) \quad (1)$$

4.1.2 Polarity Generation Phase

In this phase, we use the model (F) sequentially to generate the polarity (P_i) for each of the topic (T_i) extracted in the previous phase. We feed $X^P(T_i)$, which consists of the context C and the question prompt $Q^P_{T_i}$ for the topic $T_i \in \tau$. We then form Π , a set of topic-polarity pairs, $\Pi : [(T_1, P_1), (T_2, P_2), (T_3, P_3), \dots]$

$$X^P(T_i) = Q^P_{T_i} : C \quad ; \quad P_i = F(X^P(T_i)) \quad (2)$$

4.1.3 Verbatim Extraction Phase

In this last phase, we use the model (F) sequentially to extract verbatim (V_i) for each topic-polarity pair (T_i, P_i) produced previously. We feed

¹https://huggingface.co/datasets/amazon_us_reviews

$X^V(T_i, P_i)$ to the model, which consists of the context C and the question prompt Q_{T_i, P_i}^V for the pair $(T_i, P_i) \in \Pi$.

$$X^V(T_i, P_i) = Q_{T_i, P_i}^V : C \quad ; \quad V_i = F(X^V(T_i, P_i)) \quad (3)$$

For a review, if we get N topics in the first stage, then we subsequently use N prompts for each of the next two stages. Thus, we use a total of $2N + 1$ prompts per review, where N is the number of actionable topics present in the review.

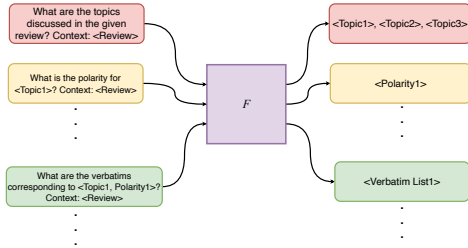


Figure 2: InsightNet Prompting

4.2 AutoTaxonomy: Semi-supervised Taxonomy Creation

We propose a bottom-up method to generate a hierarchical auto-taxonomy from reviews with weak supervision. This means we start with identifying Granular Topics from the reviews, then group them into broader (high-level) topics. This helps us preserve structure and create hierarchy for the output. We segment raw reviews (refer to Appendix section B.1 for exact segmentation steps) and assign a polarity to each segment (see equation 4). We discard segments with neutral polarity. The following steps illustrate the process of Taxonomy creation:

1. **Clustering review segments:** We cluster the positive and negative segments separately using Fast clustering², a sentence transformer (Reimers and Gurevych, 2019) method. It uses cosine similarity to cluster sentence embeddings based on a threshold value. We obtain clusters for each sentiment class, representing different aspects or topics that the reviewers mentioned in their feedback.

2. **Merging similar clusters & Cluster naming:** Human annotators merge duplicate clusters and name each cluster using pre-defined taxonomy guidelines. They name each cluster with fine-grained topic names that reflect the main idea of

²https://github.com/UKPLab/sentence-transformers/blob/master/examples/applications/clustering/fast_clustering.py

that cluster, forming the Granular Topics (Level-3) of the taxonomy.

3. **Creating hierarchy:** To structure the taxonomy, we group similar Granular Topics into Hinge Topics (Level 2) and Coarse Topics (Level 1), resulting in a multi-level hierarchical taxonomy.

4. **Keyword generation:** To get an exclusive and exhaustive set of keywords, we refined the clusters (of segments) obtained in step 2 (above). We applied the semantic similarity function (equation 5) to perform **Intra-cluster** (refer section 4.2.1) cleaning to remove redundant and semantically duplicated keywords and **Inter-cluster** (refer section 4.2.2) cleaning of keywords, to eliminate ambiguous and overlapping keywords.

4.2.1 Intra-cluster Cleaning

Given a set of keywords $K = \{k_1, k_2, \dots, k_n\}$, Algorithm 1 returns a cleaned set of non redundant keywords C such that $\forall k_i, k_j \in C, SimST(k_i, k_j) \leq \delta_a$

Algorithm 1 Intra-cluster cleaning

- 1: Let $K = \{k_1, k_2, \dots, k_n\}$ are the set of keywords of a given topic T
 - 2: Let δ_a be the intra-cluster threshold for redundancy
 - 3: Initialize $D \leftarrow \emptyset, C \leftarrow K$
 - 4: **for all** $(k_i, k_j) \in K \times K$ **do**
 - 5: **if** $SimST(k_i, k_j) > \delta_a$ **then**
 - 6: Add k_i or k_j to D
 - 7: **end if**
 - 8: **end for**
 - 9: Remove all elements in D from C
 - 10: Return C
-

4.2.2 Inter-cluster Cleaning

Algorithm 2 compares the keywords across all the topics and removes keywords which are similar to each other by converting the keywords into sentence embedding and comparing the cosine similarity between them with the ambiguity threshold δ_e .

Taxonomy derived using this approach has 91% exclusivity (uniqueness of topics) and 94% exhaustivity of topics when evaluated manually. (Table 8 in Appendix section C.2 presents sample of final taxonomy).

4.3 SegmentNet: Data Generation Mechanism

SegmentNet is a semantic matching algorithm that generates high quality training data with minimal

Algorithm 2 Inter-cluster cleaning

```
1: Let  $K = K_1, K_2, \dots, K_n$  be the initial set of
   keyword lists respectively for  $n$  topics  $T = T_1, T_2, \dots, T_n$ .
2: Let  $\delta_e$  be the inter-cluster threshold for ambiguity.
3: Create an empty hash table  $H$ .
4: for each keyword list  $K_i$  in  $K$  do
5:   for each keyword  $k_{ij}$  in  $K_i$  do
6:     if  $k_{ij}$  is not in  $H$  then
7:       Compute  $Sbert(k_{ij})$  and store it in
        $H$  with  $k_{ij}$  as the key and  $Sbert(k_{ij})$  as the
       value.
8:     end if
9:   end for
10: end for
11: for each keyword list  $K_i$  in  $K$  do
12:   for each keyword  $k_{ij}$  in  $K_i$  do
13:     for each other keyword list  $K_l$  in  $K$ 
       where  $l \neq i$  do
14:       for each keyword  $k_{lm}$  in  $K_l$  do
15:         if  $SimST(H[k_{ij}], H[k_{lm}]) >$ 
            $\delta_e$  then
16:           Remove  $k_{ij}$  from  $K_i$  and
            $k_{lm}$  from  $K_l$ 
17:         end if
18:       end for
19:     end for
20:   end for
21: end for
22: Return the final keyword lists  $K$ .
```

training. It extracts insights from reviews using the taxonomy alone. It assumes that insights are often in short phrases within a review. It produces insights at a segment level and then aggregates them at review level. This involves 3 major steps:

1. Segmentation: We use language syntax heuristics to split a review into segments. We observe that a segment typically has one sentiment and at most one topic.

2. Sentiment classification: We train a BERT-based model with two linear heads (one for +ve and one for -ve) to get the sentiment of segment

$$p, n = \text{SentimentClassifier}(S) \quad (4)$$

We use $\delta_p = 0.7$ as the classification threshold. A segment is neutral if $p < \delta_p$ and $n < \delta_p$, where p and n are the probabilities of a verbatim being positive and negative polarities respectively. We fine-tune sentiment classification model on 80k seg-

Algorithm 3 Topic matching

```
1: // Returns the most relevant topic  $T$  for a given
   segment  $S$  by applying heuristic rules.
2: Hyper-parameters:  $k = 5, \delta_h = 0.8, \delta_m = 0.3, \delta_a = 0.5$ 
3:  $T_n, S_n = \text{BestTopicAndScore}([T_i]_{i=1}^N, [SimST(S, T_i)]_{i=1}^N)$ 
4:  $T_{tkw}, S_{tkw} = \text{BestTopicAndScore}([T_i]_{i=1}^N, [\frac{1}{k} \max_p \{SimST(S, k_{i,j})\}_{j=1}^{M_i}]_{i=1}^N)$ 
5:  $T_{mkw}, S_{mkw} = \text{BestTopicAndScore}([T_i]_{i=1}^N, [\frac{1}{M_i} \sum_{j=1}^{M_i} \{SimST(S, k_{i,j})\}_{i=1}^N])$ 
6:  $A_i = \frac{1}{k} \max_p \{SimST(S, k_{i,j})\}_{j=1}^{M_i} + \frac{1}{M_i} \sum_{j=1}^{M_i} \{SimST(S, k_{i,j})\}$ 
7:  $T_{avg}, S_{avg} = \text{BestTopicAndScore}(\{A_i\}_{i=1}^N)$ 
8:
9: if  $S_{tkw} \geq \delta_h$  then
10:    $T = T_{tkw}$ 
11: else if  $S_n \geq \delta_h$  then
12:    $T = T_n$ 
13: else if  $S_{mkw} \geq \delta_h$  then
14:    $T = T_{mkw}$ 
15: else if  $T_{tkw} = T_n$  and  $S_{tkw} + S_n \geq 2 * \delta_m$ 
16:    $T = T_{tkw}$  then
17: else if  $T_{mkw} = T_{tkw}$  and  $S_{mkw} + S_{tkw} \geq 2 * \delta_m$ 
18:    $T = T_{mkw}$  then
19: else if  $T_n = T_{mkw}$  and  $S_n + S_{mkw} \geq 2 * \delta_m$ 
   then
20:    $T = T_n$ 
21: else if  $S_{avg} \geq \delta_a$  then
22:    $T = T_{avg}$ 
23: else
24:    $T = \emptyset$ 
25: end if
26:
27: return  $T$ 
```

ments with almost equal data for each label, which has 99.1% accuracy when evaluated manually.

SimST: We formulate semantic similarity function $SimST$ (equation 5) between two texts $text_i$ and $text_j$, where $sbert$ computes the SentenceBERT (Reimers and Gurevych, 2019) embedding of text.

$$SimST(text_i, text_j) = \cos(sbert(text_i), sbert(text_j)) \quad (5)$$

where $\cos(\mathbf{u}, \mathbf{v}) = \frac{\mathbf{u} \cdot \mathbf{v}}{\|\mathbf{u}\| \|\mathbf{v}\|}$ is the cosine similarity.

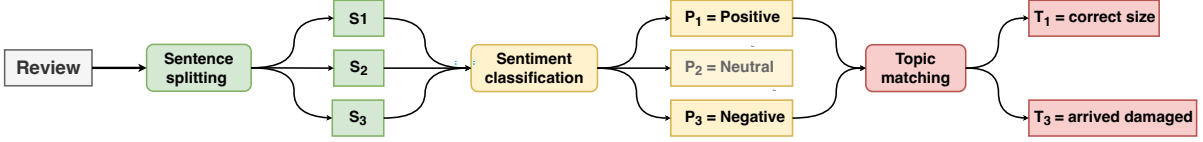


Figure 3: SegmentNet pipeline

Algorithm 4 Signalling Algorithm (*BTS*)

- 1: **procedure** BESTTOPICANDSCORE(T, X)
 - 2: // Finds the leading topic T_i as per the score values mentioned in the list X .
 - 3: **return** $T[\text{argmax}(X)], \text{max}(X)$
 - 4: **end procedure**
-

3. **Topic matching:** We devised heuristics based on the semantic matching function SimST (equation 5) and a signalling algorithm (see *BTS* Algorithm 4) to assign the best matching topic to a segment from a list of taxonomy topics. The signalling algorithm outputs the topic with the maximum similarity score and the value of that score among the given topic and similarity score pairs. Let S denote a segment and T its most relevant topic. We find T from the list of taxonomy topics (τ'), $[T_i]_{i=1}^N$ with each topic T_i has keywords $[k_{i,j}]_{j=1}^{M_i}$. We define three signals (using Algorithm 4), where the first signal (equation 6) is the semantically closest topic name and its score, the second signal (equation 7) is the topic with best mean score with the five closest keywords, and the last signal (equation 8) is the topic with the best mean score with all keywords.

$$T_n, S_n = \text{BestTopicAndScore}([T_i]_{i=1}^N, [\text{SimST}(S, T_i)]_{i=1}^N) \quad (6)$$

$$T_{tkw}, S_{tkw} = \text{BestTopicAndScore}([T_i]_{i=1}^N, [\frac{1}{5} \max_5 \{\text{SimST}(S, k_{i,j})\}_{j=1}^{M_i}]_{i=1}^N) \quad (7)$$

$$T_{mkw}, S_{mkw} = \text{BestTopicAndScore}([T_i]_{i=1}^N, [\frac{1}{M_i} \sum_{j=1}^{M_i} \{\text{SimST}(S, k_{i,j})\}_{j=1}^{M_i}]_{i=1}^N) \quad (8)$$

To identify most relevant topic, we use heuristics on the three signals for topic matching:

(a.) **High confidence match:** if any of the three signal scores is high, match with high scoring topic

($\text{score} \geq \delta_h$). This matches a segment that is very similar to a topic or keyword,

(b.) **Majority vote:** If any two signals give the same topic, match with the common topic ($\text{score} \geq \delta_m$). Since each of the three signals is an independent weak predictor of the correct topic, the fact that any two signals agree on a topic is a strong indicator of correctness,

(c.) **Best average score:** Match with the topic with the best average score across all three signals T_{avg} ($\text{score} \geq \delta_a$).

We present the topic matching algorithm (Algorithm 3) which is more robust to noisy keywords and identifies topics with higher precision than simple semantic matching.

4.4 Post-Processing

During inference, we leverage syntactic and semantic matching to tackle topics generated that are out-of-taxonomy and re redundant. We either enrich taxonomy with these topics as fine-grained subtopics (L4 topics) or as novel topics (new L3 topics).

4.4.1 Syntactic Matching

Let gT be the generated topic and τ' be the set of topics in the taxonomy. We compare gT with each topic in τ' for exact or partial match. If no match is found, we use semantic matching.

$$gT \leftarrow \begin{cases} t & \text{if } gT = t; \quad t \in \tau' \\ t & \text{if } gT \subset t; \quad t \in \tau' \\ gT & \text{otherwise} \end{cases} \quad (9)$$

4.4.2 Semantic Matching

We use a signalling algorithm (refer *BTS* Algorithm 4) to compute the best matching topics, and corresponding scores for each of the generated topic and extracted verbatim. For each topic T_i in the taxonomy topics list τ' , we find the maximum similarity with the generated topic (gT) as:

$$topic_t, score_t = BestTopicAndScore([T_i]_{i=1}^N, [SimST(gT, T_i)]_{i=1}^N) \quad (10)$$

Similarly, for each verbatim k_j in the set of verbatims K_i for each topic T_i , we find the maximum similarity with the extracted verbatim (eV) as:

$$topic_v, score_v = BestTopicAndScore([T_i]_{i=1}^N, [\max_{k \in K_i}(SimST(eV, k))]_{i=1}^N) \quad (11)$$

We use the above scores and a semantic post-processing heuristics (refer Algorithm 5) to mark the generated topic as a new topic (new L3), a fine-grained subtopic (L4) of an existing L3 topic, or an existing L3 topic.

Algorithm 5 Semantic Matching

Require: $(topic_t, score_t), (topic_v, score_v)$

- 1: **if** $score_t > 0.95$ **then**
 - 2: replace generated_topic with taxonomy topic $topic_t$
 - 3: **else if** $score_t > 0.7$ and $score_v > 0.4$ **then**
 - 4: surface the generated_topic as new granular topic (L4)
 - 5: **else**
 - 6: surface as new_topic to be added to the taxonomy
 - 7: **end if**
-

5 Experiments

5.1 Data generation ablation

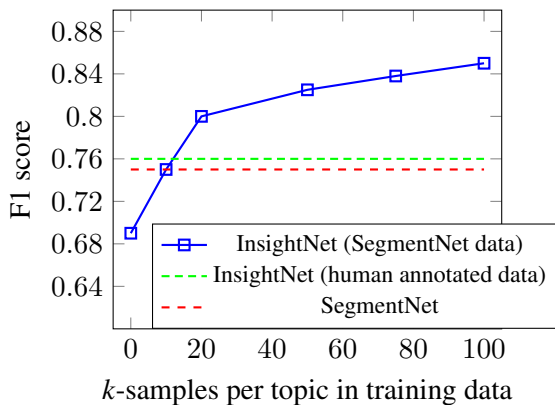


Figure 4: SegmentNet Data Ablation

We show that SegmentNet can generate training data that is better or comparable to human annotated data. Figure 4 compares the performance of InsightNet trained with SegmentNet on different dataset sizes with a *fixed* human annotated dataset

(fixed due to human bandwidth limitations). SegmentNet improves the performance by 9% over human annotated data, given only 30 samples/topic by manual annotation. This limitation is due to the heavy-tailed data (See Appendix section C.1) and the need for more data to cover the underrepresented topics. We also see that we need about three times more synthetic data to surpass the human-annotated baseline. We also show the model performance trained with k samples per topic. We find that InsightNet outperforms SegmentNet around 20 samples per topic and stabilizes around 100 samples per topic.

5.2 Prompt Engineering

The choice of prompt can significantly affect the performance of language models like ours, especially in multi-task settings. We devised different variations of decomposed prompting for our multi-task problem of extracting actionable insights from customer reviews. We experimented with different orders of prompts for verbatim extraction (V), topic identification (T_c and T_g), and polarity detection (P). We also explored different approaches for prompting for topic identification, either top-down (from coarse to granular) or bottom-up (from granular to coarse). We measured the performance of each variation using precision, recall, and F1-score metrics. We discovered that the optimal prompting strategy was to first prompt for the granular topics (T_g) from the review, then prompt for polarity (P) for each topic, and finally prompt for the verbatims (V) that correspond to the topics. This strategy achieved an F1-score of 0.80, which was considerably higher than the other variations. We also observed that using bottom-up prompting for topic identification was more efficient than using top-down prompting, as it minimized the errors in conditional prompting and enhanced the quality of topic extraction. We could deduce the coarse topics (T_c) from the granular topics (T_g) using the taxonomy. We refer to Level-1 and Level-2 topics as coarse topics and Level-3 topics as granular topics. We provide more detailed explanation of experiments in Appendix section A.2.

5.3 AmaT5: Effect of Pre-Training

We applied unsupervised pre-training (Li et al., 2021) to fine-tune a pre-trained model with unlabeled data from the target domain to enhance its transferability. We used the T5-base model (Raffel et al., 2020) with review data (20M raw reviews) and the *i.i.d. noise, replace spans* objectives to do

Model/Approach	Topic Classification (L3 + Polarity)			Verbatim Extraction	
	Precision	Recall	F1 Score	Correctness	Completeness
Multi Level Seq2seq (Liu et al., 2022)	0.34	0.38	0.36	-	-
Rule-based (Rana and Cheah, 2015)	0.56	0.61	0.58	-	-
BERT (ABSA) (Hoang et al., 2019)	0.61	0.67	0.64	-	-
DNNC - NLI (Zhang et al., 2020)	0.76	0.73	0.74	-	-
Aspect Clustering (Sircar et al., 2022)	0.70	0.79	0.74	0.70	0.97
SegmentNet	0.82	0.70	0.76	0.82	0.98
InsightNet	0.85	0.86	0.85	0.85	0.99

Table 1: InsightNet - Baselines

this. We named the resulting model **AmaT5** and it showed better performance than the original. We also tried other variations, like T5-base, along with Sentence Shuffling (see section C.4), BART (Lewis et al., 2019), FlanT5 (Chung et al., 2022) and the results are shown in Table 2.

LLM Checkpoint	Precision	Recall	F1 Score
T5-base	0.79	0.81	0.80
T5-base + Sentence Shuffling	0.80	0.81	0.80
BART	0.71	0.74	0.72
FlanT5	0.81	0.83	0.82
AmaT5	0.85	0.86	0.85

Table 2: InsightNet - Training Ablation

5.4 Experimental Results & Baselines

We conducted a comprehensive evaluation of our proposed methodology across a diverse set of 43 categories, encompassing over 2200+ distinct product types, which collectively represent more than 95% of the global volume of reviews. Our evaluation employed a dataset extracted from Amazon reviews¹. To ensure a robust assessment, both the training and test datasets were meticulously stratified at a granular topic level. Specifically, we used around 75k reviews for training and 10k reviews for testing, thereby ensuring coverage across all product categories and granular topics.

To facilitate an equitable comparison across different approaches, we carried out post-processing procedures, as outlined in Section 4.4. Our findings reveal that our approach outperforms Aspect Clustering (Sircar et al., 2022), a state-of-the-art method for topic extraction, in terms of coverage, diversity, and standardization. Specifically, our approach can generate over 1200+ unique topics that capture both positive and negative aspects of the reviews, while Aspect Clustering produces many redundant topics for the same level of coverage. Moreover, our approach ensures that the topics are consistent and coherent across reviews and product

categories, with only 12% of them being duplicates that can be easily merged in post-processing. On the other hand, Aspect Clustering approach faces some challenges in reducing high duplication rate and no standardization, meaning that many topics are redundant and suffer with duplicate entries.

Table 1 shows the comparison of our results with the baselines. We also observed around 15% new topics have emerged which were not part of taxonomy (detailed analysis in Appendix section A.3).

5.5 Why fine-tuning is required?

In contrast to existing large language models like ChatGPT/OpenAI-GPT3 (Brown et al., 2020), Llama2/Meta (Touvron et al., 2023), Bard/Google-LaMDA (Thoppilan et al., 2022), Falcon/TII (Penedo et al., 2023), which fall short in extracting structured insights from customer reviews due to issues like generating redundant topics, domain-specificity, struggle to distinguish actionable vs non-actionable verbatims, and inefficiencies due large model size and inference latency. Additionally, they lack adherence to taxonomy topics as evidenced by our experiments (see Appendix section A.1 and Table 6 for results), where only 7% of reviews produced correct outputs, 11% reasonable outputs, and the majority, 82%, yielded random results.

6 Conclusion

We have presented InsightNet, a novel multi-task model that extracts granular insights from customer reviews. InsightNet jointly performs multi-topic identification, sentiment classification, and verbatim extraction for each review, generates new topics beyond existing taxonomy, and enriches taxonomy with consistent and exhaustive topics. InsightNet surpasses the state-of-the-art methods by 11% F1-score on overall performance metrics, and achieves 85% F1-score on topic classification. Furthermore, InsightNet is scalable and can handle various tasks with a structured and hierarchical output.

References

- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2009. Multi-facet rating of product reviews. In *Advances in Information Retrieval: 31th European Conference on IR Research, ECIR 2009, Toulouse, France, April 6-9, 2009. Proceedings 31*, pages 461–472. Springer.
- Muhammad Bilal and Abdulwahab Ali Almazroi. 2022. Effectiveness of fine-tuned bert model in classification of helpful and unhelpful online customer reviews. *Electronic Commerce Research*, pages 1–21.
- Samuel Brody and Noemie Elhadad. 2010. An unsupervised aspect-sentiment model for online reviews. In *Human language technologies: The 2010 annual conference of the North American chapter of the association for computational linguistics*, pages 804–812.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Arepalli Peda Gopi, R Naga Sravana Jyothi, V Lakshman Narayana, and K Satya Sandeep. 2023. Classification of tweets data based on polarity using improved rbf kernel of svm. *International Journal of Information Technology*, 15(2):965–980.
- Mickel Hoang, Oskar Alija Bihorac, and Jacobo Rouces. 2019. Aspect-based sentiment analysis using BERT. In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 187–196, Turku, Finland. Linköping University Electronic Press.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177.
- Yin Kang and Lina Zhou. 2017. Rube: Rule-based methods for extracting product features from online consumer reviews. *Information & Management*, 54(2):166–176.
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2023. Decomposed prompting: A modular approach for solving complex tasks.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. Suichan Li, Dongdong Chen, Yinpeng Chen, Lu Yuan, Lei Zhang, Qi Chu, Bin Liu, and Nenghai Yu. 2021. Unsupervised finetuning. *arXiv preprint arXiv:2110.09510*.
- 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.
- Madhumita Guha Majumder, Sangita Dutta Gupta, and Justin Paul. 2022. Perceived usefulness of online customer reviews: A review mining approach using machine learning & exploratory data analysis. *Journal of Business Research*, 150:147–164.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: Outperforming curated corpora with web data, and web data only.
- 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. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Toqir Ahmad Rana and Yu-N Cheah. 2015. Hybrid rule-based approach for aspect extraction and categorization from customer reviews. In *2015 9th International Conference on IT in Asia (CITA)*, pages 1–5. IEEE.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Manuel J Sánchez-Franco, Antonio Navarro-García, and Francisco Javier Rondán-Cataluña. 2019. A naive bayes strategy for classifying customer satisfaction: A study based on online reviews of hospitality services. *Journal of Business Research*, 101:499–506.
- 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.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, et al. 2022. Lamda: Language models for dialog applications.
- Ivan Titov and Ryan McDonald. 2008. Modeling online reviews with multi-grain topic models. In *Proceedings of the 17th international conference on World Wide Web*, pages 111–120.

Hugo Touvron, Louis Martin, Kevin Stone, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).

Jian-Guo Zhang, Kazuma Hashimoto, Wenhao Liu, Chien-Sheng Wu, Yao Wan, Philip S. Yu, Richard Socher, and Caiming Xiong. 2020. [Discriminative nearest neighbor few-shot intent detection by transferring natural language inference](#).

Lili Zheng. 2021. The classification of online consumer reviews: A systematic literature review and integrative framework. *Journal of Business Research*, 135:226–251.

A InsightNet

A.1 Observations on usage of LLMs without fine-tuning

We experimented with different LLMs such as ChatGPT/OpenAI-GPT3, Llama-2/Meta, Bard/Google-LaMDA, Falcon/TII without any fine-tuning. We constructed the prompts using the review and the granular topics list and asked the LLM to predict the topic, polarity and verbatim for each review. The exact sequence of prompts used were:

1. **Topic generation:** *Given the review <>, identify the topics discussed in the review from the list of topics (actionable aspects) in [],*
2. **Polarity generation:** *Given the review <>, identify the polarity for each of these topics (actionable aspects) in [],*
3. **Verbatim extraction:** *Given the review <>, extract the verbatim (review segment) corresponding to the topic-polarity list []*

Table 6 shows the predictions from different LLMs. It is evident that the pre-trained LLMs do not perform well on the specific tasks, even when given the Taxonomy topics as input. In most cases, the predicted topics are not part of the Taxonomy, and are substrings of the segments. Furthermore, the extracted verbatims are non-actionable as some of them have neutral polarity.

Hence, it is not recommended using them in production systems where the stakeholders expect structured and consistent outputs.

A.2 Discussion on Prompt Engineering experiments

We tried different variations of decomposed prompting for our multi-task problem and arrived

at final working prompts. We grouped the prompts into two categories:

1. Hierarchy of topic classification:

(a) **Top-down:** We prompted the model to infer a Coarse-grained Topic (T_c) from the review (R) first. Then, we used the review and the inferred T_c as inputs to prompt the model to generate the corresponding Granular Topic (T_g).

(b) **Bottom-up:** We prompted the model to generate a Granular Topic (T_g) from the review directly. We could derive the coarse topic (T_c) from T_g using the existing taxonomy.

2. Task ordering:

We also experimented with changing the order of the tasks. These are as follows:

(a) Extracting actionable verbatim (V) first and then assigning topics to each verbatim

(b) Generating Topic (T_g) first and then extracting verbatim for each topic

(c) Extracting the polarity (P) first followed by generating topics (T_g) for each polarity (P) followed by extracting verbatim (V)

We used PT_g to denote polarity specific granular topic extraction and PT_c to denote polarity specific coarse topic extraction. We discussed the prompts, observations and conclusions for each experiment or prompting strategy in detail in Table 3.

A.3 Observations on new topic discovery

We analyzed $\sim 10k$ reviews spanning across product categories and found that our model generated $\sim 1450+$ unique topics. Out of these, $\sim 1200+$ topics matched the existing taxonomy, while $\sim 200+$ topics ($\sim 20\%$) were new and emerged from post-processing. From the new topics discovered about $\sim 15\%$ of them were fined-grained subtopics (surfaced as L4 topics) versions of the existing granular (L3) topics, and the rest were completely new topics (surfaced as new L3) identified by InsightNet, which were not present in the base taxonomy. Note: We are not revealing exact numbers to comply with company legal policy.

A.4 Post Processing Heuristics

To ensure the quality and structure of the taxonomy, the post processing heuristics evaluates the scores of the generated topics. It then determines: (1) if a

Order	P/R/F1*	Observations	Next Steps
$R \rightarrow V \rightarrow T_c \rightarrow T_g$	0.21/0.36/0.27	The model could not differentiate between review segments that need action and those that do not, and extracted both types of verbatims. This caused wrong topic assignment to verbatims that are not actionable, leading to poor precision in topic identification.	To prevent this, first identify the topics and then extract the verbatims that match them.
$R \rightarrow T_c \rightarrow T_g \rightarrow V$	0.34/0.38/0.36	The model has difficulty in distinguishing the positive and negative aspects of the review.	Introduce a polarity-based topic identification prompt.
$R \rightarrow (PT_c) \rightarrow T_g \rightarrow V$	0.37/0.51/0.43	The model often identifies topics that are contrary to the prompt’s polarity.	To avoid negative prompts generating positive topics and vice versa, first identify the topics and then assign the polarity.
$R \rightarrow T_c \rightarrow T_g \rightarrow P \rightarrow V$	0.42/0.53/0.47	The accuracy of granular topic classification is affected by the low quality of coarse topic identification. This leads to errors in conditional prompting and poor metrics for granular topic identification.	Since the topics are organized in a hierarchical taxonomy, we can improve the results by starting with granular topic identification instead of coarse topic identification. Also, to reduce the number of prompts, we can directly prompt for polarity-based topic extraction.
$R \rightarrow (PT_g) \rightarrow V$	0.63/0.78/0.70	The model identifies positive topics in negative prompts and negative topics in positive prompts. The model has difficulty in distinguishing between them and many of the topics are identified in both types of prompts, leading to poor precision.	To avoid this, first identify the actionable granular topics and then assign polarity to them.
$R \rightarrow T_g \rightarrow P \rightarrow V$	0.79/0.81/0.80	The model performed well in all three tasks.	Using the taxonomy, we can infer the higher levels: Coarse topics (T_c) and Hinge topics (T_h)

Table 3: Prompt engineering approaches. *Note: The topics are post-processed (refer section 4.4) to match the taxonomy topics before calculating the metrics. The metrics are for the Granular Topic classification task.

topic is new, or (2) if it can be a more granular topic (fine-grained subtopic L4) of an existing L3 topic, or (3) if it can be replaced by a similar topic from the taxonomy. This prevents the redundancy of topics that have the same meaning. Post processing also gives the hierarchical structure to the output. This prevents the redundancy of topics that have the same meaning. Post processing also gives the hierarchical structure to the output.

A.5 InsightNet Sample Predictions

Table 5 shows the sample predictions from InsightNet.

B SegmentNet

B.1 Segmentation Heuristics

We devised heuristics based on linguistic analysis which extracts meaningful phrases from reviews by splitting the text into sentences and then into phrases. Based on our analysis we fixed the minimum length of phrase to be 2 words to make the segment complete and meaningful.

1. **Review** \rightarrow **Sentences**: Split on { . ! ? "but" }
2. **Sentence** \rightarrow **Phrases**:
 - Split sentence on { , ; & "and" }
 - Do no split into phrases if any resulting phrases has ≤ 2 words

C Supplementary Material

C.1 Latency metrics

Latency is measured on a 10MB dataset which contains about 20k reviews. Inference is done with a batch size of 32 on a single *m5.12xlarge* instance for CPU performance and *p3.2xlarge* for GPU performance calculation.

Table 4: Inference latency comparison

	InsightNet	SegmentNet
CPU	5hrs 33mins	5hrs 40 mins
GPU	2hrs 30mins	2hrs 40mins

The aspect mentions in the reviews are heavy tail distributed, since large number of review segments are dominated by minority number of topics. We have plotted the frequency or support for each of 1200+ topics for 6 month review data spanning across all product categories, and plotted the histogram and can be seen in Figure 5.

C.2 Discussion on Taxonomy

Table 8 presents a sample of the hierarchical auto-taxonomy derived from reviews. The auto-taxonomy generated using reviews from 40+ product categories resulted in 8 L1 topics, 600+ L2 topics, and 1200+ L3 topics³.

C.3 SegmentNet examples

Table 7 shows sample data generated using SegmentNet.

C.4 Data Augmentation: Sentence Shuffling

Sentence shuffling is a data augmentation technique that we applied to the labelled reviews. We split each review into sentences based on full-stop and then randomly rearranged the sentences to form a shuffled review. The label of the shuffled review remained the same as the original review. We found that the average number of sentences in a review was between 3 and 6. Therefore, we could generate up to 6 shuffled versions of each review and add them to the training data to increase its size.

D Discussion on Future work

We intend to expand this approach to multilingual and multimodal settings. Furthermore, we plan to extend a model to perform additional tasks, such as summarising review insights at various levels of granularity.

³We are not revealing exact numbers to comply with company legal policy

Review	InsightNet - Model Prediction	InsightNet - Post Processed Output
Clear and wide angle view. Provides a clear view of a wide area.	<pre> [{'visibility', 'positive'}: ['Provides a clear view']], [{'wide angle mode', 'positive'}: ['Clear and wide angle view']] </pre>	<pre> [{'L1 Topic': 'specifications and functionality', 'L2 Topic': 'visibility', 'L3 Topic': 'visibility', 'Polarity': 'positive', 'Verbatims': ['Provides a clear view']}, {'L1 Topic': 'specifications and functionality', 'L2 Topic': 'wide angle mode', 'L3 Topic': 'wide angle mode', 'Polarity': 'positive', 'Verbatims': ['Clear and wide angle view']} </pre>
To much treble no bass. Sound was not good. To much treble no bass sounds tinny.	<pre> [{'audio quality', 'negative'}: ['Sound was not good']], [{'treble', 'negative'}: ['To much treble no bass', 'To much treble no bass sounds tinny']] </pre>	<pre> [{'L1 Topic': 'specifications and functionality', 'L2 Topic': 'audio quality', 'L3 Topic': 'audio quality', 'Polarity': 'negative', 'Verbatims': ['Sound was not good']}, {'L1 Topic': 'specifications and functionality', 'L2 Topic': 'treble', 'L3 Topic': 'treble', 'Polarity': 'negative', 'Verbatims': ['To much treble no bass', 'To much treble no bass sounds tinny']} </pre>
Short inseam. Fit was fine but theyre very short	<pre> [{'in seam fit', 'negative'}: ['Short inseam']], [{'correct fit', 'positive'}: ['Fit was fine']], [{'size smaller than expected', 'negative'}: ['theyre very short']] </pre>	<pre> [{'L1 Topic': 'design and make', 'L2 Topic': 'fit', 'L3 Topic': 'in seam fit', 'Polarity': 'negative', 'Verbatims': ['Short inseam']}, {'L1 Topic': 'design and make', 'L2 Topic': 'fit', 'L3 Topic': 'correct fit', 'Polarity': 'positive', 'Verbatims': ['Fit was fine']}, {'L1 Topic': 'design and make', 'L2 Topic': 'size', 'L3 Topic': 'size smaller than expected', 'Polarity': 'negative', 'Verbatims': ['theyre very short']} </pre>
Not comfortable. The sole inside is very hard	<pre> [{'comfort', 'negative'}: ['Not comfortable']], [{'innersole quality', 'negative'}: ['The sole inside is very hard']] </pre>	<pre> [{'L1 Topic': 'specifications and functionality', 'L2 Topic': 'ease of use comfort', 'L3 Topic': 'comfort', 'Polarity': 'negative', 'Verbatims': ['Not comfortable']}, {'L1 Topic': 'specifications and functionality', 'L2 Topic': 'material quality', 'L3 Topic': 'innersole quality', 'Polarity': 'negative', 'Verbatims': ['The sole inside is very hard']} </pre>
No fishy after taste! I take this to help control cholesterol.	<pre> [{'taste overall', 'positive'}: ['No fishy after taste']], [{'heart health', 'positive'}: ['I take this to help control cholesterol']] </pre>	<pre> [{'L1 Topic': 'specifications and functionality', 'L2 Topic': 'taste', 'L3 Topic': 'taste overall', 'Polarity': 'positive', 'Verbatims': ['No fishy after taste']}, {'L1 Topic': 'health and safety', 'L2 Topic': 'heart health', 'L3 Topic': 'heart health', 'Polarity': 'positive', 'Verbatims': ['I take this to help control cholesterol']} </pre>
It's effective. This is probably the best magnesium supplement, have tried several. It's the best thus far in helping with sleep.	<pre> [{'magnesium', 'positive'}: ['This is probably the best magnesium supplement']], [{'sleep', 'positive'}: ["It's the best thus far in helping with sleep"]] </pre>	<pre> [{'L1 Topic': 'specifications and functionality', 'L2 Topic': 'nutritional value', 'L3 Topic': 'magnesium', 'Polarity': 'positive', 'Verbatims': ['This is probably the best magnesium supplement']}, {'L1 Topic': 'health and safety', 'L2 Topic': 'sleep quality', 'L3 Topic': 'sleep', 'Polarity': 'positive', 'Verbatims': ["It's the best thus far in helping with sleep"]} </pre>
Yes! Arrived nice and intact. No melting. No old taste. Good choice!	<pre> [{'shipment or delivery', 'positive'}: ['Arrived nice and intact']], [{'storage temperature', 'positive'}: ['No melting']], [{'taste overall', 'positive'}: ['No old taste']] </pre>	<pre> [{'L1 Topic': 'shipment package and delivery', 'L2 Topic': 'delivery experience', 'L3 Topic': 'shipment or delivery', 'Polarity': 'positive', 'Verbatims': ['Arrived nice and intact']}, {'L1 Topic': 'specifications and functionality', 'L2 Topic': 'storage temperature', 'L3 Topic': 'storage temperature', 'Polarity': 'positive', 'Verbatims': ['No melting']}, {'L1 Topic': 'specifications and functionality', 'L2 Topic': 'taste', 'L3 Topic': 'taste overall', 'Polarity': 'positive', 'Verbatims': ['No old taste']} </pre>
Very pleased. I love this aloe. It's 100% organic, feels very nice on skin. I'll use this for a variety of things.	<pre> [{'organic', 'positive'}: ["It's 100% organic"]], [{'skin health', 'positive'}: ['feels very nice on skin']] </pre>	<pre> [{'L1 Topic': 'specifications and functionality', 'L2 Topic': 'nutritional value', 'L3 Topic': 'organic', 'Polarity': 'positive', 'Verbatims': ["It's 100% organic"]}, {'L1 Topic': 'health and safety', 'L2 Topic': 'health benefits', 'L3 Topic': 'skin health', 'Polarity': 'positive', 'Verbatims': ['feels very nice on skin']} </pre>
Very sturdy and shreds well. Very sturdy solid machine. Can handle 3-4 pages at a time. Shreds into small bits!!	<pre> [{'shredding quality', 'negative'}: ['Shreds into small bits']], [{'strong', 'positive'}: ['Very sturdy solid machine']], [{'paper quality', 'positive'}: ['Can handle 3-4 pages at a time']] </pre>	<pre> [{'L1 Topic': 'specifications and functionality', 'L2 Topic': 'shredding quality', 'L3 Topic': 'shredding quality', 'Polarity': 'negative', 'Verbatims': ['Shreds into small bits']}, {'L1 Topic': 'specifications and functionality', 'L2 Topic': 'durability longevity', 'L3 Topic': 'strong', 'Polarity': 'positive', 'Verbatims': ['Very sturdy solid machine']}, {'L1 Topic': 'specifications and functionality', 'L2 Topic': 'material quality', 'L3 Topic': 'paper quality', 'Polarity': 'positive', 'Verbatims': ['Can handle 3-4 pages at a time']} </pre>

Table 5: Sample predictions from InsightNet

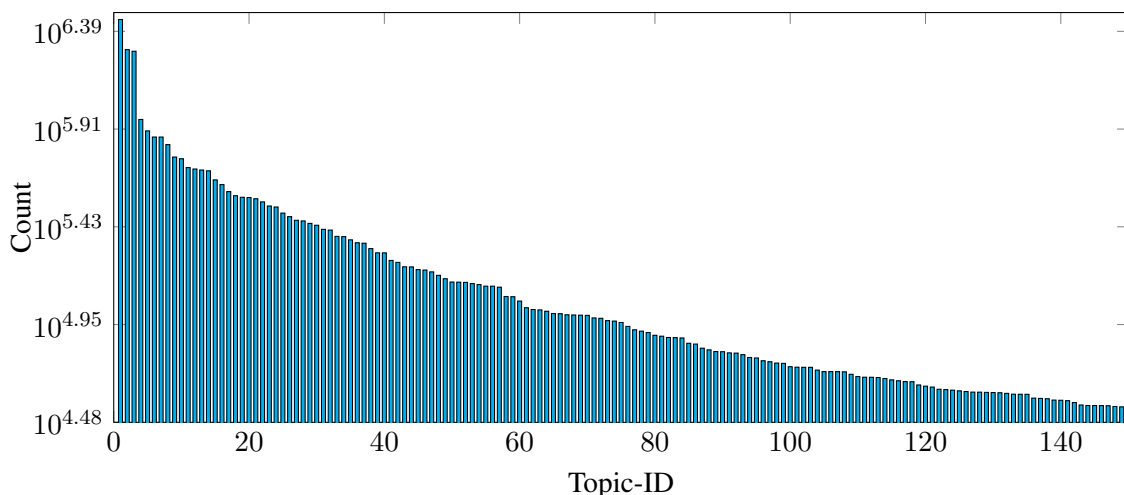


Figure 5: Heavy tailed distribution of topics

Review	LLM	Prediction
This has been a horrible experience for me. It is a 5 cup coffee maker and I was attempting to make 4,it over flowed and has taken nearly 4 hours to brew one cup of coffee	GPT-3	{ "topics": ["overflowed", "brew one cup of coffee"], "verbatim": "It over flowed and has taken nearly 4 hours to brew one cup of coffee", "polarity": "Negative" }
The opening of the bag is square, but the lid is non-existent;it's a zipper straight across the top (between two sides),which means the bag holds a lot less than it could. Disappointed.The bags were in their original flimsy plastic wrap. I don't recommend these bags.	GPT-3	{ "topic": "opening of the bag", "verbatim": "The opening of the bag is square, but the lid is non-existent; it's a zipper straight across the top (between two sides), which means the bag holds a lot less than it could.", "polarity": "negative" }
Very small. While they were cute, I had to give them to my 8 year old granddaughter. She's.the only one who could.wear them.	Llama-2	{ "topics": ["age suitability", "fit", "quality"], "verbatim": "Very small. While they were cute, I had to give them to my 8 year old granddaughter. She's the only one who could wear them.", "polarity": "neutral" }
Looks good, nice material, very tight in shoulders. I loved the material, both inside and outside. It's comfortable, but somewhat loud (like a windbreaker). It's very warm and lots of useful pockets, but it was way too tight in the shoulders to have range of motion.	Llama-2	{ "topics": ["material quality", "comfort", "warmth", "breathability", "range of motion"], "verbatim": "Looks good, nice material, very tight in shoulders. I loved the material, both inside and outside. It's comfortable, but somewhat loud", "polarities": ["positive", "positive", "neutral", "neutral", "negative"] }
Very small. While they were cute, I had to give them to my 8 year old granddaughter. She's.the only one who could.wear them.	Bard	{ "review": "Very small. While they were cute, I had to give them to my 8 year old granddaughter. She's the only one who could wear them.", "topics": ["incorrect fit"], "verbatim": "Very small. While they were cute, I had to give them to my 8 year old granddaughter. She's the only one who could wear them.", "polarity": "negative" }
Looks good, nice material, very tight in shoulders. I loved the material, both inside and outside. It's comfortable, but somewhat loud (like a windbreaker). It's very warm and lots of useful pockets, but it was way too tight in the shoulders to have range of motion.	Bard	{ "review": "Looks good, nice material, very tight in shoulders. I loved the material, both inside and outside. It's comfortable, but somewhat loud (like a windbreaker). It's very warm and lots of useful pockets, but it was way too tight in the shoulders to have range of motion", "topics": ["fit", "material", "noise"], "verbatims": ["very tight in shoulders", "loved the material, both inside and outside", "somewhat loud (like a windbreaker)"], "polarities": ["negative", "positive", "neutral"] }
This has been a horrible experience for me. It is a 5 cup coffee maker and I was attempting to make 4,it over flowed and has taken nearly 4 hours to brew one cup of coffee	Falcon	{ "topics": ["brew one cup of coffee"], "verbatim": "It over flowed and has taken nearly 4 hours to brew one cup of coffee", "polarity": "negative" }
The opening of the bag is square, but the lid is non-existent;it's a zipper straight across the top (between two sides),which means the bag holds a lot less than it could. Disappointed.The bags were in their original flimsy plastic wrap. I don't recommend these bags.	Falcon	{ "topic": "square", "verbatim": "The opening of the bag is square, but the lid is non-existent", "polarity": "neutral" }

Table 6: Prediction using pre-trained LLMs

Review	Segment	Polarity	Matched topic
Not even close. not even close to the same as the image.	Not even close	negative	no topic
	not even close to the same as the image	negative	false advertising
Color is GREAT! Have to battle the sleeve tightness. Length is great. Warmth is there. Just very tight in the arm area. Not shoulders but sleeves	Color is GREAT!	positive	color
	Have to battle the sleeve tightness	negative	size smaller than expected
	Length is great	positive	correct size
	Warmth is there	positive	warmth
	Just very tight in the arm area	negative	arm fit
	not the shoulders	neutral	-
sleeves	neutral	-	

Table 7: SegmentNet: Data Generation Examples

Coarse Topic	Hinge Topic	Granular Topic	Polarity	Keywords
customer service	responsiveness	great responsiveness	positive	replied fast, immediate response ...
customer service	responsiveness	unable to reach support	negative	no response, can't reach vendor ...
design and make	size	correct size	positive	size as expected, true to size ...
design and make	size	size larger than expected	negative	too long, bigger than expected ...
design and make	size	size smaller than expected	negative	too short, XXL fits like an L ...
health and safety	sleep quality	sleep quality	negative	not helpful for sleep, poor sleep assistance ...
health and safety	sleep quality	sleep quality	positive	sleep quality improves, good for active sleepers ...
specifications and functionality	material quality	zipper quality	negative	zipper sticks, does not zip well ...
specifications and functionality	material quality	zipper quality	positive	unzips smoothly, easy to zip ...
returns refunds and replacements	policies and initiation	cannot initiate returns	negative	no option to return, outside of return policy ...
returns refunds and replacements	policies and initiation	unclear policies	negative	no return policy, bad replacement policy ...
shipment package and delivery	packaging	good packaging	positive	safe and secure packaging, pleased with packaging quality ...
shipment package and delivery	packaging	package damaged	negative	box arrived crushed, package arrived with dents, envelope ripped open ...
shipment package and delivery	packaging	redundant packaging	negative	too much plastic in package, arrived with too many boxes, ...
shipment package and delivery	packaging	unhygienic packaging	negative	package has stains, arrived wet and soggy
used damaged expired	new used product	new product	positive	brand new product, condition is new ...
used damaged expired	new used product	refurbished product	negative	refurbished sent, it is clearly refurbished ...
used damaged expired	new used product	used product	negative	Has been used previously, was sent a used product ...
miscellaneous	advertising related	as advertised claimed	positive	works as advertised, specs match the description
miscellaneous	advertising related	false advertising	negative	pictures are deceiving, product different than expected

Table 8: Sample Taxonomy