

# Discovery-Oriented Faceting: From Coverage to Blind-Spot Discovery

Youdi Li  
ri.yutei@jp.panasonic.com  
Panasonic Connect Co., Ltd.  
Tokyo, Japan

## Abstract

When people explore large document collections to build understanding, they face a challenge: existing AI tools help them see what is central but tend to hide what is unusual. Summarization and topic modeling optimize for coverage, representing main themes while pushing minority viewpoints and edge cases out of view. This matters because discovery often depends on noticing what does not fit, such as unexpected findings, minority positions, or gaps in the literature. When tools hide this content, users may miss insights that could change their understanding. In this paper, we explore an alternative objective: blind-spot discovery, where the goal is to surface content that coverage methods suppress so that people can judge its significance for themselves. We propose three design goals and illustrate them through DOF (Discovery-Oriented Faceting), a system that organizes documents into categories with explicit boundaries, ranks categories by distinctiveness rather than size, and supports iterative refinement. Comparing DOF against coverage-based ranking across four domains, we find that the two approaches surface fundamentally different content, with DOF promoting specialized categories that coverage methods bury. We discuss how shifting from coverage to discovery may offer a complementary mode of support for people exploring large text collections.

## Keywords

Faceted Summarization, Blind-Spot Discovery, Information Gain, Cognitive Scaffolding, Human-AI Collaboration

## 1 Introduction

Making sense of large document collections is central to knowledge work. Patent examiners search archives to spot unexpected approaches. Literature researchers survey prior work to find gaps. Policy analysts review reports to identify emerging issues.

AI tools for document exploration, including faceted summarization [20], topic modeling [1, 10], and interactive navigation [14], help structure large text collections. However, current tools share a design choice that works against discovery: they optimize for coverage. FacetSum [20] and WikiAsp [12] use predefined categories that represent central content well but miss minority viewpoints. Topic models produce clusters that overlap and drift, with boundaries

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

hard to interpret [24]. LLM-guided approaches [27] improve cluster coherence but do not expose distinctive content. Once generated, most structures offer no way to refine [14].

The problem is that coverage optimization pushes unusual content out of view. The content most valuable for new understanding (specialized vocabulary, minority positions, emerging themes) is systematically deprioritized. When tools show only the central tendency, they risk reinforcing what people already know rather than exposing what they do not. Recent work addresses parts of this problem [3–5, 16], but each only partly: surfacing novelty without ranking, or accepting feedback without explicit boundaries. Ebrahimi and Peltonen [7] used constrained matrix factorization for minority topic discovery, though their method requires predefined seed words.

In this paper, we present Discovery-Oriented Faceting (DOF), a system that takes an alternative approach: blind-spot discovery, where the goal is to reveal content that standard methods suppress. DOF is built around three design goals:

**D1. Make boundaries explicit.** Each category specifies what belongs and what does not. A topic model might label a cluster “Drug Research,” leaving users unsure whether it includes clinical trials or pricing policy. DOF specifies inclusions and exclusions so users can quickly judge relevance.

**D2. Prioritize what is unusual.** DOF ranks categories by how much they differ from the corpus average, not by size, highlighting content that is both coherent and distinctive.

**D3. Support iterative refinement.** Understanding develops through exploration [22], and DOF is designed to intervene at the transition from information foraging to schema construction (see Section 5). DOF lets users merge, split, or hide categories without starting over.

We evaluate DOF across four domains, showing that distinctiveness ranking reveals specialized content that coverage methods miss. We discuss implications for tools that support discovery in knowledge work.

## 2 Related Work

### 2.1 Faceted Summarization

Latent Dirichlet Allocation [1] established probabilistic topic modeling, with BERTopic [10] extending this via neural embeddings. FacetSum [20] introduced structured summaries with predefined facets, and WikiAsp [12] extended this to 20 domains. iFacetSum [14] enabled interactive navigation via induced facets. Boytsov et al. [2] extracted aspect-sentiment pairs from reviews, though specialized to sentiment-bearing domains. These systems relied on static schemas or domain-specific pipelines, leaving open how to discover facets dynamically.

## 2.2 Novelty-Aware Summarization

DualSum [5] separated novel from redundant content, and Nouveau-ROUGE [4] proposed metrics rewarding absent content. CoCo-Sum [16] generated contrastive summaries between entity sets. STRUM [11] and Hayashi et al. [13] generated “contribution” vs. “context” summaries without predefined ontologies. Ebrahimi and Peltonen [7] incorporated seed word constraints into NMF to capture low-prevalence topics, though users must specify the seed words in advance. These approaches treated novelty as binary rather than as a ranking criterion for exploration.

## 2.3 Human-in-the-Loop Refinement

Interactive clustering with user constraints [3] improved coherence, while Smith et al. [24] found users over-trusted topic models. ClusterLLM [27] used LLM feedback to improve cluster quality, focusing on coherence rather than distinctiveness. Fang et al. [8] developed interactive topic modeling with change tracking, and Zade et al. [26] showed how visualizing disagreement aided interpretation. CollabCoder [9] and Diaz-Rodriguez [6] integrated model suggestions with collaborative coding and proposed “summaries as centroids.” However, existing systems lacked refinement mechanisms or required full regeneration after each edit.

## 3 The DOF System

Figure 1 illustrates the DOF pipeline and contrasts it with coverage-oriented approaches. The key distinction occurs *after* clustering: coverage tools rank clusters by size (surfacing typical content), while DOF ranks by distinctiveness (surfacing content that deviates from corpus norms). The pipeline has five stages, with all parameters locked before processing and held constant across all datasets.

### 3.1 Snippet Construction

Long documents are segmented into 425-word snippets with 75-word overlap. Each snippet is embedded using OpenAI’s embedding model (text-embedding-3-small, 1536 dimensions) [21]. This fixed granularity avoids the over-generalization that direct LLM application to long documents produces [19].

### 3.2 Scope-Bounded Clustering

Snippets are partitioned into  $K=15$  clusters via K-Means (random state=42). For each cluster, an LLM (GPT-4o-2024-08-06, temperature=0.7, max tokens=500) examines the snippets and generates a scope definition: a label, four positive statements (what belongs), and four negative statements (what does not belong). The LLM infers exclusions by identifying semantically related content that falls outside the cluster’s actual scope. Unlike topic models, which leave boundaries implicit, DOF explicitly states what does not belong [25]. This helps users decide whether a category fits their needs.

### 3.3 Distinctiveness Ranking

Standard tools rank by size or centrality, burying specialized content. DOF scores each cluster using Kullback-Leibler (KL) divergence between its embedding distribution and the corpus average, then sorts from most distinctive to least. A cluster about prosthetic

devices in a chemistry patent corpus scores high because its embedding distribution diverges from the corpus average.

## 3.4 Evidence Sampling

For each top-ranked category, DOF samples three snippets: a **central** snippet (closest to centroid), a **peripheral** snippet (farthest valid member), and a **transitional** snippet (median distance). Showing the range from core to edge lets users verify coherence before exploring further.

## 3.5 Iterative Refinement

Beyond ranking, DOF lets users reshape the structure incrementally: **merge** categories that overlap, **split** categories that conflate distinct ideas, or **hide** irrelevant categories. Each action updates the view in place [22], so refinements build on each other.

## 4 Evaluation

We evaluate DOF on *blind-spot discovery*: whether the system exposes content that deviates from corpus-wide patterns, making it available for human judgment rather than hiding it behind coverage rankings.

### 4.1 Datasets

We selected four datasets spanning distinct domains (Table 1): **GovReport** [15] (U.S. government agency reports), **BillSum** [17] (congressional legislation), **BigPatent** [23] (patent applications), and **BookSum** [18] (literary works). All were processed with identical locked parameters ( $K=15$ , 425-word snippets, 75-word overlap). This diversity allows us to observe how DOF generalizes without domain-specific tuning.

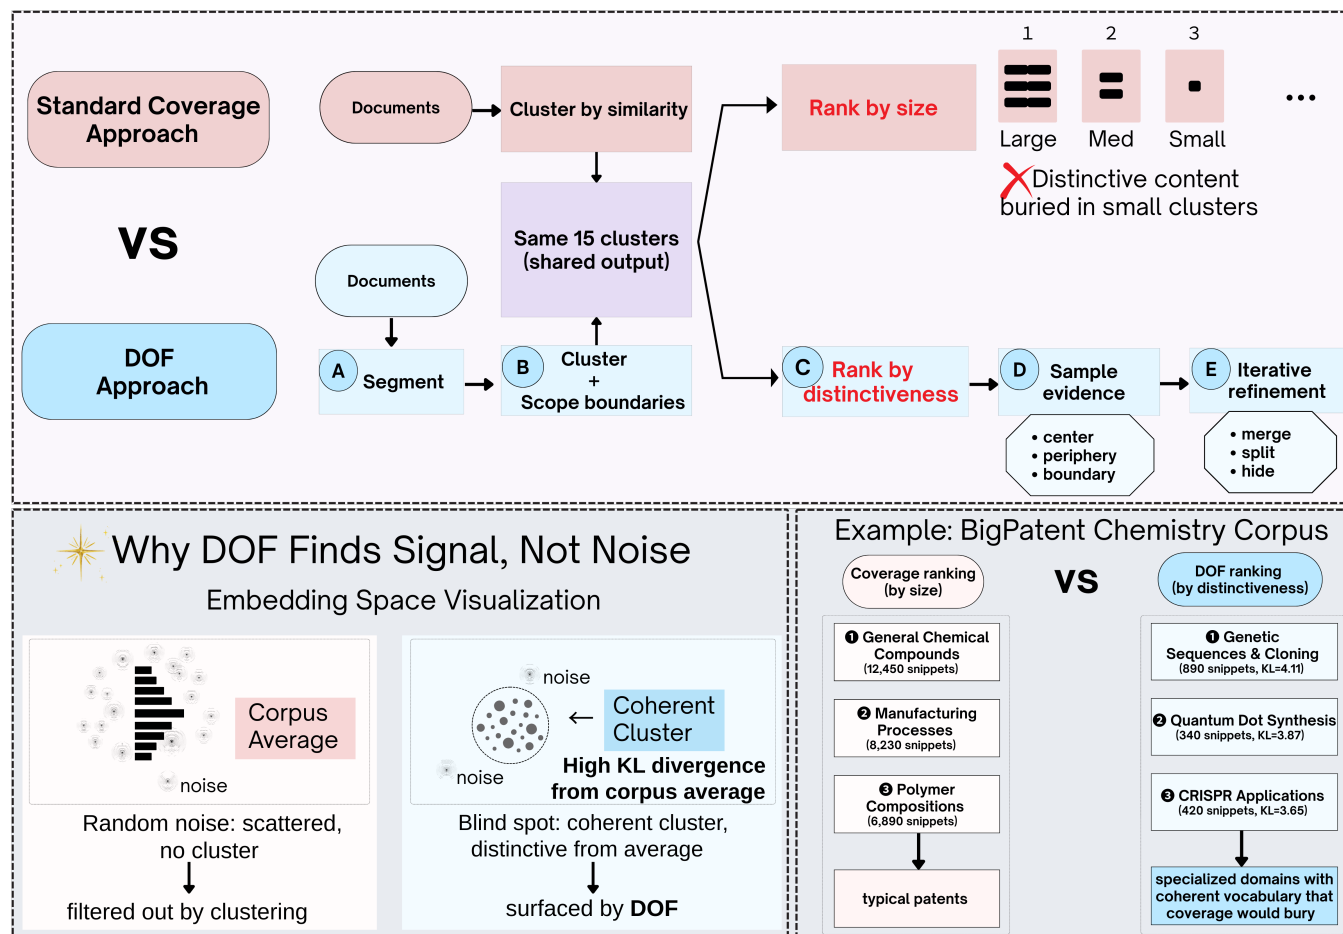
Table 1: Dataset overview.

Dataset	Domain	Docs	Snippets
GovReport	Government	973	21,001
BillSum	Legislation	3,269	12,872
BigPatent	Patents	10,764	150,000
BookSum	Literature	1,431	17,321
<b>Total</b>		16,437	201,194

### 4.2 Ranking Strategy Comparison

To show that DOF brings different content into view than coverage-based approaches, we compare two ranking strategies on the same clustering results: (1) **Coverage**: rank by cluster size; (2) **DOF**: rank by KL divergence. Top-5 overlap counts how many facets appear in both rankings’ top five, where lower overlap means users see different content. Spearman  $\rho$  measures whether rankings go in the same direction, and  $\rho = -1$  means the largest clusters are the least distinctive, so coverage and DOF show opposite content.

In technical domains (BigPatent, BookSum), the correlation is perfectly negative ( $\rho = -1.0$ ): the largest clusters have the lowest distinctiveness scores. Coverage ranking foregrounds typical content (Innovative Personal Care Products, Advanced Personal Care



**Figure 1: DOF contrasted with coverage approaches.** Upper: Both approaches cluster documents, but the key difference is ranking. Coverage ranks by size (large clusters first), burying distinctive content. DOF ranks by KL divergence (distinctive first), then adds scope boundaries, evidence sampling, and user refinement. Lower left: Signal vs. noise in embedding space. Random noise scatters without forming clusters, while coherent blind spots form distinct clusters with high KL divergence, which DOF prioritizes. Lower right: Illustrative BigPatent example showing the same clusters ranked differently. Coverage foregrounds large, typical clusters while DOF foregrounds small, specialized ones.

**Table 2: Coverage vs. DOF ranking comparison.** Lower overlap and negative  $\rho$  indicate different content surfaced.

Dataset	Top-5 Overlap	Spearman $\rho$
GovReport	2/5	-0.06
BillSum	2/5	+0.31
BigPatent	0/5	-1.00
BookSum	0/5	-1.00

Devices), while DOF highlights specialized domains (Advanced Personal Hygiene Products, Advanced Prosthetic Devices). In BookSum, DOF reveals facets like “Strategies of Political Leadership” (*The Prince*) and “Colonial Confrontations” (*Lord Jim*), which coverage methods rank lower in favor of more common literary themes.

Government and legislative text shows weaker divergence because formulaic language means even distinctive facets tend to be moderately sized.

### 4.3 Cross-Domain Results

**Table 3: Distinctiveness results across domains.**

Dataset	$KL_A$	Max KL	Top Facet
GovReport	0.98	1.74	Trade
BillSum	1.08	1.97	Tax
BigPatent	2.78	3.59	Hygiene
BookSum	2.78	3.28	Politics

Facet distinctiveness varies systematically across domains (Table 3). Government and legislative text shows lower distinctiveness

( $KL_A \approx 1.0$ ) due to formulaic language, while technical and literary text shows higher distinctiveness ( $KL_A \approx 2.8$ ) due to specialized vocabulary. In BigPatent, “Hygiene Products” (Max  $KL=3.59$ ) has an embedding distribution far from the corpus average, while in GovReport, “Trade Policy” (Max  $KL=1.74$ ) remains closer, reflecting the shared language of government text.

#### 4.4 Scope Boundaries Aid Interpretation

Scope boundaries help users understand what belongs in each facet. For example, a BigPatent facet labeled “Advanced Prosthetic Devices” explicitly states what it includes (bionic limb technologies, neural interface control systems, adaptive mobility prosthetics, sensory feedback mechanisms) and what it excludes (textile manufacturing, food preservation, water purification, kitchen utensil design). Without these exclusions, users might confuse this facet with broader biomedical or manufacturing content. In GovReport, a facet labeled “Medicaid and Medicare Administration” excludes private insurance, international health systems, and pharmaceutical research, helping analysts quickly confirm it matches their task.

#### 4.5 Preference Refinement

We simulated 10 rounds of refinement operations across all four datasets to evaluate whether the merge, split, and hide operations help users reach a stable organization (Table 4). In each round, we applied rule-based operations: merging facets with centroid similarity above 0.8, and hiding facets with  $KL$  divergence below the dataset median. We track three metrics: **inter-facet similarity** (cosine similarity between centroids) measures overlap,  $KL_A$  measures distinctiveness, and **edits per round** tracks convergence. On average, similarity decreased by 28%,  $KL_A$  increased by 15%, and edits dropped by 71%, suggesting refinement reduces redundancy and focuses on distinctive content.

**Table 4: Simulated preference refinement (10 rounds per dataset).**

Dataset	Similarity $\Delta$	$KL_A$ $\Delta$	Edits $\Delta$
GovReport	-23%	+12%	-68%
BillSum	-27%	+14%	-71%
BigPatent	-29%	+16%	-72%
BookSum	-31%	+18%	-74%
<b>Average</b>	<b>-28%</b>	<b>+15%</b>	<b>-71%</b>

#### 4.6 Parameter Robustness and Scope Quality

Table 5 reports three automated checks on GovReport. The inverse relationship between cluster size and distinctiveness holds across  $K=5$  to 30, confirming that DOF’s ranking is not an artifact of a single  $K$  setting. We define a snippet as boundary-ambiguous if its cosine distance ratio to the nearest versus second-nearest centroid exceeds 0.95 (i.e., it nearly belongs to two clusters). At this threshold, 37.4% of snippets are ambiguous overall, but distinctive clusters have fewer such cases ( $\rho = -0.90$ ,  $p < 0.0001$ ), suggesting DOF prioritizes content that is more cleanly separable. An LLM coherence check (GPT-4o-mini) rated all 15 scope definitions at 4.7/5

on average, with label clarity and inclusion coherence at ceiling (5.0/5) and keyphrase alignment lowest (4.0/5).

**Table 5: Automated validation (GovReport,  $K=15$ ).**

Analysis	Metric	Result
K sensitivity ( $K=5-30$ )	$\rho$ (size vs. $KL$ )	-.73 to -.90
	Top-5 overlap, $K \geq 10$	$\leq 1/5$
Boundary membership	Ambiguous snippets	37.4%
	$\rho$ ( $KL$ vs. boundary%)	-.90 <sup>***</sup>
Scope quality (GPT-4o-mini)	Label clarity	5.0/5
	Inclusion coherence	5.0/5
	Exclusion usefulness	4.9/5
	Keyphrase alignment	4.0/5

<sup>\*\*\*</sup>  $p < 0.0001$

## 5 Discussion and Future Work

DOF reframes document exploration around what is unusual rather than what is common. Here we situate DOF within sensemaking theory, discuss its intended users, and outline current limitations.

*Where DOF Intervenes in Sensemaking.* Pirolli and Card’s [22] sensemaking model describes a loop from information foraging (searching, filtering, extracting) to synthesis (building schemas, drawing conclusions). DOF intervenes at the foraging-to-synthesis transition. During foraging, coverage-based tools narrow attention toward frequently occurring themes, potentially causing users to construct schemas that reflect only central content. DOF expands the foraging stage by exposing unfamiliar yet coherent clusters before the user commits to a schema. DOF makes unfamiliar clusters visible and interpretable, but whether users act on this exposure, and how it shapes their final schemas, requires user studies to establish.

*Target Users and Use Contexts.* DOF’s value likely varies across use contexts. We see the strongest fit for *exploratory* professionals who face open-ended questions over large corpora: a patent examiner searching for prior art in an unfamiliar technology area benefits from seeing specialized clusters (e.g., prosthetic innovations in a chemistry corpus) that a size-ranked view would bury. A literature researcher surveying a new field may use DOF to identify minority positions that challenge mainstream narratives. Users with well-defined search goals may find distinctiveness ranking less useful or even disorienting, since DOF deliberately foregrounds content outside the familiar center. In such cases, coverage-based tools remain more appropriate. DOF is designed to complement, not replace, them.

*Limitations and Future Work.* The preference refinement used rule-based simulations rather than real user studies. The fixed parameters ( $K=15$ , 425-word snippets) are robust across  $K=5$  to 30 (Table 5), but K-Means with cosine similarity can produce clusters that are mathematically coherent but semantically coarse; LLM scope generation partially compensates but does not guarantee human-meaningful boundaries. User studies and soft clustering are natural next steps.

## References

- [1] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022.
- [2] Ilya Boytsov, Vinny DeGenova, Mikhail Balyasin, Joseph Walt, Caitlin Eusden, Marie-Claire Rochat, and Margaret Pierson. 2025. End-to-End Aspect-Guided Review Summarization at Scale. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track*. Association for Computational Linguistics. doi:10.18653/v1/2025.emnlp-industry.31
- [3] David Cohn, Rich Caruana, and Andrew McCallum. 2003. *Semi-supervised Clustering with User Feedback*. Technical Report TR2003-1892. Cornell University. <https://ecommons.cornell.edu/items/bcf991e7-2079-48c1-9e29-1b69a60bd1ea>
- [4] John M. Conroy, Judith D. Schlesinger, and Dianne P. O’Leary. 2011. Nouveau-ROUGE: A Novelty Metric for Update Summarization. *Computational Linguistics* 37, 1 (2011), 1–8. doi:10.1162/coli\_a\_00033
- [5] Jean-Yves Delort and Enrique Alfonseca. 2012. DualSum: A Topic-Model Based Approach for Update Summarization. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Avignon, France, 214–223. <https://aclanthology.org/E12-1022/>
- [6] Jairo Diaz-Rodriguez. 2025. Summaries as Centroids for Interpretable and Scalable Text Clustering. arXiv:2502.09667
- [7] Seydeh Fatemeh Ebrahimi and Jaakko Peltonen. 2025. Constrained Non-negative Matrix Factorization for Guided Topic Modeling of Minority Topics. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Suzhou, China, 35573–35598. doi:10.18653/v1/2025.emnlp-main.1802
- [8] Zheng Fang, Lama Alqazlan, Du Liu, Yulan He, and Rob Procter. 2023. A User-Centered, Interactive, Human-in-the-Loop Topic Modelling System. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Dubrovnik, Croatia, 505–522. doi:10.18653/v1/2023.eacl-main.37
- [9] Jie Gao, Yuchen Guo, Gionnieve Lim, Tianqin Zhang, Zheng Zhang, Toby Jia-Jun Li, and Simon Tangi Perrault. 2024. CollabCoder: A Lower-barrier, Rigorous Workflow for Inductive Collaborative Qualitative Analysis with Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM. doi:10.1145/3613904.3642002
- [10] Maarten Grootendorst. 2022. BERTopic: Neural Topic Modeling with a Class-based TF-IDF Procedure. arXiv:2203.05794
- [11] Beliz Gunel, Sandeep Tata, and Marc Najork. 2023. STRUM: Extractive Aspect-Based Contrastive Summarization. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*. ACM, Austin, TX, USA, 28–31. doi:10.1145/3543873.3587304
- [12] Hiroaki Hayashi, Prashant Budania, Peng Wang, Chris Ackerson, Raj Neervannan, and Graham Neubig. 2021. WikiAsp: A Dataset for Multi-domain Aspect-based Summarization. *Transactions of the Association for Computational Linguistics* 9 (2021), 211–225. doi:10.1162/tacl\_a\_00362
- [13] Hiroaki Hayashi, Wojciech Kryściński, Bryan McCann, Nazneen Rajani, and Caiming Xiong. 2023. What’s New? Summarizing Contributions in Scientific Literature. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Dubrovnik, Croatia, 978–991. doi:10.18653/v1/2023.eacl-main.72
- [14] Eran Hirsch, Alon Eirew, Ori Shapira, Avi Caciularu, Arie Cattan, Ori Ernst, Ramakanth Pasunuru, Hadar Ronen, Mohit Bansal, and Ido Dagan. 2021. iFacetSum: Coreference-based Interactive Faceted Summarization for Multi-Document Exploration. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 283–297. doi:10.18653/v1/2021.emnlp-demo.33
- [15] Luyang Huang, Shuyang Cao, Nikolaus Parulian, Heng Ji, and Lu Wang. 2021. Efficient Attention for Long Document Summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, 1419–1436. doi:10.18653/v1/2021.naacl-main.112
- [16] Hayate Iso, Xiaolan Wang, Stefanos Angelidis, and Yoshihiko Suhara. 2022. Comparative Opinion Summarization via Collaborative Decoding. In *Findings of the Association for Computational Linguistics: ACL 2022*. Association for Computational Linguistics, Dublin, Ireland, 3307–3324. doi:10.18653/v1/2022.findings-acl.261
- [17] Anastassia Kornilova and Vladimir Eidelman. 2019. BillSum: A Corpus for Automatic Summarization of US Legislation. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*. Association for Computational Linguistics, Hong Kong, China, 48–56. doi:10.18653/v1/D19-5406
- [18] Wojciech Kryściński, Nazneen Rajani, Divyansh Agrawal, Caiming Xiong, and Dragomir Radev. 2021. BookSum: A Collection of Datasets for Long-form Narrative Summarization. arXiv:2105.08209
- [19] Michael X. Liu, Tongshuang Wu, Tianying Chen, Franklin Mingzhi Li, Aniket Kittur, and Brad A. Myers. 2024. Selenite: Scaffolding Online Sensemaking with Comprehensive Overviews Elicited from Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM. doi:10.1145/3613904.3642149
- [20] Rui Meng, Khushboo Thaker, Lei Zhang, Yue Dong, Xingdi Yuan, Tong Wang, and Daqing He. 2021. Bringing Structure into Summaries: A Faceted Summarization Dataset for Long Scientific Documents. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 1080–1089. doi:10.18653/v1/2021.acl-short.137
- [21] OpenAI. 2024. Text Embedding 3 Small. <https://platform.openai.com/docs/models/text-embedding-3-small>
- [22] Peter Pirolli and Stuart Card. 2005. The Sensemaking Process and Leverage Points for Analyst Technology as Identified through Cognitive Task Analysis. In *Proceedings of International Conference on Intelligence Analysis*, Vol. 5. McLean, VA, 2–4.
- [23] Eva Sharma, Chen Li, and Lu Wang. 2019. BIGPATENT: A Large-Scale Dataset for Abstractive and Coherent Summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 2204–2213. doi:10.18653/v1/P19-1212
- [24] Alison Smith, Varun Kumar, Jordan Boyd-Graber, Kevin Seppi, and Leah Findlater. 2018. Closing the Loop: User-Centered Design and Evaluation of a Human-in-the-Loop Topic Modeling System. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces (IUI '18)*. ACM, 293–304. doi:10.1145/3172944.3172965
- [25] Sitong Wang, Samia Menon, Dingzeyu Li, Xiaojuan Ma, Richard Zemel, and Lydia B. Chilton. 2025. Schemex: Interactive Structural Abstraction from Examples with Contrastive Refinement. arXiv:2504.11795
- [26] Himanshu Zade, Margaret Drouhard, Bonnie Chinh, Lu Gan, and Cecilia Aragon. 2018. Conceptualizing Disagreement in Qualitative Coding. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, Article 159, 11 pages. doi:10.1145/3173574.3173733
- [27] Yuwei Zhang, Zihan Wang, and Jingbo Shang. 2023. ClusterLLM: Large Language Models as a Guide for Text Clustering. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Singapore, 13903–13920. doi:10.18653/v1/2023.emnlp-main.858