

When Numbers Tell Half the Story: Human-Metric Alignment in Topic Model Evaluation

Thibault Prouteau¹, Francis Lareau^{3,4}, Nicolas Dugué⁵,
Jean-Charles Lamirel^{1,2}, Christophe Malaterre⁴

¹Université de Lorraine, CNRS, LORIA, Nancy (France)

²Université de Strasbourg, Strasbourg (France)

³Université de Sherbrooke, Dept of Philosophy and Applied Ethics, Sherbrooke (Canada)

⁴Université du Québec à Montréal, Dept of Philosophy & CIRST, Montréal (Canada)

⁵Le Mans Université, LIUM, Le Mans (France)

{thibault.prouteau, jean-charles.lamirel}@loria.fr, {francis.lareau, christophe.malaterre}@uqam.ca,
nicolas.dugue@univ-lemans.fr

Abstract

Topic models uncover latent thematic structures in text corpora, yet evaluating their quality remains challenging, particularly in specialized domains. Existing methods often rely on automated metrics like *topic coherence* and *diversity*, which may not fully align with human judgment. Human evaluation tasks, such as *word intrusion*, provide valuable insights but are costly and primarily validated on general-domain corpora. This paper introduces *Topic Word Mixing* (TWM), a novel human evaluation task assessing inter-topic distinctness by testing whether annotators can distinguish between word sets from single or mixed topics. TWM complements *word intrusion*'s focus on intra-topic coherence and provides a human-grounded counterpart to diversity metrics. We evaluate six topic models—both statistical and embedding-based (LDA, NMF, Top2Vec, BERTopic, CFMF, CFMF-emb)—comparing automated metrics with human evaluation methods based on nearly 4,000 annotations from a domain-specific corpus of philosophy of science publications. Our findings reveal that word intrusion and coherence metrics do not always align, particularly in specialized domains, and that TWM captures human-perceived distinctness while appearing to align with diversity metrics. We release the annotated dataset and task generation code. This work highlights the need for evaluation frameworks bridging automated and human assessments, particularly for domain-specific corpora.

Keywords: topic model, human evaluation, word intrusion, topic coherence, topic diversity

1. Introduction

Topic models are unsupervised methods that aim to uncover latent thematic structures in document collections, representing topics as a distribution over words. Such methods can provide valuable insight into large corpora (Noichl, 2021; Malaterre and Lareau, 2023). As topic models evolve—from classical *Latent Dirichlet Allocation* (LDA, Blei et al. (2003)) to embedding-based and neural variants—evaluating and comparing their quality remains a fundamental challenge.

The field has largely converged on two evaluation paradigms: automated metrics and human judgments. Automated metrics, such as *topic coherence* and *topic diversity* are most often based on word co-occurrence patterns, and enable low-cost comparisons across methods and hyperparameters. However, their relationship with human evaluations of topic quality remains debated (Hoyle et al., 2021). Human evaluation methods, most notably Chang et al. (2009)'s *word intrusion*, offer a valuable starting point by directly assessing topic coherence, though their validation has largely focused on general-domain corpora. This raises a two-sided question: *to what extent do au-*

tomated metrics capture the full range of human judgments, beyond coherence, and in particular in domain-specific contexts?

A critical limitation of existing human evaluation tasks is their focus on *intra-topic* coherence—whether words within a topic are semantically, or distributionally related, through the lens of *word intrusion*. This provides an incomplete picture: a topic model can yield internally coherent topics that overlap significantly, failing to effectively partition the corpus. While automated diversity metrics attempt to measure *inter-topic* distinctness, their relationship to human judgment remains unclear. Moreover, specialized domains that involve jargon and abstract concepts pose unique challenges. In academic fields, like philosophy of science, evaluating whether topics capture meaningful conceptual distinctions requires careful consideration of domain-characteristics, often falling back on domain experts.

This paper makes three main contributions. First, we introduce *Topic Word Mixing* (TWM), a human evaluation task that assesses *inter-topic* distinctness by testing whether annotators can identify if a word-set originates from a single topic or is a mixture of two topics. This task complements

word intrusion's (TWI) focus on *intra-topic* coherence acting like a human pendant to topic diversity, providing direct human judgment on topic boundaries.

Second, we conduct a comprehensive evaluation comparing six topic models (LDA, NMF, Top2Vec, BERTopic, CFMF, CFMF-emb) using both automated metrics (coherence, diversity) and human evaluation (*word intrusion*, *topic word mixing*) totaling almost 4,000 annotations with 10 annotators on a domain-specific corpus spanning decades of philosophy of science publications. We release the dataset of annotated *word intrusion* and *topic word mixing* tasks, along with the code to generate such tasks, to foster the development of new evaluation metrics.

Third, we analyze the relationship between automated and human evaluation methods. Results reveal that *word intrusion* and *topic coherence* do not always align well, raising doubts about the reliability of coherence metrics in specialized corpora. This suggests that automated coherence measures based solely on domain-specific references may be insufficient and that broader, generalist contexts—such as those captured by large-scale embeddings—could improve evaluation. Consequently, these findings point to the need for developing embedding-based automated measures of coherence. Importantly, we find that constructing intrusion tasks in specialized domains requires additional considerations—particularly regarding the semantic similarity between intruder words and topic words—that have not been traditionally addressed. The human annotation results also indicate that *word mixing* appears to track topic diversity reasonably well, at least in the present corpus.

Section 2 reviews related work on topic model evaluation. Section 3 describes our corpus, automated metrics, *word intrusion* and our newly introduced *topic word mixing* human evaluations. Section 4 details our experimental setup. Section 5 reports results from both automated and human-centered evaluations. Section 6 discusses implications and limitations, Section 7 concludes.

2. Related Work

Evaluating topic models remains a fundamental challenge. Early work relied primarily on perplexity and likelihood measures as assessments of topic quality, which were shown to correlate poorly with human judgment (Blei et al., 2003; Wallach et al., 2009). *Topic coherence* metrics were thus introduced, measuring the semantic similarity between top words within a topic. Several coherence variants have been introduced with diverse reference similarity corpus and aggregation functions. NPMI (*Normalized Pointwise Mutual Information*)

measures word co-occurrence patterns and was shown to correlate with human judgment (Newman et al., 2010; Lau et al., 2014), while C_v (Röder et al., 2015) incorporates word vectors and sliding windows. UMass coherence uses document co-occurrence from the corpus itself. These metrics are widely used for model selection.

Beyond coherence, researchers have proposed additional metrics capturing other aspects of topic quality. Notably, *topic diversity* measures the proportion of unique words across topics, as a way to identify topic overlap and redundancy (Dieng et al., 2020; Tran et al., 2013). *Coverage* metrics measure how well topics represent the entire range of documents present in the corpus. However, the relationship between these metrics and human perception of topic quality is not established. Technical vocabulary and abstract concepts can present a challenge for automated metrics.

Given the limitations of automated metrics, human evaluation remains essential for validating topic model quality. Chang et al. (2009) introduced two influential tasks: *word intrusion* (TWI) and *topic intrusion*. In TWI, annotators identify an intruder among a topic's top words; success correlates with human perceived coherence. In *topic intrusion*, annotators are asked to identify which topic does not fit a given document. When conducted on two generalist corpora, *New York Times* and *Wikipedia* documents, these tasks revealed that models optimized for likelihood demonstrate less interpretable topics for humans.

Subsequent work has extended these evaluation paradigms. Newman et al. (2010) investigated topic quality by asking human to rate topics. Lau et al. (2014) have extended and automated *word intrusion* comparing it to automated coherence metrics, finding moderate but imperfect correlation; more recent work proposed hybrid methods bridging human and machine evaluation (Yang et al., 2025). Such paradigms have also been used to assess the interpretability of word embeddings (Murphy et al., 2012; Faruqui et al., 2015; Subramanian et al., 2018; Prouteau et al., 2022). However, these studies have predominantly focused on topic coherence and on general-domain corpora (news, Wikipedia), leaving open questions about evaluation and correlation of human-judgment with other automated metrics in domain-specific settings.

3. Evaluation Framework

3.1. Corpus

The corpus represents an extensive and comprehensive collection of the full texts from eight leading philosophy of science journals, spanning from

1931–the publication date of the earliest journal–through 2017 (*British Journal for the Philosophy of Science*, *European Journal for Philosophy of Science*, *Erkenntnis*, *International Studies in the Philosophy of Science*, *Journal for General Philosophy of Science*, *Philosophy of Science*, *Studies in History and Philosophy of Science Part A* and *Synthese*).

With a total of 16,917 articles and approximately 65 million word occurrences (23.7k types), this corpus includes both English and non-English items (6% of primarily German, Dutch, and French). Following [Malaterre and Lareau \(2022\)](#), foreign language documents were machine-translated into English. No further preprocessing was required by embedding-based methods. The input for remaining methods was tokenized, part-of-speech tagged and tokens were lemmatized. Low-frequency terms were removed by excluding words that appeared in fewer than 50 sentences across the corpus. To focus topic induction on semantically salient lexical items, only content words (nouns, verbs, adjectives, adverbs) were retained to produce vector representations needed by other topic models, such as word-frequency Term \times Document Matrix (TDM); normalized TF-IDF TDM; and Doc2Vec vectors.

3.2. Topic Models

Our goal is to evaluate topics across different modeling paradigms. We selected methods involving diverse approaches: probabilistic (*LDA*, *NMF*, *CFMF*), and embedding-based (*BERTopic*, *CFMF-emb*, *Top2Vec*). **Latent Dirichlet Allocation** (*LDA*) ([Blei et al., 2003](#)) treats each document as a mixture of topics, where each topic is a distribution over words, estimated via Gibbs sampling on a word-frequency term-document matrix (TDM). **CFMf** combines Feature Maximization ([Lamirel et al., 2016](#)), which selects topic-terms using an F-measure scheme on normalized TF-IDF vectors, with an adapted version of Growing Neural Gas clustering ([Fritzke, 1994](#)) using prototype renormalization ([Lamirel et al., 2025](#)) and exploiting normalized TF-IDF document vectors as input. **CFMf-emb** uses the same Feature Maximization but applies GNG with renormalization on normalized transformer-derived document embeddings. **Top2Vec** ([Angelov, 2020](#)) embeds words and documents using Doc2Vec ([Le and Mikolov, 2014](#)), discovers clusters via HDBSCAN ([Campello et al., 2013](#)) without prespecified K , and extracts top terms via cosine proximity to cluster centroids. **BERTopic** ([Grootendorst, 2022](#)) uses HDBSCAN on transformer embeddings, and applies class-based TF-IDF to rank words per topic. **Nonnegative Matrix Factorization** (*NMF*) ([Lee and Seung, 1999](#)) factorizes a TF-IDF matrix into

nonnegative term-topic and topic-document matrices via alternating least squares.

Models were built with $K = 10, 25$, and 50 topics. For *LDA*, *CFMF*, *CFMF-emb*, and *NMF*, K was set explicitly. For *Top2Vec* and *BERTopic*, effective K was indirectly controlled by tuning HDBSCAN’s minimum cluster size; this resulted in models with $K = 10, 25, 48$ for *Top2Vec* and $K = 10, 25, 54$ for *BERTopic*. Document embeddings are extracted (except for *Top2Vec* that uses *Doc2Vec*) using `stella_en_1.5B_v5` ([Zhang et al., 2025](#)), due to its performances for long texts (131,072+ tokens) ([Muennighoff et al., 2023](#)). All models produce ranked *top-W* word descriptors to describe each topic.

3.3. Automated Metrics

Coherence measures the interpretability of a topic by quantifying the semantic relatedness of its top words; higher coherence indicates that top terms tend to co-occur in the reference corpus. Coherence metrics include *PMI*, *NPMI*, *U-Mass*, and *Coherence C_V* . Here we used *Coherence C_V* because it is supposed to correlate well with human judgments on topic quality ([Röder et al., 2015](#)). Using a sliding window of size s (typically 110), C_V builds for each top word a smooth *NPMI*-based context vector over the set of N top words. It then aggregates these vectors into a topic vector, computes cosine similarity between each word vector and the topic vector (via segmentation), and reports the mean cosine across all words and topics of the corpus under investigation. The C_V score ranges from $[-1, 1]$, typically normalized to $[0, 1]$, with 1 indicating more coherent topics.

Topic diversity quantifies the distinctness of topics by measuring the overlap of their top terms; higher diversity means topics share fewer common top words. Given K topics and N top words for each topic, let $TopW_k$ be the set of top N words of topic k and $U = \bigcup_{k=1}^K TopW_k$ the set of unique top words across all topics. Topic diversity is: $\mathbf{TD} = \frac{|U|}{K \cdot N}$, it ranges from 0 (all topics share the same top words) to 1 (no top word is shared between topics).

3.4. Human Evaluation

Human evaluation is critical for assessing topic model quality, though scarcely implemented compared to automated metrics. [Chang et al. \(2009\)](#)’s seminal work introduced *word intrusion* and *topic intrusion* tasks. [Mimno et al. \(2011\)](#) proposed coherence measures correlating with human judgments, while [Röder et al. \(2015\)](#) systematically compared automated coherence metrics against human evaluations. [Lau et al. \(2014\)](#); [Yang et al. \(2025\)](#) have revisited these paradigms and provide

classifier and LLM-based approaches of *word intrusion* and *topic coherence rating*. While coherence metrics have been shown to correlate with human judgment across multiple studies, Hoyle et al. (2021) demonstrated that this alignment is imperfect and inconsistent. Furthermore, most validation work has been conducted on general-domain corpora, raising questions about generalization to specialized domains.

(Topic) Word Intrusion. (TWI) was originally introduced in Chang et al. (2009). Aimed at evaluating the coherence of topics, TWI adopts a contrastive approach by presenting the top ranked words (head) of a topic along with an intruder taken randomly, following certain conditions. The motivating idea is that a topic is coherent if humans are able to discriminate the intruder hidden among the top words of the topic. In cases where topics are hard to interpret, evaluators will typically choose intruders at random, which implies a low coherence for the topics presented. Although (Lau et al., 2014) found a correlation between human judgment in word intrusion and NPMI on classical corpora (*New York Times*, *Wikipedia*), Hoyle et al. (2021) have suggested that this correlation may not hold consistently across datasets or experimental setups.

Word intrusion has been frequently used to assess the quality of topic models. However, to the best of our knowledge, word intrusion tests have mainly been conducted on generalist corpora and rarely, if at all, on domain-specific corpora. One of the main obstacles to word intrusion evaluation is the cost of human annotations. This partly explains why researchers favor automated metrics. In addition, when expert knowledge is required, the search for annotators can be even more difficult.

Despite these obstacles, we conducted a *word intrusion* evaluation on the present domain-specific corpus to investigate how human evaluation of topic coherence relates to automated metrics in this context.

Topic Word Mixing. As noted above, topic diversity provides a complementary metric, reflecting how distinct topics are from one another. Unlike coherence, there is currently no human evaluation framework for assessing topic diversity. Automated measures of diversity typically rely on lexical overlap or similarity between topic-word distributions, making diversity an intrinsic property of the model rather than a perceptual one. We therefore propose to ground quantitative diversity metrics in human judgment, by viewing topic distinctness as a human-perceived proxy for topic diversity.

We thus introduce **Topic Word Mixing** (TWM), a human evaluation task aimed at assessing inter-topic distinctness. The task probes whether human annotators can determine if a given set of words originates from a single topic or from a mixture of two. Intuitively, if a topic model learns thematically diverse topics, humans should easily recognize mixed sets as heterogeneous. Conversely, when topics overlap semantically, this distinction becomes more difficult to make.

Formally, let $\mathcal{T} = \{\beta_1, \dots, \beta_K\}$ be a topic model and $W_i^{(m)}$ denote the set of the top- m words of topic i as ordered in β_i . For an even integer n , a *topic word mixing task* of size n is defined as

$$\mathcal{M}(n) = \begin{cases} (W_i^{(n/2)} \cup W_j^{(n/2)}, y = 2), & i \neq j, \\ (W_k^{(n)}, y = 1), \end{cases}$$

where $y \in \{1, 2\}$ is the class label: $y = 1$ indicates a single-topic set of n words, and $y = 2$ indicates a mixture of two topics with $n/2$ words each.

This formulation allows for a direct human-based assessment of topic distinctness, thus providing an interpretable measure of how separable the topics are in semantic space.

4. Experimental Setup

Automated Metrics We chose to measure topic model quality with two metrics: *coherence* C_V (Röder et al., 2015) and *diversity* (section 3.3). While these metrics enable efficient comparison across models and hyperparameters, their alignment with human judgment—particularly in specialized domains—remains an open question that our human evaluation tasks aim to address.

(Topic) Word Intrusion. In the TWI evaluation we conducted here, each word intrusion task consisted of **5 words**: the 4 top-words of a given topic (the head) and 1 intruder word. The intruder was selected according to the following criteria:

- It belongs to the lower 50% of the evaluated topic’s ranked words while appearing in the top 10% of words in at least one other topic. This ensures that the intruder is semantically distinct but still plausible in the corpus.
- Its average frequency in the corpus had to match that of the evaluated topic’s top words. This constraint ensures that annotators could not rely on word frequency as a cue.

To verify annotator attentiveness, we included simple control tasks in which the intruder word was highly salient. Finally, the word order within each task was randomized.

Topic Word Mixing. To assess topic distinctness, we propose to focus the TWM tasks on the most confusable topic pairs within each model. Such pairs can be identified by computing the cosine similarity between topic embeddings, where each topic embedding corresponds to the mean vector of the top 50 topic words encoded here with the `stella_en_1.5B_v5` model. For each topic, the most similar counterpart is selected to form candidate pairs for the TWM task. This procedure directs the evaluation toward subtle, fine-grained distinctions between semantically proximate topics rather than trivially distinct ones.

Each TWM consisted of **8 words**. For mixed-topic tasks involving 2 topics, these 8 words consisted of the top 4 words of each topic; for single-topic tasks, the 8 words consisted of the top 8 words of the sampled topic. The number of words (8) was chosen to balance informativeness and cognitive feasibility by providing enough context for judgment, while remaining close to short-term memory limits (Miller, 1956). Half of the tasks presented a mixture of two topics, and half presented words from a single topic, ensuring a balanced classification setting.

To mitigate cognitive load and ensure that the task primarily tested topic-level distinctness rather than annotator ability for grouping terms, half of the words in each task were bolded: in single-topic tasks, these bolded words were randomly chosen; in mixed-topic tasks, one of the two topics was represented by bold words. This visual cue relieves participants from inferring topic clusters and focuses their effort on determining whether the set contains one or two topics.

Word order within each task was randomized to prevent participants from relying on adjacency cues. For each model, topic pairs were sampled uniformly among the most similar pairs, ensuring diversity of topics across the evaluation set. Easily solvable control tasks were included to check for annotator attentiveness.

Annotators. Ten annotators—undergraduate philosophy students with training in epistemology and philosophy of science—were recruited to evaluate the topic models using TWI and TWM. All possessed relevant knowledge in the corpus domain and received financial compensation for their participation. The annotation process was distributed across the annotators, with five individuals assigned to TWI and five to TWM. To mitigate fatigue and maintain annotation quality, each task was divided into four tracks, allowing annotators to take breaks between segments.

Annotation Platform & Track Descriptions. The annotation tasks were conducted using Label-

How many topics are present amongst the following words ?

time
space
body
galileo
relativity
theory
newton
motion

1st

2nd

Figure 1: Example of a two-topic Word Mixing task in *LabelStudio*.

Studio’s online platform (Tkachenko et al., 2020-2025). Annotators accessed four separate tracks for each task, with all items randomized prior to presentation to prevent order effects. For the word intrusion task (TWI), annotators were shown five words and asked to identify the single intruder word that did not belong with the others. The word mixing task (TWM) interface presented fig. 1 has the annotators select “1” or “2” perceived topics.

Human Evaluation Dataset. Based on these criteria, a total of 18 models—six topic modeling methods trained with approximately 10, 25 and 50 topics each) were evaluated using both TWI and TWM. For TWI, all topics from models of size 10 were annotated, while for models of size 25 and 50, 60% of topics were randomly sampled for annotation. This resulted in 330 distinct tasks, amounting to 1650 annotations across five annotators. For TWM, both single-topic and mixed-topic tasks needed to be sampled. Accordingly, all models with 10 topic were fully included, while 50% of tasks were uniformly randomly sampled for models with 25 and 50 topics. In total, 467 tasks were evaluated across the 18 models, yielding 2335 annotations. The code and annotation interface used to generate and sample intrusion and mixing tasks from topic models are open-sourced on Github¹.

5. Results

Automated Metrics. Table 1 details the coherence and diversity score for each model to be annotated in TWI and TWM. Regarding coherence, Top2Vec achieves the best results regardless of the granularity of the topic model, ranging from 0.72 to 0.80 depending on K . Feature maximization-based methods (CFMF, CFMF-emb) follow with slightly more moderate coherence (0.67-0.72) while NMF shows lower scores (0.59-0.67).

¹https://github.com/thibaultprouteau/LREC_2026_Proutreau_Human_Metrics_Topic_Models

BERTopic and LDA exhibit the weakest coherence, with LDA scoring the lowest across all configurations (0.47-0.55).

Concerning topic diversity, Top2Vec leads again, maintaining near-perfect scores (0.98-1), indicating minimum overlap in top words across topics. CFMF, CFMF-emb, and NMF show relatively high diversity for models with 10 topics (0.90-0.96) but experience notable decline with a greater number of topics, dropping to 0.73-0.81 for $K = 50$. This pattern—decreasing diversity with increasing K —is consistent across models, suggesting that as the number of topics grows, they tend to exhibit greater overlap. LDA and BERTopic show moderate diversity (0.66-0.80). Despite being based on embeddings like CFMF-emb, BERTopic does not produce highly diverse topics.

$K \approx$	Coherence-10			Diversity-10		
	10	25	50	10	25	50
Top2Vec	0,72	0,79	0,80	1,00	1,00	0,98
CFMF	0,69	0,72	0,72	0,96	0,85	0,81
CFMF-emb	0,67	0,69	0,70	0,90	0,79	0,73
NMF	0,59	0,65	0,67	0,90	0,88	0,78
BERTopic	0,52	0,58	0,62	0,79	0,70	0,66
LDA	0,47	0,54	0,55	0,80	0,74	0,66

Table 1: Coherence C_V and diversity (10 words) of each topic model.

Topic Word Intrusion. The *inter-annotator agreement* reaches an average of 80.3% and a Fleiss’ Kappa (Fleiss and Cohen, 1973) of 0.688 which indicates a *substantial agreement* (Landis and Koch, 1977).

Figure 2 presents the average of annotation task success rates across all model configurations (model_K). Model performance varies considerably within model families depending on model size K . BERTopic_10 achieves perfect accuracy (1.00) while BERTopic_54 drops to 0.82. Similarly, CFMF-emb_10 reaches 0.90, but CFMF-emb_51 falls to 0.75. This pattern of higher accuracy for smaller models seems to hold for most models (except lda_10), indicating that annotators find topics more coherent when K is small. However, LDA_54 outperforms many models at similar topic numbers, and Top2Vec consistently underperforms regardless of K , with Top2Vec_25 scoring lowest at 0.61. It is worth noting that confidence intervals remain large despite the overall substantial agreement between annotators. Yet, Top2Vec is subpar, in comparison to BERTopic_10 and LDA_54.

The aggregated results per model family are presented in Figure 3. As shown, BERTopic achieves the highest performance on TWI (0.900) followed by CFMF-emb, NMF and LDA. The CFMF

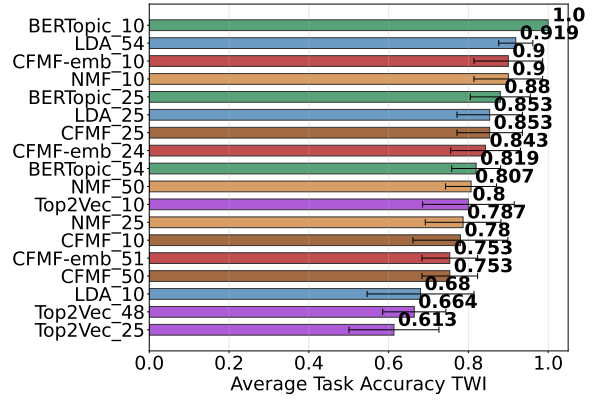


Figure 2: Macro average annotation accuracy for TWI per model (95% confidence intervals).




Figure 3: Macro average annotation accuracy for TWI per model family.

variant, which does not use embeddings, attains a lower accuracy (0.796) compared to its embedding-based counterpart. Finally, Top2Vec yields the lowest overall performance among all model families (0.693).


Traditionally, TWI tasks are constructed by choosing words from the topic top-terms and, most importantly, by selecting an intruder based on word ranking constraints. However, relying solely on word ranking for intruder selection appears imperfect. Indeed, a posteriori examination of the distribution of average cosine similarity between top words and intruder (computed using embeddings from stella_en_1.5B_v5), reveals clear discrepancies, as depicted in Figure 4a. This matters since, intuitively, if an intruder is semantically close to the topic top terms, the task should become more difficult. In this respect, models such as BERTopic and Top2Vec are likely advantaged, whereas CFMF-emb may have received more challenging tasks overall. Average task accuracy is inversely correlated with the mean similarity between top-words and intruder ($r = -0.145, p = 0.0079$).

Results are then adjusted, based on the residuals of tasks to the fitted slope, to correct for task difficulty. In the adjusted results fig. 4b, CFMF-emb, LDA and CFMF gain accuracy by alleviating the impact of intruder to topic head similarity. NMF and LDA decrease slightly and BERTopic remains the


leader, but loses 1.35 percentage point of accuracy. Though such adjustment likely corrects for task difficulty, the overall ranking of model families does not change here.

When comparing TWI annotation accuracy with coherence C_V , we can see that $Top2Vec$ human perceived coherence doesn't align with automated measurement. A similar finding holds for $BERTopic$ that has comparatively lower C_V coherence but achieves higher accuracy TWI . Thus, automated co-occurrence based coherence measures may not fully capture human interpretability, at least not as well in specialized domains as in more generic corpora. This divergence underscores the complimentary nature of human evaluation to automated metrics in topic modeling.

Topic Word Mixing The *inter-annotator agreement* for TWM reaches an average of 77.1% and a Fleiss' Kappa of 0.270, which is a *fair agreement* and is a testament to the difficulty of the task. Indeed, we chose to mix topics that are the closest for which the decision boundary may be thin. Figure 5 presents the mixing results per model. $Top2Vec$ maintains good performances across granularities (0.698-0.942) with $Top2Vec_{10}$ scoring highest. The $CFMF-emb_{10}$ and NMF_{10} variants follow closely (0.87). $BERTopic$ and $CFMF$ exhibit more variabil-



(a) Cosine similarity between topic top-4 words and intruder for TWI per model family.



(b) Adjustment of macro average annotation accuracy scores for TWI per model family taking into account the similarity between intruders and topic top-4 words.

Figure 4: Adjusting results to account for task difficulty (intruder-top-words similarity).

ity across model sizes. Confidence intervals remain significant for some models for which detecting two-topic tasks was lower.

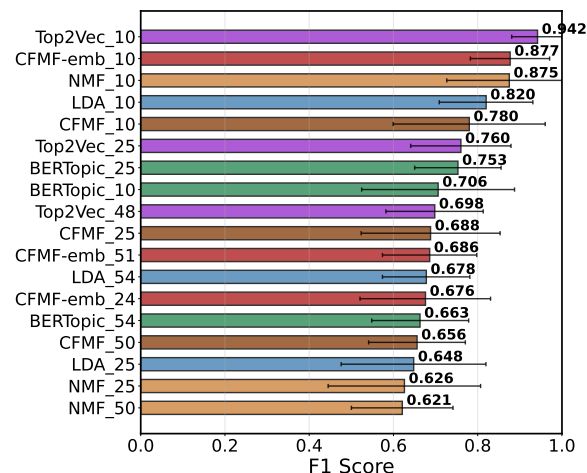


Figure 5: Macro average task F1 Score for TWM per model (with 95% confidence intervals).

When considering families of models (Figure 6), $Top2Vec$ achieves the best score (0.800) followed by $CFMF-emb$ (0.746). LDA , $CFMF$, NMF and $BERTopic$ obtain similar results.




Figure 6: Macro average task F1 Score for TWM per model family.

Compared with automated measures of *diversity*, the TWM results show a consistent pattern. $Top2Vec$ has near perfect diversity scores (0.98-1.00) and outperforms all other models on TWM . Interestingly, while $CFMF-emb$ shows lower quantitative diversity than $CFMF$, it achieves higher human-perceived diversity. Nonetheless, the contrast that was apparent between models in terms of diversity (table 1) seems less pronounced in the TWM evaluation.

Note that the TWM results also align with coherence measures for $Top2Vec$ and $CFMF-emb$. This suggests that *Topic Word Mixing* may complement *Topic Word Intrusion* in more ways than just adding an inter-topic dimension. Indeed, for two mixed topics to be perceived as distinct, each must first be internally coherent; conversely, when a task involves a single topic whose words lack coherence, it becomes more difficult for annotators to recognize it as such. Thus, TWM indirectly captures as-

pects of both topic distinctness and internal coherence, reinforcing its value as a complementary human evaluation method.

Another more general remark is that when considering TWI and TWM scores jointly and producing an average ranking of the methods based on that combination, $CFMF-emb$ distinguishes itself from others.

6. Discussion

Misalignment between automated metrics and human judgment. Our results reveal misalignment between automated coherence and human annotation with the *word intrusion* on this domain-specific corpus. $Top2Vec$ achieved the highest coherence ($C_V=0.72-0.80$) and near perfect diversity (0.98-1.00), yet performs poorest on TWI (0.693). Conversely, $BERTopic$ excels at TWI (0.900) despite having the second-lowest automated coherence (0.52-0.62). These discrepancies suggest that co-occurrence metrics may not reliably capture human-perceived coherence in specialized domains, where technical vocabulary and abstract concepts pose unique challenges.

Performance under joint human evaluation. Large confidence intervals limit firm conclusions about model performance. With this caveat in mind, $CFMF-emb$ achieves the best overall performance (mean rank = 2) with balanced performance across both tasks. NMF and LDA follow (mean rank = 4 and 3.5) with moderate performance. $CFMF$ ranks lower (mean rank = 4.5), suggesting potential benefits from incorporating word embeddings. Finally, $BERTopic$ and $Top2Vec$ show somewhat different profiles, with $BERTopic$ appearing to favor coherence and $Top2Vec$ appearing to favor diversity, though both show discrepancies with their respective metric predictions.

What does topic word mixing capture? The TWM task appears to capture topic distinctness, at least partially aligning with the diversity metric. However, it likely probes richer aspects of topics. We hypothesize that TWM indirectly captures both diversity and some coherence: if topics are incoherent, distinguishing mixed top-words becomes difficult; conversely, coherent and distinct topics yield clearer mixing judgments.

Granularity effects remain unclear. Although we relied on annotators with domain expertise, our annotations provide insufficient evidence to draw strong conclusions about the effects of model granularity (K) on human-perceived quality (e.g. LDA_{54} outperforms LDA_{10}). Far more annotations would be required to establish clear trends.

Consequently, the choice of K should be guided by research goals and human appraisal. For instance, studies seeking to map the general thematic structure of a corpus—such as the overall landscape of philosophy of science—may be adequately served by models with relatively few topics (e.g. 10–25). Conversely, investigations targeting narrowly defined themes, such as specific philosophical theories of scientific explanation, may require models with a larger number of topics to ensure that at least one captures the phenomenon of interest.

Task difficulty considerations. Task difficulty should be considered when designing human-based evaluations. Our a posteriori analysis revealed that intruder-to-topic similarity (measured via word embedding) significantly affects TWI . Future work should control for such similarity a priori to ensure balanced task difficulty across models. This would help prevent confounding task difficulty with model quality.

Toward embedding-based measures. The embedding-based similarity we use to assess task difficulty creates an opportunity to revisit prior evaluations (Aletras and Stevenson, 2013) and to develop new coherence metrics grounded in distributional semantics. Traditional coherence relies on corpus-internal-coherence (whether it is measured on the corpus at hand or a more general corpus). While they are appropriate for capturing domain-specific patterns, they are limited by corpus boundaries. Large pre-trained embeddings encode broader world knowledge from vast corpora, potentially offering richer semantic representations. An externalized coherence measure based on embedding similarity between topic words could complement corpus-internal metrics like C_V coherence leveraging the broader embedded lexical and conceptual knowledge. Since our annotators possess, to some degree world knowledge and domain expertise, embedding may operationalize coherence in ways that better align with human judgment.

Scope limitations. Coherence and diversity likely capture only two aspects of topic quality. Models can yield topics that are both highly coherent and clearly distinct, yet fail to represent the variety of topics present in the corpus by focusing on narrow, salient themes. For this reason, coverage measures—assessing how well topics capture the full range of documents—should complement coherence and diversity evaluations, whether they are metric or human-based. Unfortunately, coverage is difficult to assess qualitatively through human evaluation due to the usual scale

of corpora. Consequently, quantitative coverage measures remain indispensable complements to human-centered evaluations.

7. Conclusion

Our results challenge the assumption that automated topic coherence metrics and human word intrusion evaluations are strongly correlated. In our domain-specific setting, coherence scores (e.g., C_V) and word intrusion accuracy often diverged; standard metrics may not reliably reflect human perception of coherence in specialized fields. This calls for caution when relying solely on automated coherence measures, especially outside general-domain corpora. To complement word intrusion, we introduced *Topic Word Mixing* (TWM), a new human evaluation probing inter-topic distinctness. TWM offers a direct, human-centered view of topic diversity, complementing diversity metrics. Our results show that automated diversity metrics partially align with human judgments. Looking forward, our study underscores the importance of combining multiple evaluation perspectives—automated and human-centered—to obtain a more complete picture of topic model quality. We release our annotated datasets of almost 4,000 examples and our code to support further research in this direction. Future work should explore embedding-based coherence measures to better align with human evaluations.

8. Acknowledgements

J.-C.L. & T.P. acknowledge funding from AID-RAPID SUMINO. F.L. acknowledges funding from Canada Social Sciences and Humanities Research Council (Postdoctoral Fellowships 756-2024-0557). C.M. acknowledges funding from Canada Social Sciences and Humanities Research Council (Grant 430-2018-00899) and Canada Research Chairs (CRC-950-230795). N.D. acknowledges support from Agence Nationale de la Recherche (ANR) through the DIG-ING project (ANR-21-CE23-0010). The authors acknowledge that they were granted a free academic license of “Label Studio” by HumanSignal, available at <https://labelstud.io>. Experiments presented in this paper were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>).

9. Bibliographical References

- Nikolaos Aletras and Mark Stevenson. 2013. Evaluating topic coherence using distributional semantics. In *Proceedings of the 10th international conference on computational semantics (IWCS 2013)—Long Papers*, pages 13–22.
- Dimo Angelov. 2020. Top2vec: Distributed representations of topics. *arXiv preprint arXiv:2008.09470*.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Ricardo JGB Campello, Davoud Moulavi, and Jörg Sander. 2013. Density-based clustering based on hierarchical density estimates. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 160–172. Springer.
- Jonathan Chang, Sean Gerrish, Chong Wang, Jordan Boyd-graber, and David Blei. 2009. Reading Tea Leaves: How Humans Interpret Topic Models. In *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc.
- Adji B Dieng, Francisco JR Ruiz, and David M Blei. 2020. Topic modeling in embedding spaces. *Transactions of the Association for Computational Linguistics*, 8:439–453.
- Manaal Faruqi, Yulia Tsvetkov, Dani Yogatama, Chris Dyer, and Noah A. Smith. 2015. *Sparse Overcomplete Word Vector Representations*. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1491–1500, Beijing, China. Association for Computational Linguistics.
- Joseph L Fleiss and Jacob Cohen. 1973. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and psychological measurement*, 33(3):613–619.
- Bernd Fritzke. 1994. A growing neural gas network learns topologies. *Advances in neural information processing systems*, 7.
- Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.
- Alexander Hoyle, Pranav Goel, Andrew Hian-Cheong, Denis Peskov, Jordan Boyd-Graber, and Philip Resnik. 2021. Is automated topic model evaluation broken? the incoherence of

- coherence. *Advances in neural information processing systems*, 34:2018–2033.
- Jean-Charles Lamirel, Nicolas Dugué, and Pascal Cuxac. 2016. New efficient clustering quality indexes. In *2016 International joint conference on neural networks (IJCNN)*, pages 3649–3657. IEEE.
- Jean-Charles Lamirel, Francis Lareau, and Christophe Malaterre. 2025. Old but not obsolete: Bag-of-words vs. embeddings in topic modeling. In *20th International Conference on Scientometrics & Informetrics*, pages 2275–2282.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Jey Han Lau, David Newman, and Timothy Baldwin. 2014. [Machine Reading Tea Leaves: Automatically Evaluating Topic Coherence and Topic Model Quality](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 530–539, Gothenburg, Sweden. Association for Computational Linguistics.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196. PMLR.
- Daniel D Lee and H Sebastian Seung. 1999. Learning the parts of objects by non-negative matrix factorization. *nature*, 401(6755):788–791.
- Christophe Malaterre and Francis Lareau. 2022. The early days of contemporary philosophy of science: Novel insights from machine translation and topic-modeling of non-parallel multilingual corpora. *Synthese*, 200(3):242.
- Christophe Malaterre and Francis Lareau. 2023. The emergence of astrobiology: a topic-modeling perspective. *Astrobiology*, 23(5):496–512.
- George A Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81.
- David Mimno, Hanna Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. 2011. Optimizing semantic coherence in topic models. In *Proceedings of the 2011 conference on empirical methods in natural language processing*, pages 262–272.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive text embedding benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia. Association for Computational Linguistics.
- Brian Murphy, Partha Talukdar, and Tom Mitchell. 2012. Learning Effective and Interpretable Semantic Models using Non-Negative Sparse Embedding. In *Proceedings of COLING 2012*, pages 1933–1950, Mumbai, India. The COLING 2012 Organizing Committee.
- David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin. 2010. Automatic evaluation of topic coherence. In *Human language technologies: The 2010 annual conference of the North American chapter of the association for computational linguistics*, pages 100–108.
- Maximilian Noichl. 2021. Modeling the structure of recent philosophy. *Synthese*, 198(6):5089–5100.
- Thibault Prouteau, Nicolas Dugué, Nathalie Camelin, and Sylvain Meignier. 2022. Are embedding spaces interpretable? results of an intrusion detection evaluation on a large french corpus. In *LREC 2022*.
- Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. Exploring the space of topic coherence measures. In *Proceedings of the eighth ACM international conference on Web search and data mining*, pages 399–408.
- Anant Subramanian, Danish Pruthi, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Eduard Hovy. 2018. SPINE: SParse Interpretable Neural Embeddings. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Maxim Tkachenko, Mikhail Malyuk, Andrey Holmanyuk, and Nikolai Liubimov. 2020–2025. [Label Studio: Data labeling software](#). Open source software available from <https://github.com/HumanSignal/label-studio>.
- Nam Khanh Tran, Sergej Zerr, Kerstin Bischoff, Claudia Niederée, and Ralf Krestel. 2013. Topic cropping: Leveraging latent topics for the analysis of small corpora. In *International Conference on Theory and Practice of Digital Libraries*, pages 297–308. Springer.
- Hanna M Wallach, Iain Murray, Ruslan Salakhutdinov, and David Mimno. 2009. Evaluation methods for topic models. In *Proceedings of the 26th annual international conference on machine learning*, pages 1105–1112.

Xiaohao Yang, He Zhao, Dinh Phung, Wray Buntine, and Lan Du. 2025. [LLM Reading Tea Leaves: Automatically Evaluating Topic Models with Large Language Models](#).

Dun Zhang, Jiacheng Li, Ziyang Zeng, and Fulong Wang. 2025. [Jasper and stella: distillation of sota embedding models](#).