

# CATCH: A Controllable Theme Detection Framework with Contextualized Clustering and Hierarchical Generation

Rui Ke<sup>1</sup>, Jiahui Xu<sup>1</sup>, Shenghao Yang<sup>1</sup>, Kuang Wang<sup>1</sup>, Feng Jiang<sup>2\*</sup>, Haizhou Li<sup>1,3,4</sup>

<sup>1</sup>Shenzhen Research Institute of Big Data, The School of Data Science, The Chinese University of Hong Kong, Shenzhen

<sup>2</sup>Artificial Intelligence Research Institute, Shenzhen University of Advanced Technology

<sup>3</sup>The School of Artificial Intelligence, The Chinese University of Hong Kong, Shenzhen

<sup>4</sup>Department of Electrical and Computer Engineering, National University of Singapore  
jiangfeng@suat-sz.edu.cn

## Abstract

Theme detection is a fundamental task in user-centric dialogue systems, aiming to identify the latent topic of each utterance without relying on predefined schemas. Unlike intent induction, which operates within fixed label spaces, theme detection requires cross-dialogue consistency and alignment with personalized user preferences, posing significant challenges. Existing methods often struggle with sparse, short utterances for accurate topic representation and fail to capture user-level thematic preferences across dialogues. To address these challenges, we propose CATCH (Controllable Theme Detection with Contextualized Clustering and Hierarchical Generation), a unified framework that integrates three core components: (1) context-aware topic representation, which enriches utterance-level semantics using surrounding topic segments; (2) preference-guided topic clustering, which jointly models semantic proximity and personalized feedback to align themes across dialogue; and (3) a hierarchical theme generation mechanism designed to suppress noise and produce robust, coherent topic labels. Experiments on a multi-domain customer dialogue benchmark (DSTC-12) demonstrate the effectiveness of CATCH with 8B LLM in both theme clustering and topic generation quality.

## Introduction

In real-world customer service domains such as banking, finance, travel, and insurance, accurately identifying the underlying theme of each user utterance is essential for enhancing service efficiency, understanding user needs, and retrieving contextually relevant knowledge. Unlike intent induction (Gung et al. 2023), which typically maps utterances to a predefined label space (Pu et al. 2022; Costa et al. 2023), theme detection aims to uncover latent and potentially novel topics without prior knowledge. Effective theme detection requires preliminary precise topic assignment within a single dialogue (Nguyen et al. 2022; Du, Buntine, and Johnson 2013a), but more importantly should be consistent across multiple dialogues and align with user preferences (Mendonça et al. 2025), which regularizes inter-dialogue theme consolidation, as illustrated in Figure 1.

\*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

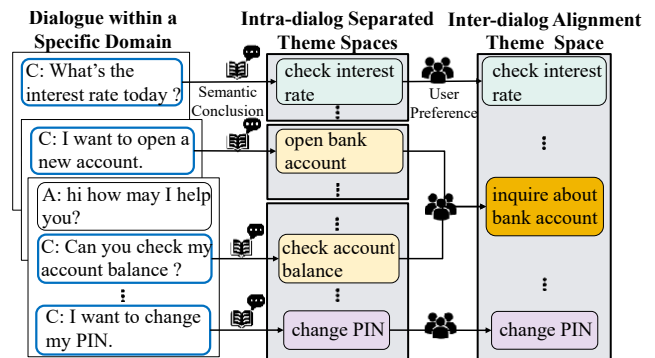


Figure 1: Illustration of the controllable theme detection task. Given a set of dialogues with unlabeled utterances, a theme is generated for each utterance. The theme granularity is influenced by auxiliary inputs such as user preferences (Mendonça et al. 2025).

These challenges underscore the need for models that can generalize beyond surface-level semantics and adapt to diverse real-world conversational scenarios.

However, existing approaches fail to address the real-world controllable theme detection for three key challenges. First, short utterances often lead to sparse and ambiguous semantic signals, making it difficult for conventional topic modeling methods (Blei, Ng, and Jordan 2003; Pham et al. 2024) to construct reliable topic representations. Second, while topic clustering methods (Chatterjee and Sengupta 2020; Gung et al. 2023) group utterances based on surface-level semantics, they typically overlook user-specific preferences, resulting in inconsistent clustering across dialogues even when the underlying intent is similar. Moreover, most previous work lacks a structured and controllable theme generation mechanism (Perkins and Yang 2019; Zeng et al. 2021), causing the generated topic labels to vary arbitrarily between contexts and limiting their applicability in downstream applications.

To address these challenges, we propose **CATCH** (Controllable And Thematic Clustering with Hierarchy), a controllable theme detection framework that combines intra-dialogue context modeling with inter-

dialogue user preference alignment. Specifically, CATCH comprises three key components: (1) a context-aware topic representation module that enhances semantic expressiveness by leveraging dialogue-level topic segmentation; (2) a preference-guided clustering mechanism that jointly models semantic similarity and user-specific signals to ensure thematic consistency across dialogues; and (3) a hierarchical theme generation process, inspired by Chain-of-Thought prompting and stabilized via majority voting, that yields robust, structured, and domain-adaptive topic labels.

Experimental results on the Controllable Dialogue Theme Detection benchmark (DSTC-12) show that our CATCH consistently outperforms strong baselines in both in-domain and cross-domain settings, even under a few preference supervisions. Notably, our framework ranks **second** in the official blind evaluation, achieving strong performance in both automatic and human assessments with a lightweight design (about an 8B-level model). Extensive ablation studies and qualitative analyses further confirm the effectiveness and generalizability of each component. The main contributions of this work are as follows:

- We propose **CATCH**, a novel controllable theme detection framework that integrates topic segmentation, preference-enhanced clustering, and hierarchical label generation. It enables fine-grained and user-adaptive theme modeling across dialogues, bridging structured representation with user-centered controllability.
- We introduce a contextualized clustering mechanism that jointly models intra-dialogue semantic flow and inter-dialogue user preference patterns, thereby aligning topic grouping not only with content coherence but also with personalized thematic expectations.
- We propose a hierarchical theme generation strategy guided by controllable LLM prompting, where local semantic candidates are first derived within fine-grained clusters, and subsequently consolidated into coherent themes through consensus-driven aggregation, mitigating label drift and enhancing cross-context consistency.
- Our framework achieves **second place** in the DSTC-12 Controllable Theme Detection track under both automatic and human evaluations, with a lightweight design and high sample efficiency. Extensive ablation studies and qualitative analyses confirm its effectiveness, particularly in low-resource and cross-domain settings.

## Related Work

### Intra-dialogue Theme Detection

Intra-dialogue theme detection focuses on identifying the thematic affiliation of each utterance within a single conversation. It generally involves two subtasks: *topic segmentation* and *topic generation*.

**Topic segmentation.** Dialogue Topic Segmentation aims to divide a dialogue into semantically coherent segments. Early unsupervised approaches relied on surface-level signals such as word co-occurrence (Hearst 1997; Eisenstein and Barzilay 2008) or topic distributions (Riedl and Bieemann 2012; Du, Buntine, and Johnson 2013b). More recent methods construct contrastive utterance-pair data and

fine-tune pretrained models like BERT (Vuong et al. 2025; Devlin et al. 2019; Xing and Carenini 2021; Gao et al. 2023) and LLMs (Fan et al. 2024; Xu et al. 2025). However, these methods apply uniform segmentation strategies across datasets with different granularities, often leading to inconsistent performance due to domain mismatch and dialogue fragmentation.

**Topic generation.** Topic generation methods, often based on probabilistic or neural topic modeling (Blei, Ng, and Jordan 2003; Pham et al. 2024), aim to infer high-level abstract themes from dialogue utterances. Yet, the brevity and sparsity of dialogue turns pose challenges for robust topic inference (Bach et al. 2021; Lin et al. 2024). To address this, several methods attempt to augment short texts via auxiliary signals (Nguyen et al. 2022; Tuan et al. 2020; Jiang et al. 2024). Despite their success in extracting global themes, most existing methods fail to ensure intra-dialogue consistency or control the granularity of the generated labels.

### Inter-dialogue Theme Detection

Inter-dialogue theme detection mainly concerns clustering and aligning topics across multiple dialogues, which is critical to understanding user intent on scale and constructing reusable dialogue themes.

A common strategy is to cluster utterances based on their semantic representations (Nguyen et al. 2025; Grootendorst 2022; Zhang et al. 2022; Sia, Dalmia, and Mielke 2020). While these methods are efficient and scalable, they often assume a fixed set of topics and do not support theme discovery (Perkins and Yang 2019). Other works explore the complexity of theme space via multiview clustering (Vu et al. 2025; Perkins and Yang 2019), density-based methods (Chatterjee and Sengupta 2020), or structured prediction signals (Liu et al. 2021; Zeng et al. 2021). However, most of these approaches focus solely on semantic similarity and overlook the influence of user preferences on thematic variation. This leads to inconsistent topic grouping across dialogues, even when the underlying intent is similar.

Furthermore, none of the above approaches provides explicit or controllable mechanisms for theme label generation, which limits their applicability in settings that require interpretable and stable outputs across domains.

## Methodology

### Task Definition

We define the controllable theme detection (TD) task as a structured theme generation problem over dialogue utterances. Given a set of utterances  $U = \{u_1, \dots, u_m\}$  extracted from dialogues of a specific domain with low resource user preference annotation (as shown in Table 1), the goal is to assign each utterance  $u_i \in U$  a theme label  $L_i$  that is both preference-aligned and contextually consistent across dialogues. Formally, the output of the theme detection process is a finite family of theme labels  $\mathbb{L} = \{L_1, \dots, L_n\}$  aligned with user preferences  $P$ , where each label corresponds to a subset of utterances:

$$\{L_1 : \{u_j, \dots, u_k\}, \dots, L_n\} \leftarrow TD(U | P) \quad (1)$$




Figure 2: The overall workflow of CATCH.

To achieve this goal, as illustrated in Figure 2, we propose **CATCH**, a controllable theme detection framework that incorporates both intra- and inter-dialogue modeling <sup>1</sup>.

### Context-Aware Intra-Dialogue Theme Representation

To address the semantic sparsity and ambiguity commonly observed in short utterances, we design a context-aware intra-dialogue theme representation module. It leverages a dual-branch topic segmentation framework to infer latent segment boundaries and construct thematically coherent spans. Each utterance representation is then enriched by aggregating contextual information from its corresponding segment, enabling more robust and consistent theme encoding across dialogue turns.

**Dual-branch Topic Segmentation Framework.** Inspired by DialSTART (Gao et al. 2023), our topic segmentation module adopts a dual-branch structure composed of two parallel encoders: (1) a topic encoder to measure topical similarity between utterances in a dialogue and (2) a coherence encoder to model discourse continuity between utterance-pairs. These two components collaborate to detect topic boundaries in a multi-turn dialogue. Specifically, for each adjacent utterance-pair  $(u_i, u_{i+1})$ , the segmenter computes a boundary score  $r_i$  by aggregating the outputs of both encoders:

$$B = \{r_i = t_i + c_i \mid i = 1, \dots, n - 1\} \quad (2)$$

Here,  $t_i$  is the cosine similarity between the topic representations  $\mathbf{h}_i$  and  $\mathbf{h}_{i+1}$  from the Topic Encoder, while  $c_i$  is the coherence score generated by the Coherence Encoder.

**Contrastive Objectives for Encoder-Specific Learning.** To optimize the segmenter, we formulate two contrastive losses based on the outputs of the two encoders. These losses enable the model to distinguish thematically or coherently similar utterances from unrelated ones.

(1) **Topic Encoder Contrastive Loss  $L_1$ :** This loss supervises the Topic Encoder to promote topical consistency in a dialogue. For each utterance  $u_i$ , positive samples  $S_i$  are drawn from a window of neighboring utterances, and nega-

tive samples  $S_i^*$  are sampled from outside this window:

$$L_1(u_i) = \sum_{u^+ \in S_i} \sum_{u^- \in S_i^*} \max(0, \eta + e_i^- - e_i^+) \quad (3)$$

where  $e_i^+$  and  $e_i^-$  represent the similarity scores between  $u_i$  and the positive ( $u^+$ ) and negative ( $u^-$ ) samples, respectively, and  $\eta$  is a margin hyperparameter.

(2) **Dual Encoder Contrastive Loss  $L_2$ :** This loss supervises both the Coherence and Topic Encoder to capture discourse-level continuity in different dialogues. For each utterance-pair  $P_i = (u_i, u_{i+1})$ , we define positive utterance-pair segments  $\mathbb{P}_i$  (neighboring segment pairs) and negative pairs  $\mathbb{P}_i^*$  (randomly sampled from the other dialogues). The loss is defined as:

$$L_2(P_i) = \sum_{P^+ \in \mathbb{P}_i} \sum_{P^- \in \mathbb{P}_i^*} \max(0, \eta + r_i^- - r_i^+) \quad (4)$$

with the segment-level relevance score  $r_i^\pm$  computed by  $P_i = (u_i, u_{i+1})$  and  $P^\pm = (u, v)^\pm$ :

$$r_i = \text{sim} \left( \frac{h_i + h_{i+1}}{2}, \frac{h^u + h^v}{2} \right) + c_i \quad (5)$$

where  $h$  is the topic representations of utterance provided by Topic Encoder, and the  $\pm$  of  $r_i$  is the same with  $\pm$  of  $P$ .

The final optimization combines both encoder-specific contrastive learning objectives:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N L_1(u_i) + \frac{1}{M} \sum_{i=1}^M L_2(P_i) \quad (6)$$

where  $N$  and  $M$  denote the number of samples for each loss.

**Two-Stage Adaptation Strategy.** We further enhance the model’s ability to generalize and align with the target domain through a two-stage adaptation process:

**Stage 1: Conversation-Level Adaptation.** This stage aims to adapt the model to the target dialogue domain by constructing training samples from every utterance  $u_i$  in the target dataset. For each  $u_i$ , both the Topic Encoder and Coherence Encoder are supervised using the aforementioned contrastive objectives  $L_1$  and  $L_2$ .

**Stage 2: Utterance-Level Adaptation.** To further specialize the model for real-world usage scenarios, we conduct fine-grained adaptation based on utterances annotated with

<sup>1</sup><https://github.com/SUAT-AIRI/CATCH>

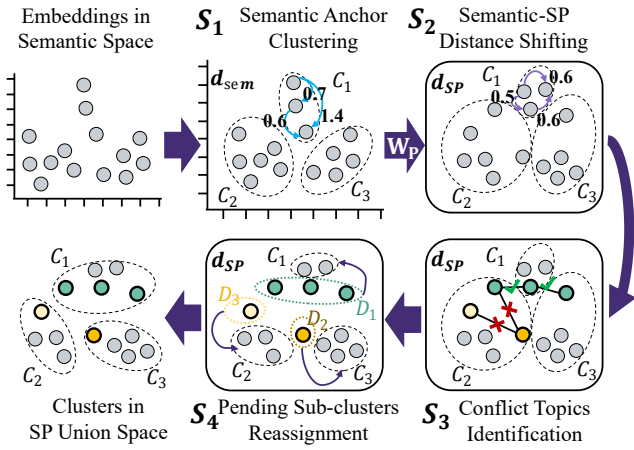


Figure 3: Illustration of the Preference-Enhanced Topic Clustering Algorithm. “✓” denotes a strong linking preference, while “✗” denotes a strong splitting preference.

themes. We assume these annotated utterances are treated as likely indicators of topic boundaries and closer to the real-world using scenarios.

After adaptation, we apply the TextTiling algorithm (Hearst 1997) on the computed boundary scores  $B$  to segment the dialogue. Given a target utterance, we extract its surrounding segment as its contextual block and use it to construct a semantically grounded theme representation  $T$ , which serves as the input to the following topic clustering.

### Preference-Enhanced Topic Clustering

To acquire the theme family  $\mathbb{L}$ , we aim to group similar topics across different dialogues into coherent theme clusters leveraging on their theme representations  $T$  obtained by the topic segmentor. To this end, we propose a preference-enhanced topic clustering algorithm that jointly considers semantic similarity and user preference signals, guided by a trained preference reward model (PRM).

We first design a Semantic-Preference (SP) distance kernel to dynamically fuse semantic similarity with user preference. For each topic pair  $(x, y)$ , the SP is defined as:

$$d_{SP}(x, y) = w_{x,y} \cdot d_{sem}(x, y) \quad (7)$$

where  $d_{sem}$  denotes the Euclidean distance between topic embeddings, and  $w_{x,y} = PRM(x, y)$  is a preference scalar produced by the reward model to indicate the degree to which two topics belong to the same theme.

Because the true joint space combining semantic and preference information is latent and not explicitly constructed, we propose a four-step algorithm grounded in semantic space but progressively incorporating preference signals, as illustrated in Figure 3.

- **Step 1 (S1):** Topics are clustered solely based on semantic similarity to form initial anchor clusters, such as  $C_I$ .
- **Step 2 (S2):** The mutual SP distance  $d_{SP}$  of topics are calculated and assigned to substitute semantic distance

---

### Algorithm 1: Hierarchical Theme Generation

---

**Data:**  $K$  topic clusters  $C_K$ , where each  $C_i \in C_K$  containing topic segment  $P_j$  as elements; Prompt-based LLM  $LLM(prompt|input)$

**Result:** Themes of topic clusters  $\{T_K\}$

```

1 for each  $C_i$  do
2   Randomly split  $C_i$  into several groups  $S_N$  with
    $S_n = \{P_1^n, \dots, P_{10}^n\}$  and  $\cup_{n \in N} S_n = C_i$ ;
3   for each  $S_n$  do
4      $l_n \leftarrow LLM(\text{Label}|S_n)$ ;
5   end
6   Initialize the cleaned labels set  $L_i$ ;
7    $l^{core} \leftarrow LLM(\text{VoteCore}|\{l_n\}, n \in N)$ ;
8   for each  $l_n$  do
9     if  $LLM(\text{IsRelevant}|\{l_n, l^{core}\})$  then
10       $L_i.append(l_n)$ ;
11    end
12  end
13   $T_i \leftarrow LLM(\text{Conclude}|L_i)$ ;
14 end

```

---

$d_{sem}$  by multiplying Preference scalars, such as  $W_p$  delegating the full matrix of  $w_{x,y}$ .

- **Step 3 (S3):** Preference scalars are further used to detect *conflict* topic pairs whose preference signals contradict their current cluster assignments. These conflict topics are then extracted and re-clustered in the SP distance space to form new pending sub-clusters, such as  $D_I$ .
- **Step 4 (S4):** Each pending sub-cluster is reassigned to the most compatible anchor cluster by minimizing the aggregated SP distance, thereby producing final preference-consistent clusters.

Furthermore, since annotating user preferences across all topic pairs is costly and inefficient, we train a preference reward model to estimate human preferences between utterances. The model is trained on a small number of preference-aligned samples ( $S$ ), each labeled as either belonging to the same theme or not, and is used to assign compatibility scores to topic pairs.

$$S = \{((t_i, t_j), y_{ij}) \mid y_{ij} \in \{0, 1\}\} \quad (8)$$

We regard the Preference Reward Model (PRM) as a classification model that predicts a continuous preference score  $w_{ij} = PRM(t_i, t_j) \in [0, 1]$  indicating the degree of linking tendency between the pair. The model is optimized using mean squared error (MSE) loss:

$$\mathcal{L}_{PRM} = (w_{ij} - y_{ij})^2 \quad (9)$$

### Hierarchical Theme Label Generation

In this final step, each topic cluster  $C_i$  is assigned a structured label  $T_i$ . To support downstream applications requiring domain-adaptive and interpretable labeling, we employ a prompt-based large language model (LLM) to generate concise, actionable theme names in the form of verb phrases that reflect user intent.

Type	Domain	# Dialog.	# Utterance	# Preference
Offline	Banking	1634	58418 (980)	164/164
Offline	Finance	1725	196764 (3000)	173/173
Offline	Insurance	836	60352 (1333)	155/126
Online	Travel	765	72010 (999)	77/77

Table 1: Data Statistics of the DSTC-12 Dataset. The numbers in parentheses indicate the number of sampled utterances with annotated themes. In the *Preference* column, the values denote the number of should-link / cannot-link utterance pairs, respectively.

As illustrated in Algorithm 1, the three-stage hierarchical generation framework, integrating fine-grained local labeling with semantic consolidation, is an autonomous data negotiation process allowing automatic theme emergence.

**Intra-group Labeling.** Following a divide-and-conquer strategy, each topic cluster is partitioned into smaller groups to satisfy LLM input length constraints and enhance generation quality. The diversity across groups also helps retain semantic variation within the cluster. Then, each group is labeled through the LLM prompt (*IsRelevant*).

**Noisy Label Cleaning.** Group-level labels may include inconsistent or noisy descriptions. To filter these, we prompt LLM to identify a core label  $l^{core}$  via majority voting (*VoteCore*) which selects the most frequent group label, and retain only those labels deemed semantically relevant (*IsRelevant*) to  $l^{core}$ , resulting in a cleaned label set  $L_i$ .

**Inter-group Theme Consolidation.** Finally, the cleaned label set  $L_i$  is summarized using a consolidation prompt (*Conclude*) to yield the final theme label  $T_i$ , which captures the dominant intent of the cluster while maintaining semantic consistency across utterances.

## Experiments

### Experimental Settings

**Datasets.** We conduct experiments on the multi-domain customer support dialogue datasets (Banking, Finance, Insurance, and Travel) provided by DSTC-12 (Mendonça et al. 2025), as summarized in Table 1. Each dataset contains two key types of annotations of the themed utterances: (1) Theme Annotation: Each target utterance is annotated with its corresponding theme label. (2) Preference Annotation: A binary relation (*should-link* and *cannot-link*) of a pair of target utterances indicating whether they should be grouped under the same theme (should-link) or not (cannot-link).

In the offline evaluation, we use the Banking dataset for training and the Finance and Insurance domains as the valid dataset. For online evaluation, we deploy our model to predict theme labels on the Travel dataset, which lacks golden annotations. The predicted results are submitted to the organizers of DSTC-12 for blind evaluation. Throughout training, CATCH is trained solely based on preference annotations without accessing the ground-truth theme labels.

**Metrics.** To comprehensively evaluate the effectiveness of CATCH, we follow the DSTC-12 (Mendonça et al. 2025) evaluation protocol, which assesses two core aspects: (1) the

quality of theme segmentation (i.e., utterance clustering), and (2) the quality of generated theme labels.

**Offline Evaluation.** For theme segmentation quality, we use two standard clustering metrics: **Normalized Mutual Information (NMI)** (Vinh, Epps, and Bailey 2010), which quantifies the mutual dependence between predicted and reference clusters normalized by their entropies, and **Clustering Accuracy (Acc)**, computed via the Hungarian algorithm to align clusters optimally. For theme label quality, we evaluate the semantic and textual correspondence between predicted and reference labels using: **Cosine Similarity (Cos-Sim)** based on Sentence-BERT embeddings, **ROUGE** (Lin 2004) for n-gram overlap, and an **LLM-based score** that assesses label format and informativeness via vicuna-13B evaluation guided by human-crafted criteria.

**Online Evaluation.** For the held-out test set without golden labels, the DSTC-12 organizers perform additional evaluations including both automatic metrics and manual human judgments.

**Baselines.** We compare CATCH against three representative baselines: (1) **GRand**: the official DSTC-12 baseline, which performs utterance-level clustering using randomly sampled should-link and cannot-link pairs, followed by direct theme generation. (2) **GRM**: an enhanced variant of GRand that conducts topic-level clustering and uses a preference reward model to propagate preference annotations to all topic pairs before generating labels. (3) **GSP**: a variant of GRM that employs the SP distance kernel for clustering without adopting the two-step algorithm.

### Implementation Details

In the topic segmentation stage, we follow the previous work (Gao et al. 2023), using *bert-base-uncased* and *sup-simcse-bert-base-uncased* as our coherence encoder and topic encoder, respectively. During the pre-training and fine-tuning process, we both set the learning rate to  $5e-6$  and the epoch to 3. The relevance threshold is set to be 0.5 during boundary inference.

In the preference-enhanced topic clustering, we employ *all-mpnet-base-v2* to obtain the sentence transformer embeddings and use **UMAP** to reduce embedding dimension. For semantic clustering, we employ **Spectrum** clustering method with the default cluster number  $K$  being 30, following the common design. As for PRM, we use *bert-base-uncased* as the default reward model with a learning rate of  $2e-5$  and an epoch of 3. During preference inference, we set the tendency threshold of linking  $\theta_l$  to be 0.85 and the tendency threshold of splitting  $\theta_s$  to be 0.15.

In the theme label generation, we employ *LLaMA3-8B-Instruct* as the default LLM for prompt generation. The number of topics in each subgroup is set to be 25.

In the evaluation part, offline LLM-based evaluation is conducted on Vicuna-13B as DSTC-12 suggests, while the online test result is provided by DSTC-12.

The computational requirement of three stages is shown in Table 3. CATCH is a lightweight design that is only trained on models with less than 500M parameters in total, and uses all substitutable open-source models.

Method	Finance					Insurance				
	Clustering Metrics		Theme Label Quality			Clustering Metrics		Theme Label Quality		
	Acc	NMI	Rouge-1/2/L	CosSim	LLM-Score	Acc	NMI	Rouge-1/2/L	CosSim	LLM-Score
GRand	24.6	28.2	5.0 / 3.5 / 5.0	13.8	87.0	41.5	42.2	12.3 / 0.0 / 12.3	47.8	86.6
GRM	39.1	51.5	21.6 / 6.4 / 21.1	42.8	82.9	39.6	51.7	27.1 / 8.4 / 26.2	<b>57.5</b>	96.5
GSP	23.3	28.0	19.1 / 4.1 / 19.0	48.5	85.8	23.5	30.1	20.8 / 8.3 / 20.7	44.6	87.2
<b>CATCH</b>	<b>55.8</b>	<b>67.1</b>	<b>42.4 / 24.5 / 42.4</b>	<b>59.3</b>	<b>97.3</b>	<b>54.5</b>	<b>62.6</b>	<b>41.8 / 16.1 / 41.8</b>	57.0	<b>100.0</b>

Table 2: Performance comparison on Finance and Insurance domains, grouped by theme distribution and theme label quality. All values are in percentages.

Stage	Model	Size	Training Cost
TopSeg	2*Bert-base	220M	2*GPU-h
PeC	all-mpnet-base-v2	110M/109M	0.05/0*GPU-h
HieGen	LLaMA-3-8B	8B	0*GPU-h

Table 3: Computational cost (4090). TopSeg: Topic segmentation stage; PeC: preference-enhanced clustering stage; HieGen: hierarchical theme generation stage.

## Offline Experimental Results

We train CATCH on the Banking dataset and conduct the experiment in two data scenarios: in-domain data, out-of-domain data. In the in-domain task, we test on Banking set. For the out-of-domain task, we respectively evaluate the model on the Finance and Insurance sets.

Method	Acc	NMI	Rouge-1/2/L	Cos	LLM
GRand	36.8	33.4	11.1 / 2.9 / 11.1	30.8	82.0
GRM	46.9	51.6	22.0 / 3.8 / 20.4	37.3	86.8
GSP	15.4	4.4	6.9 / 0.6 / 6.7	52.8	90.7
<b>CATCH</b>	<b>56.7</b>	<b>65.4</b>	<b>35.3 / 10.0 / 35.3</b>	<b>58.5</b>	<b>95.9</b>

Table 4: Performance on the Banking Domain.

**The Performance in the In-domain Scenario.** Table 4 demonstrates that CATCH achieves the best performance across all evaluation metrics, showing consistent improvements in both theme segmentation and label generation quality. For **theme distribution**, CATCH yields the highest Clustering Accuracy (Acc = 56.7%, NMI = 65.4%), significantly outperforming the strongest baseline GRM (Acc = 46.9%, NMI = 51.6%). This indicates that the preference-enhanced clustering provides more accurate and consistent theme segmentation. For **theme label quality**, CATCH also leads by a large margin with ROUGE-1 of 35.3% and a Cosine Similarity (CosSim) of 58.5%, surpassing GRM. Furthermore, the 95.9% LLM-based score indicates high consistency with predefined label formats and semantics.

Interestingly, GSP obtains high CosSim (52.8%), but with extremely low Acc (15.4%) and NMI (4.4%), suggesting that it generates plausible-sounding labels but poor underlying topic segmentation. This confirms that directly clustering in SP union space without a grounding space results in misleading theme segmentation.

**The Performance in the Out-of-domain Scenario.** Follow DSTC-12 suggest settings, two out-of-domain datasets

(Finance and Insurance) are used to evaluate the generalization ability of CATCH. As shown in Table 2, CATCH consistently outperforms all baselines across all metrics in both domains, demonstrating strong robustness and cross-domain adaptability. In terms of **theme distribution**, preference-enhanced clustering helps CATCH achieve the highest clustering scores in both Finance (Acc = 55.8, NMI = 67.1) and Insurance (Acc = 54.5, NMI = 62.6), significantly surpassing GRM and other baselines. This also indicates that our clustering mechanism remains effective even when applied to unseen domains. For **theme label quality**, CATCH again leads with strong ROUGE scores (e.g., ROUGE-L = 42.4 in Finance, 41.8 in Insurance) and the highest CosSim and LLM-Score due to the hierarchical theme generation paradigm. Notably, these values surpass those on the in-domain test, underscoring the strong generalization capability of our generation paradigm across domains.

In addition, contrasting all baselines, CATCH performs the best on the Finance domain, whose dialogues have a more ambiguous theme. We attribute this to both the topic-level representation that makes it easier to infer latent thematic structures, and the preference-enhanced clustering that introduces auxiliary input to correctly classify vague points.

## Online Blind Evaluation

We present the official blind test results in Table 5 and in Table 6. Our submission (Team E), based on a lightweight LLaMA-3-8B model without any reliance on proprietary models (e.g., GPT-4o), achieved **second place** overall in both evaluations, and especially, performed stably high in theme label quality. The other two competitive works, Team C (Kim et al. 2025) and Team D (Lee and Lee 2025), both leverage proprietary LLMs more or less.

**Automatic Evaluation.** As shown in Table 5, our system ranked second overall with an average score of 67.5%, closely following Team C. Even with a smaller model, CATCH demonstrated comparable capability to Team C in the final **label generation** task with 41.2 ROUGE-L, 62.5% Cosine Similarity, and consistently high BERTScores (F1 = 93.3%), though for *theme distribution*, the 35.8% Acc and 47.7% NMI are relatively low. This indicates the robustness of CATCH in providing high-quality themes even though clustering quality is low. Notably, CATCH continuously generates structurally clean and stylistically consistent labels as shown by an average of 94.6% in *style scores*.

**Human Evaluation.** Our model also ranked second in human evaluation with an average score of **71.8%**, being consistent with automatic evaluation, highlighting its prac-

Team ID	LLM	Acc	NMI	Rouge-1/2/L	CosSim	BERTScore (P/R/F1)	Sec-1	Sec-2	Avg	Overall
Team C	API	<b>68.0</b>	<b>70.4</b>	<b>45.2 / 23.8 / 45.1</b>	<b>69.9</b>	<b>95.0 / 94.7 / 94.7</b>	<b>100.0</b>	<b>99.5</b>	<b>99.7</b>	<b>75.5</b>
Team E (ours)	<30B	35.8	47.7	42.3 / 16.5 / 41.2	62.5	93.9 / 92.8 / 93.3	93.5	95.7	94.6	67.5
Team D	<30B	51.8	47.7	34.6 / 21.3 / 34.3	55.9	92.5 / 91.5 / 91.9	80.4	76.6	78.5	63.1
Team A	API	48.4	42.0	32.7 / 4.6 / 29.8	59.5	89.8 / 91.2 / 90.4	46.0	56.5	51.2	53.5
Team F	<30B	26.7	9.1	23.1 / 0.8 / 21.1	46.0	85.7 / 89.3 / 87.2	4.1	3.5	3.8	33.4
Team B	API	17.9	2.0	5.0 / 0.0 / 5.0	37.1	85.2 / 88.0 / 86.5	12.0	0.1	6.1	28.8

Table 5: Automatic evaluation results on the blind test set (Travel domain). All values are percentages. LLM: API indicates usage of proprietary models via API; <30B denotes open models smaller than 30B.

tical effectiveness and robustness. We received high marks in both *semantic relevance* (86.3%) and *domain specificity* (91.1%), showing our labels are both meaningful and context-aware. In terms of linguistic quality, we achieved **93.7%** in both grammatical structure and word choice, affirming our system’s fluency and readability.

### Ablation Study

To better understand the contribution of each component in our framework, we conduct comprehensive ablation studies from two perspectives: Module-level ablation and Setting-level ablation.

**Module-Level Ablation.** We evaluate the full model and its three ablated variants on the Finance dataset (see Table 7): **w/o-PeC**: clustering only relies on semantics. **w/o-TopSeg**: theme representation is based solely on raw utterance embeddings. **w/o-HieGen**: removes hierarchical generation and adopts flat label generation instead.

All three modules contribute substantially to overall performance. Discarding **PeC** causes dis-alignment with user preferences, shown by -7.8% decrease in CosSim. Removing **TopSeg** significantly decreases clustering quality, with -14.2% Acc and -13.8% NMI, demonstrating its importance in capturing topical coherence across utterances. Simplifying **HieGen** in flat generation leads to the greatest loss in label generation quality, particularly in ROUGE-L (-2.8%), confirming the effectiveness of hierarchical modeling.

Furthermore, w/o-HieGen greatly declines in theme distribution with -19% Acc and -18.9% NMI, because HieGen is capable of assigning a correct label for the majority of topics within a mingled cluster, where flat label generation usually encounters malfunction, thus provides meaningless labels. Section: Correlation Analysis of PeC and HieGen further probes the relationship between PeC and HieGen.

**Setting-Level Ablation.** To evaluate the stability and design sensitivity of each module, we further examine alternative settings while keeping the overall framework intact, as shown in Table 8.

We first examine the effectiveness of different training depths for the topic segmentation module by either using the original DialSTART without adaptation or only on the conversation-level adaptation dataset. It exhibits a clear trend: deeper domain adaptation leads to better performance.

We also substitute the original settings in the other two modules with stronger ones: replace the clustering algorithm (Spectral Clustering) with ITER-DBSCAN (Chatterjee and

Sengupta 2020)(DBS) and ClusterLLM (Zhang, Wang, and Shang 2023)(ClusLLM), and LLaMA3-8B with GPT-4o for label generation. Both substitution results show the improvements in clustering and theme label quality. Moreover, the combination of both strong substitutions (ClusLLM+GPT-4o) achieves the best performance, demonstrating that our framework consistently benefits from stronger submodules, leading to comprehensive performance gains.

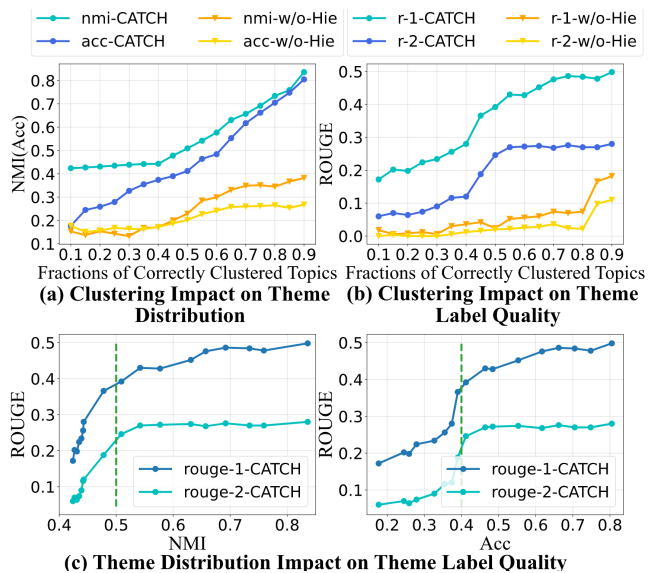


Figure 4: Correlation analysis of PeC and HieGen. Figure 4(a) illustrates the relationship between FCC and cluster-level metrics (NMI and Accuracy).

### Analysis

Though the ablation study suggests that HieGen can partially compensate for the instability introduced by PeC, the mutual effect between components remained unclear. Thus, we further conduct cascade error analysis to expose the error delivery process, and then analyze specifically the profound correlation between PeC and HieGen.

**Cascade Error Analysis.** To decompose the aggregation error to each component, we respectively set the accuracy of TopSeg, HieGen, and their combination to be 100%, creating three ideal models: w/-TopSeg=1, w/-PeC=1, and Both=1. Moreover, the ideal upper limit (IUL) is set to be the combination of w/-GPT4o and w/-TopSeg, PeC=1.

Team ID	Per-Utterance Functional Metrics					Per-Cluster Structural Metrics		Per-Cluster Functional (TD)	Overall Avg.
	SR	AU	GR	ACT	DR	CWC	GS		
Team C	<b>89.7</b>	<b>82.8</b>	<b>47.8</b>	<b>74.8</b>	<b>98.8</b>	<b>100.0</b>	<b>100.0</b>	<b>91.1</b>	<b>85.6</b>
Team E (ours)	86.2	54.6	22.5	54.5	91.1	93.7	93.7	78.3	71.8
Team D	68.8	63.7	26.4	60.3	94.3	91.7	66.7	90.9	70.3
Team A	77.3	63.7	22.8	56.2	79.8	83.3	100.0	75.8	69.8
Team F	45.2	41.6	7.7	41.6	67.5	95.0	100.0	72.6	58.9
Team B	65.0	12.9	0.0	4.1	97.8	100.0	33.3	0.0	39.1

Table 6: Human evaluation results on the blind test set (Travel domain). All values are percentages. Metrics: Semantic Relevance (SR), Analytical Utility (AU), Granularity (GR), Actionability (ACT), Domain Relevance (DR), Conciseness & Word Choice (CWC), Grammatical Structure (GS), and Thematic Distinctiveness (TD).

Model	Acc	NMI	Rouge-1/2/L	Cos	LLM
CATCH	<b>55.8</b>	<b>67.1</b>	<b>42.4 / 24.5 / 42.4</b>	<b>59.3</b>	97.3
w/o-PeC	48.8	59.6	40.7 / 26.7 / 40.7	51.5	<b>98.4</b>
w/o-TopSeg	41.6	53.3	23.6 / 10.1 / 23.6	45.3	87.8
w/o-HieGen	36.8	48.2	19.6 / 9.1 / 19.6	30.3	82.3

Table 7: Module-Level ablation results on the Finance set.

Model	Acc	NMI	Rouge-1/2/L	Cos	LLM
CATCH	55.8	67.1	42.4 / 24.5 / 42.4	59.3	97.3
w/o-Utte. Ada	50.6	59.7	35.6 / 20.0 / 35.6	58.7	95.3
w/o-Utte. + Conv. Ada	44.3	56.6	32.5 / 5.9 / 29.4	53.0	90.2
w/-DBS	60.4	69.2	45.7 / 27.8 / 45.7	61.5	98.1
w/-ClusLLM	70.1	73.4	46.7 / 28.2 / 46.7	62.1	97.5
w/-Qwen-3-8B	56.0	67.3	43.6 / 25.4 / 43.6	59.3	97.3
w/-GPT4o	57.0	68.1	45.7 / 26.2 / 44.9	61.1	98.3
w/-DBS+GPT4o	61.8	70.1	46.8 / 28.6 / 45.9	63.7	98.6
w/-ClusLLM+GPT4o	<b>70.4</b>	<b>74.1</b>	<b>49.9 / 30.3 / 48.8</b>	<b>71.0</b>	<b>98.6</b>

Table 8: Setting-level ablation results on the Finance set.

Model	1-pk	Acc	NMI	Rouge-1/2/L	Cos	LLM
CATCH	0.37	55.8	67.1	42.4 / 24.5 / 42.4	59.3	97.3
w/-TopSeg=1	1	56.9	66.3	46.9 / 26.3 / 45.3	66.1	98.4
w/-PeC=1	0.37	100	100	51.6 / 32.4 / 51.6	73.7	98.4
Both=1	1	100	100	56.0 / 34.5 / 55.5	77.9	98.7
IUL	1	100	100	<b>59.0 / 37.4 / 58.1</b>	<b>80.2</b>	<b>99.5</b>

Table 9: Cascade error analysis on the Finance set. 1-pk indicates the topic boundary accuracy.

The results in Table 9 align with the ablation findings: reducing each partial error steadily improves final performance. Notably, the framework is more sensitive to cluster quality, yielding a larger gain (+9.2% w/-PeC=1) compared with topic segmentation (+4.5% w/-TopSeg=1) on Rouge-1.

**Correlation Analysis of PeC and HieGen.** To further investigate the interplay between PeC and HieGen, we analyze how clustering quality affects the final theme labeling performance. Specifically, we evaluate both CATCH and its ablated variant w/o-HieGen under varying *fractions of correctly clustered topics* (FCC), by randomly reassigning a portion of topics to incorrect clusters.

As shown in Figure 4(a) and (b), CATCH consistently outperforms w/o-HieGen in both theme distribution accuracy and label quality. HieGen exhibits strong intra-cluster denoising capabilities: the cluster-level quality remains closely proportional to FCC, while in contrast, the performance of

w/o-HieGen plateaus at a low level due to its limited capacity in generating coherent theme labels.

Furthermore, Figure 4(b) shows that HieGen substantially reduces the system’s dependency on clustering quality. For CATCH, theme label quality improves sharply even with a large amount of clustering errors, where FCC exceeds only 0.4, while the improvement for w/o-HieGen delays to FCC surpasses 0.8. This trend is further extended in Figure 4(c), which demonstrates the relationship between theme distribution consistency and label quality. Notably, HieGen maintains strong performance even under suboptimal clustering: ROUGE-1 >0.4 and ROUGE-2 >0.26, despite limited NMI (0.5) and Accuracy (0.4). These results show that HieGen’s divide-and-conquer denoising can fix clustering errors as long as about half of the topics are correctly grouped.

## Conclusion

In this paper, we propose CATCH, a novel theme detection framework that significantly enhances the automatic discovery and consistency of themes within a latent topic space aligned with user preferences. By treating the entire architecture as a theme generation pipeline, CATCH jointly models intra-dialogue theme representation and inter-dialogue preference-aware alignment, leading to coherent and user-aligned theme labels. Extensive experiments demonstrate the robustness and generalizability of CATCH across diverse tasks, while ablation studies further reveal the complementary roles and coordination of its three key modules. In future work, we plan to continuously improve the framework with cutting-edge techniques and make it more adaptive and dynamic, enabling its application to a broader range of downstream scenarios, such as proactive dialogue systems, dialogue control, and fine-grained dialogue analysis.

## Acknowledgments

This work was supported by the project of National Natural Science Foundation of China (Grant No. 62271432), Shenzhen Science and Technology Program (Grant No. ZDSYS20230626091302006), and Program for Guangdong Introducing Innovative and Entrepreneurial Teams (Grant No. 2023ZT10X044), Key Project of Shenzhen Higher Education Stability Support Program (Grant No. 2024SC0009), and Shenzhen Science and Technology Program (Grant No. RCBS20231211090538066).

## References

- Bach, T. X.; Anh, N. D.; Van Linh, N.; and Than, K. 2021. Dynamic transformation of prior knowledge into Bayesian models for data streams. *IEEE TKDE*, 35(4): 3742–3750.
- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *JMLR*, 3(Jan): 993–1022.
- Chatterjee, A.; and Sengupta, S. 2020. Intent Mining from past conversations for Conversational Agent. In *COLING*, 4140–4152.
- Costa, R.; Martins, B.; Viana, S.; and Coheur, L. 2023. Towards a Fully Unsupervised Framework for Intent Induction in Customer Support Dialogues. *CoRR*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 4171–4186.
- Du, L.; Buntine, W.; and Johnson, M. 2013a. Topic segmentation with a structured topic model. In *NAACL-HLT*, 190–200.
- Du, L.; Buntine, W.; and Johnson, M. 2013b. Topic Segmentation with a Structured Topic Model. In *NAACL-HLT*, 190–200.
- Eisenstein, J.; and Barzilay, R. 2008. Bayesian Unsupervised Topic Segmentation. In *EMNLP*, 334–343.
- Fan, Y.; Jiang, F.; Li, P.; and Li, H. 2024. Uncovering the Potential of ChatGPT for Discourse Analysis in Dialogue: An Empirical Study. In *REC-COLING 2024*, 16998–17010.
- Gao, H.; Wang, R.; Lin, T.-E.; Wu, Y.; Yang, M.; Huang, F.; and Li, Y. 2023. Unsupervised dialogue topic segmentation with topic-aware contrastive learning. In *ACM SIGIR*, 2481–2485.
- Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.
- Gung, J.; Shu, R.; Moeng, E.; Rose, W.; Romeo, S.; Gupta, A.; Benajiba, Y.; Mansour, S.; and Zhang, Y. 2023. Intent Induction from Conversations for Task-Oriented Dialogue Track at DSTC 11. In *DSTC 11*, 242–259.
- Hearst, M. A. 1997. Text Tiling: Segmenting Text into Multi-paragraph Subtopic Passages. *Computational Linguistics*, 23(1): 33–64.
- Jiang, F.; Liu, W.; Chu, X.; Li, P.; Zhu, Q.; and Li, H. 2024. Advancing Topic Segmentation and Outline Generation in Chinese Texts: The Paragraph-level Topic Representation, Corpus, and Benchmark. In *REC-COLING 2024*, 495–506.
- Kim, S.; Jeong, T.; Hong, S.; Kim, S.; Lee, J.; Chang, D.-S.; and Koo, M.-W. 2025. KSTC: Keyphrase-driven Sentence embedding and Task independent prompting for filling slot in the Generation of theme label. In *Proceedings of the Twelfth Dialog System Technology Challenge*, 44–73.
- Lee, J.; and Lee, G. 2025. The Limits of Post-hoc Preference Adaptation: A Case Study on DSTC12 Clustering. In *Proceedings of the Twelfth Dialog System Technology Challenge*, 36–43.
- Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, 74–81.
- Lin, Y.; Ma, X.; Gao, X.; Li, R.; Wang, Y.; and Chu, X. 2024. Combating Label Sparsity in Short Text Topic Modeling via Nearest Neighbor Augmentation. In *Findings of ACL*, 13762–13774.
- Liu, P.; Ning, Y.; Wu, K. K.; Li, K.; and Meng, H. 2021. Open intent discovery through unsupervised semantic clustering and dependency parsing. *arXiv preprint arXiv:2104.12114*.
- Mendonça, J.; Zhang, L.; Mallidi, R.; Lavie, A.; Trancoso, I.; D’Haro, L. F.; and Sedoc, J. 2025. Overview of dialog system evaluation track: Dimensionality, language, culture and safety at dstc 12. In *Proceedings of the Twelfth Dialog System Technology Challenge*, 27–35.
- Nguyen, H.; Pham, H.; Nguyen, S.; Van Linh, N.; and Than, K. 2022. Adaptive infinite dropout for noisy and sparse data streams. *Machine Learning*, 111(8): 3025–3060.
- Nguyen, Q. D.; Nguyen, T.; Nguyen, D. A.; Van, L. N.; Dinh, S.; and Nguyen, T. H. 2025. GloCOM: A Short Text Neural Topic Model via Global Clustering Context. In *NAACL 2025*, 1109–1124.
- Perkins, H.; and Yang, Y. 2019. Dialog Intent Induction with Deep Multi-View Clustering. In *EMNLP-IJCNLP*, 4016–4025.
- Pham, D.-T.; Vu, T. T. N.; Nguyen, T.; Ngo, L.; Nguyen, D.; and Nguyen, T. 2024. NeuroMax: Enhancing Neural Topic Modeling via Maximizing Mutual Information and Group Topic Regularization. In *Findings of EMNLP*, 7758–7772.
- Pu, J.; Chen, G.; Chang, Y.; and Mao, X. 2022. Dialog intent induction via density-based deep clustering ensemble. *arXiv preprint arXiv:2201.06731*.
- Riedl, M.; and Biemann, C. 2012. TopicTiling: A Text Segmentation Algorithm based on LDA. In *Proceedings of ACL 2012 Student Research Workshop*, 37–42.
- Sia, S.; Dalmia, A.; and Mielke, S. J. 2020. Tired of Topic Models? Clusters of Pretrained Word Embeddings Make for Fast and Good Topics too! In *EMNLP*, 1728–1736.
- Tuan, A. P.; Tran, B.; Nguyen, T. H.; Van, L. N.; and Than, K. 2020. Bag of biterms modeling for short texts. *Knowledge and Information Systems*, 62(10): 4055–4090.
- Vinh, N. X.; Epps, J.; and Bailey, J. 2010. Information Theoretic Measures for Clusterings Comparison: Variants, Properties, Normalization and Correction for Chance. *Journal of Machine Learning Research*, 11(95): 2837–2854.
- Vu, T.; Do, M.; Nguyen, T.; Van, L. N.; Dinh, S.; and Nguyen, T. H. 2025. Topic Modeling for Short Texts via Optimal Transport-Based Clustering. In *Findings of the Association for Computational Linguistics: ACL 2025*, 7666–7680.
- Vuong, H. T.; Le, T.; Vu, T.; Nguyen, T.; Van, L. N.; Dinh, S.; and Nguyen, T. H. 2025. HiCOT: Improving Neural Topic Models via Optimal Transport and Contrastive Learning. In *Findings of ACL 2025*, 13894–13920.
- Xing, L.; and Carenini, G. 2021. Improving Unsupervised Dialogue Topic Segmentation with Utterance-Pair Coherence Scoring. In *SIGDIAL*, 167–177.

Xu, J.; Jiang, F.; Gao, A.; D’Haro, L. F.; and Li, H. 2025. Unsupervised Mutual Learning of Discourse Parsing and Topic Segmentation in Dialogue. *IEEE Transactions on Audio, Speech and Language Processing*.

Zeng, Z.; Ma, D.; Yang, H.; Gou, Z.; and Shen, J. 2021. Automatic intent-slot induction for dialogue systems. In *WWW*, 2578–2589.

Zhang, Y.; Wang, Z.; and Shang, J. 2023. ClusterLLM: Large Language Models as a Guide for Text Clustering. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 13903–13920.

Zhang, Z.; Fang, M.; Chen, L.; and Namazi-Rad, M.-R. 2022. Is Neural Topic Modelling Better than Clustering? An Empirical Study on Clustering with Contextual Embeddings for Topics. In *NAACL-HLT*, 3886–3893.