

What Matters to an LLM? Behavioral and Computational Evidences from Summarization

Yongxin Zhou, Changshun Wu, Philippe Mulhem, Didier Schwab, Maxime Peyrard

Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, 38000, Grenoble, France

{firstname.lastname}@univ-grenoble-alpes.fr

Abstract

Large Language Models (LLMs) are now state-of-the-art at summarization, yet the internal notion of importance that drives their information selections remains hidden. We propose to investigate this by combining behavioral and computational analyses. Behaviorally, we generate a series of length-controlled summaries for each document and derive empirical importance distributions based on how often each information unit is selected. These reveal that LLMs converge on consistent importance patterns, sharply different from pre-LLM baselines, and that LLMs cluster more by family than by size. Computationally, we identify that certain attention heads align well with empirical importance distributions, and that middle-to-late layers are strongly predictive of importance. Together, these results provide initial insights into *what* LLMs prioritize in summarization and *how* this priority is internally represented, opening a path toward interpreting and ultimately controlling information selection in these models.¹

1 Introduction

Large Language Models (LLMs) are increasingly entrusted with the management of information: they filter, select, and summarize our textual information. This shift raises a fundamental question: *what do LLMs consider important?* Summarization offers a natural entry point to this question, since the core challenge of summarization is the identification and selection of *important information*.

Before the rise of LLMs, summarization systems were designed around simple, surface-level heuristics that correlated with importance as reflected in benchmark datasets. For example, centrality measures (Zopf et al., 2016), lexical overlap (Luhn, 1958), and frequency of mention (Edmundson, 1969; Nenkova and Vanderwende, 2005)

¹Our code and data are available at <https://github.com/yongxin2020/llm-importance-summarization>.

were strong predictors of which content would be selected. Remarkably, the crude baseline of taking the first few sentences of a news article remained one of the hardest to beat for years (See et al., 2017; Xing et al., 2021). LLMs radically changed this landscape. Today, summarization is largely accomplished by prompting an LLM, yielding fluent, flexible, and adaptable summaries with little or no task-specific engineering (Zhang et al., 2024). Yet this success comes at the cost of transparency: unlike extractive methods that explicitly score information units (McDonald, 2007; Gillick et al., 2008; Li et al., 2013; Peyrard and Eckle-Kohler, 2016), LLM-based summarization provides no clear account of the internal notion of importance driving its selections. Understanding this hidden notion of importance is both scientifically and practically critical, as it speaks to how LLMs structure and prioritize knowledge (Carter et al., 2019; Cammarata et al., 2020; Meng et al., 2022; Geva et al., 2023; Monea et al., 2024) in ways that increasingly shape human access to information.

In this work, we take a two-pronged interpretability approach to this question visualized in Figure 1. With **behavioral analysis**, we first study model outputs themselves. For any given input document, we generate 10 summaries of varying lengths, then track how frequently each information unit (e.g., token) is selected across summaries. This yields an empirical importance distribution: an estimate of what the model consistently prioritizes when forced to compress information at different lengths. With the **computational analysis**, we then turn inward to the model’s computation, asking whether and where this empirical importance is represented. Attention mechanisms, central to the Transformer architecture, are themselves distributions over tokens and can be interpreted as mechanisms of importance weighting (Kobayashi et al., 2020). We therefore test whether specific attention heads align with the behavioral importance distributions. Be-

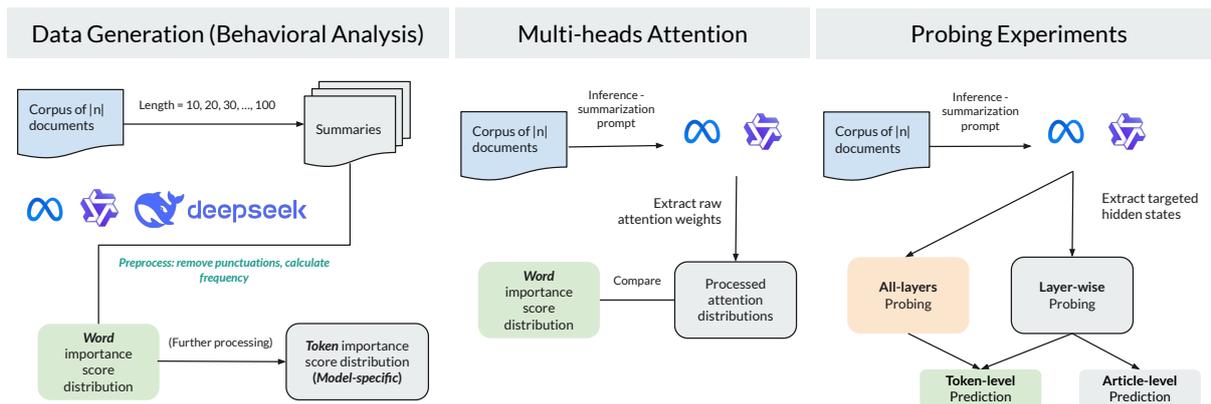


Figure 1: Analytical framework for modeling *information importance*. (1) **Behavioral Analysis:** We generate length-variant summaries with LLMs across three datasets (CNN/DailyMail, SAMSum, DECODA-French). The *importance distribution* $I_M(D)$ is derived as summary persistence. (2) **Attention Analysis:** Raw attention weights are aggregated and normalized to obtain token-level distributions. (3) **Probing:** Hidden states are used to train probes in three scenarios (S1: Layer-wise/Token, S2: All-layers/Token, S3: Layer-wise/Article) to predict $I_M(D)$.

yond attention, we probe layer activations: using a low-capacity linear probe specifically designed to output importance scores for all tokens simultaneously. This setup mitigates common risks of overfitting in probing studies (Ravichander et al., 2021; Kumar et al., 2022; Teney et al., 2022). Furthermore, to avoid false discoveries we include a “dead salmon” probe baseline by comparing probe predictions against a probe trained on a randomized version of the models (Méloux et al., 2025).

Contributions. We release a dataset of 274,330 summaries generated by 7 LLMs, preprocessed to obtain the empirical importance distributions. Our behavioral analysis shows that LLMs exhibit broadly similar summarization behavior, sharply diverging from pre-LLM baselines. Similarity is influenced more by model family than by scale. Our computational analysis demonstrates that certain attention heads alone capture substantial aspects of the importance distribution, and that middle-to-late layers of the network are highly predictive of behavioral importance. This work initiates a research direction to better understand, and ultimately control, the latent notion of information importance encoded within the computational structure of LLMs.

2 Related Work

2.1 Information Importance in Summarization

The core task of summarization is to select and condense salient information from a source document,

a process that fundamentally requires learning a latent representation of *information importance* – an abstract, multi-faceted, and context-dependent concept (Narayan et al., 2018; Kedzie et al., 2018).

Historically, *importance* was operationalized through surface-level statistical features for extractive methods, such as term frequency and sentence position (Luhn, 1958), or measured post-hoc by n -gram overlap metrics like ROUGE against a reference (Lin, 2004). While these approaches facilitated practical solutions, they captured signals that merely correlate with human intuition. For instance, structural features like centrality and repetitions remain common proxies (Kedzie et al., 2018), but their weaknesses are exposed by simple adversarial attacks (Zopf et al., 2016).

Theoretical efforts provided a more formal foundation. Early work treated *importance* as a latent variable optimized indirectly via system performance (Nenkova and McKeown, 2012). A deeper formalization grounds it in information theory, representing texts as probability distributions over semantic units (Bao et al., 2011). This view, compatible with distributional embeddings, allows information-theoretic tools to operate at a semantic level (Carnap and Bar-Hillel, 1952; Zhong, 2017). Peyrard (2019) crystallizes this framework by formally defining *importance* through the unified concepts of redundancy, relevance, and informativeness.

Empirical studies investigate how models operationalize this latent construct. For instance, behav-

ioral probes using length-controlled summarization show that LLMs develop a nuanced, hierarchical understanding of salience, though this internal representation shows weak correlation with human judgment and is not directly accessible via introspection (Trienes et al., 2025). The effective definition of *importance* is not universal but is contingent on the summarization task and its communicative goals (Zhou et al., 2025). It is therefore shaped by conversational dynamics and action-oriented objectives in dialogues (Zhou et al., 2024; Ghebriout et al., 2025) and by event centrality in news (Li et al., 2016; Zhang et al., 2024).

A key open question remains how this abstract concept is mechanistically *encoded* within models. Our work addresses this gap through a multi-faceted analysis: (1) a **behavioral** study of summary outputs, (2) an examination of **attention** mechanisms, and (3) **probing** of hidden state representations to trace how a model’s own output-based importance signal is constructed across layers.

2.2 Behaviors and Biases in Summarization

Research on summarization behaviors and biases investigates what content models prioritize and how faithfully they reproduce it. A well-documented behavioral bias is **positional bias**, where models under-attend to middle content, creating a “U-shaped” trend in faithfulness for long-form summarization (Wan et al., 2025) that also impacts conversational summarization by affecting how models handle information based on its location in a dialogue (Sun et al., 2025).

Beyond content selection, studies analyze biases in how content is generated. This includes analyzing **social biases** in summaries, where controlled evaluations isolate model bias from source bias to measure issues like demographic or gender skews in generated content (Steen and Markert, 2024). It also includes the analysis of **faithfulness and factuality** (Maynez et al., 2020; Goyal and Durrett, 2021), a core challenge where LLMs exhibit distinct error patterns, such as generating plausible “circumstantial” inferences in dialogues unsupported by direct evidence (Ramprasad et al., 2024).

While existing research documents *what* models output, including systematic biases like positional effects, the internal computational origins of these behaviors remain unclear. Our work investigates these internal mechanisms alongside a behavioral analysis of summarization outputs.

2.3 Interpretability and Probing for Analysis

The multi-head attention mechanism, introduced by Vaswani et al. (2017), enables transformers to attend to information from different representation subspaces. Analyses indicate that attention heads often specialize in specific linguistic phenomena (Voita et al., 2019), and their distributions have consequently been interpreted as a form of token importance weighting and used as a window into model behavior and decision-making (Kobayashi et al., 2020; Li et al., 2022).

A parallel research direction seeks to understand the rich information encoded within LLMs’ hidden representations, which frequently surpasses what is expressed in their explicit outputs (Burns et al., 2023). For instance, probing classifiers can decode latent knowledge such as pre-encoded plans for future response (Dong et al., 2025), internal signals of factual inaccuracies (Orgad et al., 2025), or layer-wise capabilities for simulating personality traits, which can subsequently be leveraged to edit the personality expressed by LLMs during inference (Ju et al., 2025).

While existing work has investigated attention specialization and applied probing to tasks like syntax analysis or fact-checking, the mechanistic representation of *summarization-specific* importance remains less explored. Our work bridges these areas by analyzing internal representations to trace how a model’s own operational definition of *importance* is computed and encoded throughout the transformer architecture.

3 Methodology

Let M be a fixed transformer-based language model and let $D = (u_1, \dots, u_n)$ denote a document decomposed into a set of atomic *information units* (e.g., sentences, discourse units, or tokens in the rest of the paper). We introduce a model-dependent *importance distribution* $I_M(D)$, which captures the relative importance assigned by model M to the information units in D .

Formally, $I_M(D)$ is a probability distribution over the units of D , where $I_M(D)_j$ reflects the propensity of model M to include unit u_j when constrained to produce a summary of limited length. This distribution is intended to encode the model’s implicit trade-offs and preferences when selecting a small subset of units to represent the document.

Because $I_M(D)$ is not directly observable, we estimate it empirically via repeated conditional gen-

eration. Given a document D , we prompt M to generate k summaries of varying target lengths, with $k = 10$ in our experiments. Let $S^{(\ell)} = \{s_1^{(\ell)}, \dots, s_k^{(\ell)}\}$ denote the set of summaries generated at length constraint ℓ .

For each information unit u_j , we compute its empirical selection frequency

$$\hat{I}_M(D)_j = \frac{1}{k} \sum_{i=1}^k \mathbb{I}[u_j \in s_i], \quad (1)$$

where $\mathbb{I}[\cdot]$ is the indicator function and s_i ranges over all generated summaries across lengths. Intuitively, units that appear consistently in shorter summaries and persist in longer ones receive higher importance scores. This principle is supported by Pyramid-based evaluation methods, which also count the frequency of information in human reference summaries as a proxy for importance (Nenkova et al., 2007).

Using the estimated importance distribution $I_M(D)$, we pursue three complementary objectives: (i) To analyze the behavioral patterns of summarization models by studying the properties of $I_M(D)$ across datasets (Section 4); (ii) To investigate the relationship between empirical importance $I_M(D)$ and the model’s internal attention mechanisms (Section 5); (iii) To assess whether importance information is encoded in the model’s hidden state representations via probing methods (Section 6).

3.1 Datasets

We conduct experiments on two core summarization datasets to derive and analyze the empirical importance distribution. Their contrasting characteristics allow us to evaluate whether patterns of importance encoding are consistent across different genres and summary styles.

- **CNN/DailyMail** (See et al., 2017): An English news dataset containing over 300k journalist-written articles. We use version 3.0.0², originally designed for reading comprehension and question answering but widely adopted for extractive and abstractive summarization.
- **SAMSum** (Gliwa et al., 2019): A dataset of $\sim 16k$ messenger-style English conversations with linguist-created summaries. It includes 14,732 training, 818 validation, and 819 test

²https://huggingface.co/datasets/ccdv/cnn_dailymail

samples, providing a testbed for abstractive summarization of informal, interactive text.

Multilingual Extension: To assess the generalizability of our findings across languages, we additionally evaluate on the **DECODA** corpus (Favre et al., 2015), a French call center dialogue dataset. The corresponding experiments and analyses are provided in Appendix A due to space constraints.

3.2 Models

We evaluate the following models: two open-weight model families and one commercial model.

-  Llama-3.2-1B-Instruct, Llama-3.1-8B-Instruct
-  Qwen2.5-1.5B-Instruct, Qwen2.5-3B-Instruct, Qwen2.5-7B-Instruct, Qwen2.5-14B-Instruct
-  deepseek-chat³

The models and their corresponding links are detailed in Appendix C.2.

3.3 Data Generation

We generated data using 3,000 samples from the CNN/DailyMail test set, the full SAMSum test set (819 samples), and the full DECODA test set (100 samples). For each input, we generated $k = 10$ summaries of varying lengths using length-variant prompts (see Table 1, where $N \in \{10, 20, 30, \dots, 100\}$). We then compute the empirical importance distribution for each model and each document using the formula described in Equation 1.

3.4 Metric Selection and Validation

To compare importance distributions, several classes of metrics are applicable. These include (i) distributional measures such as Kullback–Leibler divergence, Jensen–Shannon divergence, and Wasserstein distance; (ii) list comparison measures such as Pearson and Spearman correlations; and (iii) information retrieval metrics that emphasize agreement among top-ranked units, including nDCG@ k . In the main experiments, we focus on two complementary metrics: Spearman correlation and nDCG@10. When a specific reference importance distribution is available (e.g., in attention-based and probing analyses), we report nDCG@10. In settings without a designated reference distribution, we use the symmetric Spearman correlation as a notion of similarity. However, we evaluate a

³Experiments conducted in September 2025, the model points to non-thinking mode of DeepSeek-V3.1.

total of 14 comparison metrics, and Appendix B provides more details on our choice of Spearman and nDCG. Importantly, our findings are qualitatively robust to the choice of metric. Results obtained with alternative metrics are reported in the Appendix for behavioral analysis (Appendix D), multi-head attention analysis (Appendix E), and probing experiments (Appendix F).

4 Behavioral Analysis

This section analyzes how the *summary persistence* importance distribution $I_M(D)$ varies across model families and scales. We quantify inter-model similarities and identify recurring statistical patterns in model outputs to characterize their shared notion of importance. Additional analyses of positional bias and entropy are provided in Appendix D.2 and Appendix D.3, respectively.

4.1 Summarization Baselines

To compare LLM summarization behavior, we implement three pre-LLM baselines:

Baseline 1: First-N-Words Frequency. Simulates lead bias by calculating word frequency across ten document truncations (first 10, 20, ..., 100 words). This tests if LLMs prioritize early content beyond simple heuristics.

Baseline 2: Token Frequency. Estimates importance using raw word counts normalized by the document’s maximum word frequency, serving as a basic statistical baseline.

Baseline 3: TextRank. Extracts and scores keywords using the TextRank algorithm (Mihalcea and Tarau, 2004), with scores normalized to the range [0, 1] via min-max normalization.

We also introduce **Human Frequency**, which assigns importance scores based on word presence in ground-truth summaries⁴.

4.2 Model Behavioral Similarity Analysis

To analyze similarity in importance distributions, we construct a union vocabulary from all words appearing in any model’s summaries for each document. For each document, words present in one model’s importance distribution but absent from another’s were assigned an importance score of zero, ensuring valid probability distributions.

We calculate pairwise model *dissimilarity* using Spearman rank correlation distance, defined as

⁴Note: This serves as a **proxy** for human-annotated importance derived from reference summaries, rather than direct human annotation of source word importance.

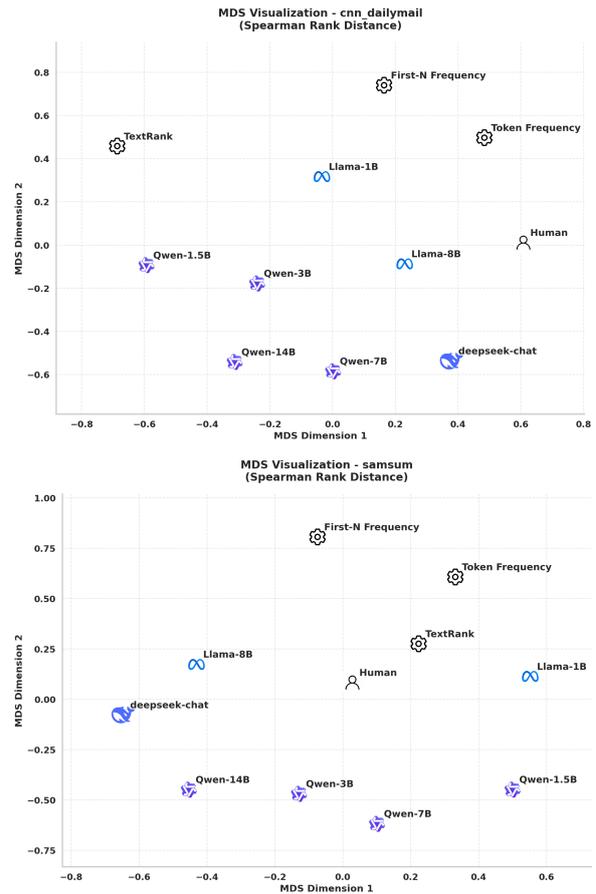


Figure 2: Pairwise model similarity based on Spearman rank correlation distance for importance distributions, visualized via two-dimensional Multidimensional Scaling (MDS). Results are shown for the CNN/DailyMail (top) and SAMSum (bottom) datasets.

$d = 1 - \rho$, where ρ is Spearman’s correlation coefficient. For each document, we compute the Spearman distance between models’ word importance rankings, then average across all common documents to obtain a single pairwise distance. These distances are visualized via Multidimensional Scaling (MDS) in Figure 2, which projects the dissimilarity values into two-dimensional Euclidean space where inter-model distances reflect their behavioral similarity. Our analysis reveals two key patterns in summarization behavior:

- Distinct LLM Clustering:** LLMs exhibit broadly similar behavior, forming a cluster distinct from pre-LLM baselines. *Human Frequency* occupies a central position between LLM and pre-LLM summarization models.
- Architectural Bias:** Models tend to cluster by family (e.g., Llama vs. Qwen). This family-based clustering appears more pronounced on the CNN/DailyMail dataset than on SAMSum.

The distributional similarity measured by NDCG@10 is provided in Figure 8 in the Appendix.

4.3 Qualitative Insights

The importance distribution $I_M(D)$ across all models and datasets (Figure 6, Appendix) shows that most words receive low scores, with approximately 50% of word-importance pairs assigned a score of 0.1. Only a small fraction of words are marked as highly important (≥ 0.8).

In both datasets, **named entities and core concepts** (e.g., main actors, central events) consistently receive the highest importance scores. **Function words and general connectors** are uniformly assigned low importance. Supporting details typically fall within a medium importance range. A more detailed analysis of score distributions is provided in Appendix C.3.

5 Attention–Importance Alignment

We now examine internal computations to address the second goal: *How closely do a model’s multi-head attention patterns align with the output-based importance distribution $I_M(D)$?*

For each sample (300 from CNN/DailyMail, 819 from SAMSum), we extract the raw attention weights from the model’s multi-head attention layers when processing the input prompt (before generating the summary). Then, we compute a token-level *attention received* score by summing, for each token, the incoming attention from all other positions in the sequence. For words composed of multiple subword tokens, we averaged the attention scores of their constituent tokens. Both the resulting attention distribution and the empirical importance distribution $I_M(D)$ are normalized.

5.1 Top- k Ranking Alignment with Importance Scores

We evaluate the top- k ranking consistency of attention heads by visualizing the average NDCG@10 per head across the CNN/DailyMail, SAMSum, and DECODA datasets (Figs. 11–13 in the Appendix). To examine performance trends by depth, we also plot the average NDCG@10 across layers for each dataset (Fig. 17 in the Appendix). Our analysis reveals several key observations:

Dataset Effect. The alignment strength varies significantly by dataset. SAMSum exhibits strong alignment, whereas CNN/DailyMail shows weak

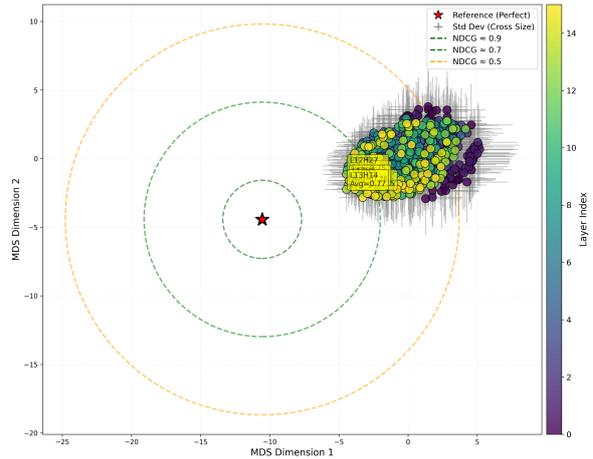


Figure 3: Multi-Dimensional Scaling (MDS) projection of attention heads for Llama-3.2-1B-Instruct on SAMSum. Points represent heads positioned by their per-sample NDCG@10 profiles, reflecting similarity in top- k ranking quality with importance distribution. Point color indicates layer depth (lighter = deeper). The red star marks the ideal point of perfect ranking alignment (NDCG@10 = 1.0); dashed contours indicate similarity thresholds (e.g., NDCG@10 ≥ 0.9).

alignment. This disparity likely stems from dataset properties such as abstractiveness, document length, and summary style, which influence how well attention mechanisms encode the underlying importance ranking. The shorter, more extractive nature of SAMSum dialogues appears to align more naturally with attention patterns.

Layer-Wise Specialization. Performance peaks in different layers across models. While early layers show some alignment, the most proficient ranking frequently occurs in middle to late layers, indicating depth-wise functional specialization.

5.2 MDS Visualization

We employ Multidimensional Scaling (MDS) to visualize the clustering of attention heads based on their performance profiles. Figure 3 presents this visualization for the Llama-3.2-1B-Instruct model on the SAMSum dataset, embedding each head in a two-dimensional space based on the similarity of their per-sample NDCG@10 score vectors.

The MDS projection reveals that 98.6% of attention heads (505/512) achieve NDCG@10 ≥ 0.5 on SAMSum. Heads cluster progressively closer to the reference distribution (NDCG@10 = 1.0) in deeper layers, with layer 13 containing some of the highest performers (e.g., L13H14, NDCG@10 = 0.769 \pm 0.114). The tight clustering of high-performing heads, together with small

per-head standard deviations (indicated by minimal cross-markers), demonstrates that the model’s top- k ranking behavior is both effective and consistent across inputs. Corresponding results using other metrics (e.g., Spearman correlation) for the best-performing head are provided in Appendix E.3.

Overall, these results suggest that inspection of attention distribution alone is already predictive of the distributional behavior of the models when generating summaries.

6 Probing Experiments

We train surrogate models (probes) to predict the empirical *importance distribution* $I_M(D)$ from model-internal hidden states. Each token in the input document is assigned an importance score corresponding to the average importance score of the information unit (word) to which it belongs, as defined by $I_M(D)$. As input to the probe, we use the hidden state vectors associated with individual tokens at a given transformer layer. For words that occur multiple times within a document, we aggregate their representations by averaging their hidden state vectors, yielding a single representation per word. All probes are implemented as a one-hidden layer multi-layer perceptrons (MLPs). They were trained for 20 epochs using the Kullback–Leibler (KL) divergence loss on a 60:20:20 train/validation/test split, with early stopping (patience=3). Additional details regarding training and dataset statistics are provided in Appendix C.4.

6.1 Three Probe Training Scenarios

We define three probing scenarios to investigate which components of model M encode information predictive of *importance distribution* $I_M(D)$. They are defined as follows:

1. **Layer-wise Probing.** A separate probe is trained on the hidden states of *each individual layer* (including the embedding layer) to predict I_i *per token*. This isolates the predictiveness of individual hidden layers.
2. **All-layers Probing.** A single probe is trained on features from the *concatenated hidden states of a token across all layers*, predicting I_i *per token*. This tests if combining cross-layer information improves token-level prediction.
3. **Article-level Probing.** A separate probe is trained *per layer*, but for all tokens in parallel with a single KL loss over *all tokens*.

6.1.1 Probing Baselines

TextRank baseline. We compute TextRank scores as an unsupervised, content-based reference for probe evaluation. For each document, TextRank scores (Section 4.1) are aligned with the empirical *importance distribution* $I_M(D)$ using the union vocabulary and normalized to sum to one. They are then compared against the target empirical importance distribution as a measure of the agreement between a parameter-free salience signal, providing a baseline for probe performance.

Randomized-weights baseline. To further control for false discoveries, we evaluate probing performance on models with randomized weights. This *dead salmon* baseline tests whether predictive performance arises from pretrained representations rather than probe capacity alone on random projections of the embeddings. Following Méroux et al. (2025), we reinitialize all model parameters using the architecture’s native initialization scheme, extract hidden states using the same procedure as for pretrained models, and train probes to predict $I_M(D)$.

Performance is reported on the test set using Spearman’s rank correlation and NDCG@10. Further results for all model–dataset pairs are reported in Appendix Tables 5 and 6.

6.2 Experimental Results for Article-level Probing

This section presents the results for **Scenario 3: Article-level Probing**, where the probing loss is aggregated across all tokens in a document. This evaluates the probe’s ability to reconstruct the relative importance distribution for an entire document context. Results for **Scenario 1 (Layer-wise Probing)** and **Scenario 2 (All-layers Probing)** are provided in Appendix F.3 and Appendix F.4, respectively.

6.2.1 Quantitative Results

This section reports NDCG@10 results on CNN/DailyMail and SAMSum datasets (Figure 4); complementary Spearman results are provided in Appendix F.2. NDCG@10 values are generally high, indicating that probes capture locally consistent word-importance rankings. Performance rises in early layers and peaks in middle-to-late layers across all datasets, suggesting that the most discriminative information for top- k ranking is encoded at intermediate depths. Unlike Spearman correlation, which remains stable on SAMSum and DECODA,

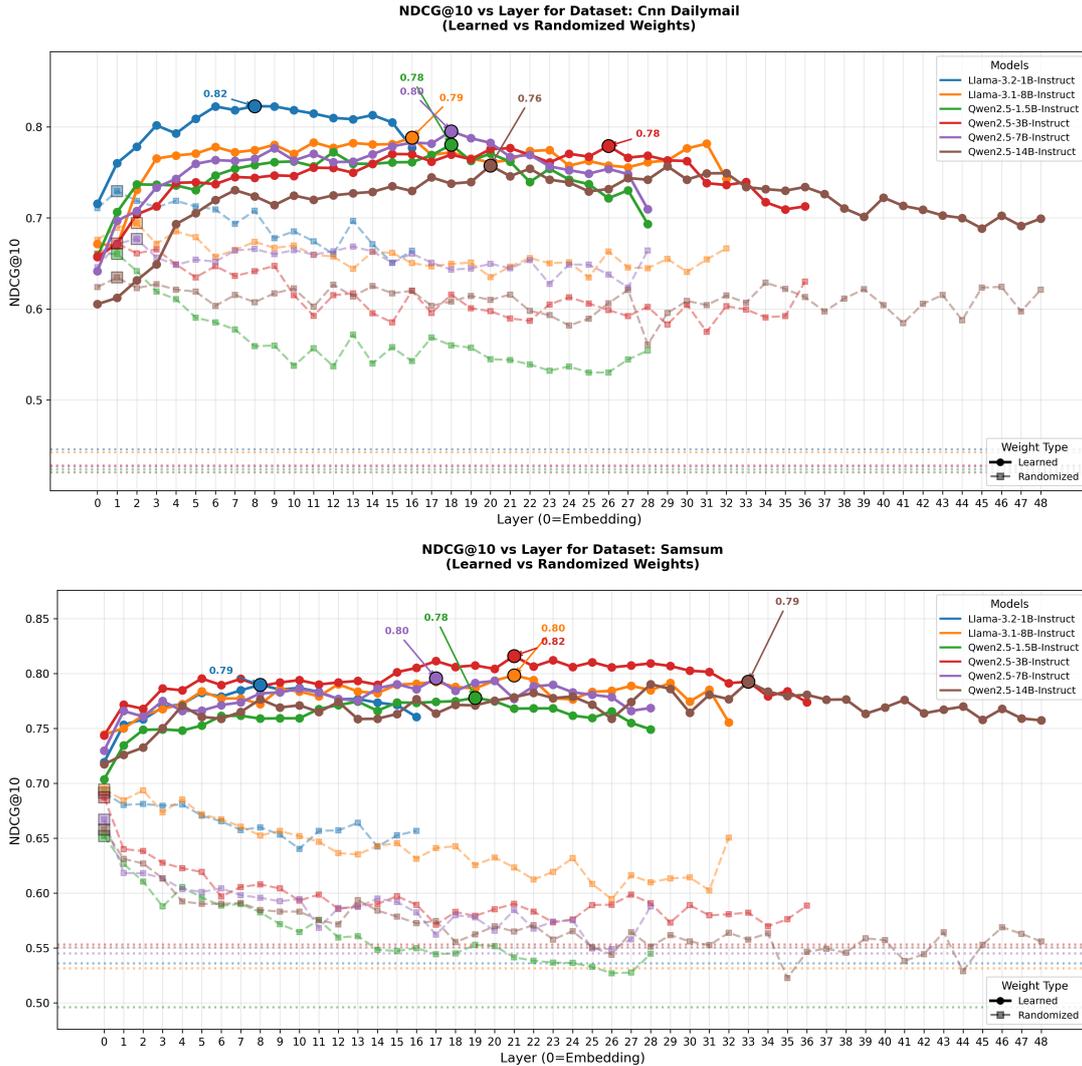


Figure 4: Article-level probing NDCG@10 across layers for CNN/DailyMail (top) and SAMSUM (bottom). Round dots show learned model performance; square dots show the *Randomized Weights Baseline*. The best-performing layers are annotated. Horizontal dashed lines show the TextRank baseline for each model.

NDCG@10 exhibits pronounced layer-wise variation across all three datasets.

Probes trained on learned model weights substantially outperform those trained on randomized-weight models, achieving improvements of 0.1–0.2 NDCG@10 points, and consistently exceed the TextRank baseline. In contrast, randomized-weight probes show degraded performance and high inter-layer variance. This demonstrates that the LLMs do indeed construct representations that are predictive of the importance distributions.

6.2.2 Example Analysis

We select Qwen2.5-3B-Instruct for a case study because of its balanced performance: it achieves the highest NDCG@10 score on SAMSUM while ranking in the middle on CNN/DailyMail (Sec-

tion 6.2.1). Figure 5 visualizes article-level probe predictions via heatmaps, illustrating the layer-wise encoding of importance for representative samples from both datasets. The heatmaps (Fig. 5) plot predicted importance across transformer layers (vertical axis) versus the first 50 input tokens (horizontal axis), with intensity indicating encoding strength. Token-level analysis reveals consistent high importance across layers for: (1) CNN/DailyMail: named entities (e.g., “Palestinian”, “Authority”, “the international criminal court”); (2) SAMSUM: person names (e.g., “Hannah”, “Betty”, “Amanda”, “Larry”) and core dialogue tokens (“number”, “it”, “her”). These patterns reflect dataset characteristics: *importance* is broadly distributed across content in news articles, but focuses more heavily on participant references in dialogues.

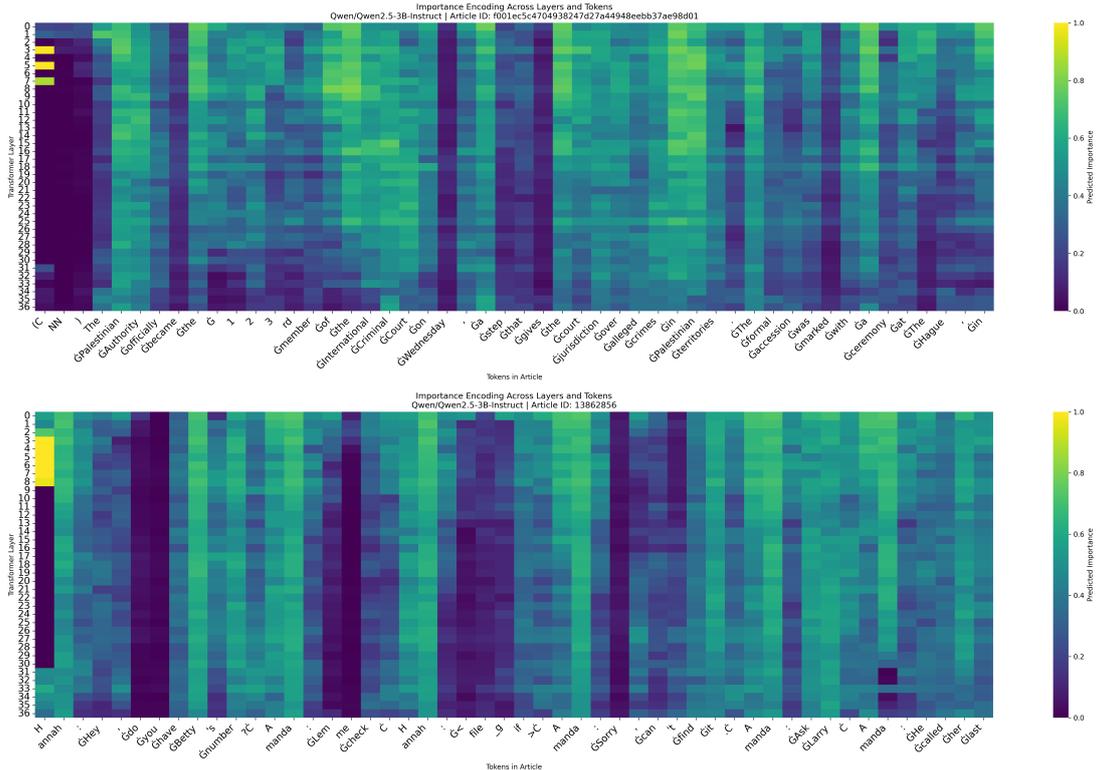


Figure 5: Heatmap visualization of layer-wise importance predictions for Qwen2.5-3B-Instruct, showing probe outputs across all layers and the first 50 tokens for representative samples (top: CNN/DailyMail, bottom: SAMSum).

7 Discussion and Conclusion

This study provides a computational and behavioral investigation into the latent *information importance* that guides content selection in LLMs. We derived an empirical importance distribution $I_M(D)$ by analyzing information persistence across length-controlled summaries. Our behavioral analysis shows that LLMs display consistent importance patterns distinct from pre-LLM baselines, with models clustering more strongly by architecture than by size.

To predict $I_M(D)$, we employed attention analysis and hidden-state probing. Both approaches identify the middle-to-late layers as critical for encoding importance, with specific attention heads serving as effective predictors. These results show that inspection of the internal computation of the models while processing a document can be highly predictive of the distributional importance of information units. However, the strength of the predictions is dataset-dependent (strongest on SAMSum, weakest on CNN/DailyMail) and model-dependent (with Llama models outperforming Qwen). Our central findings are twofold. First, the encoding of importance is **insensitive to model scale**, showing no consistent improvement with increased parameters.

Second, it is **highly task-specific**, being far more predictable in conversational dialogues (SAMSum) than in long-form news (CNN/DailyMail). This suggests that a model’s capacity for importance ranking is less a function of its size and more a product of its architectural design, training data and the inherent nature of the source text.

Our work provides initial insights into LLMs’ summarization priorities and their internal representations, paving the way toward interpreting and controlling information selection. A key gap remains: while models maintain consistent internal hierarchies, we lack direct methods to access or control them. Future work should bridge this gap through causal manipulation experiments for output control.

Limitations

This study has several limitations. First, our data processing pipeline involves numerous choices (e.g., handling of function words) that could be refined to ensure high importance scores correspond more directly to substantive content. Second, despite using length-constraint prompts, not all generated summaries strictly adhered to the specified token counts, introducing variance. A third limi-

tation stems from the data characteristic detailed in Appendix C.4: the imperfect alignment between annotated summary words and source tokens results in a limited set of word–hidden-state pairs for probe training. Future work could improve this mapping by incorporating semantic similarity, moving beyond exact lexical matching. In general, other information units more semantically meaningful can be inspected in the future.

Acknowledgments

This work was partially supported by the “Intelligent Systems for Data, Knowledge, and Humans” axis of the Grenoble Computer Science Laboratory (LIG). This work was conducted within the French research unit UMR 5217 and was supported by CNRS (grant ANR-22-CPJ2-0036-01) and by MIAI@Grenoble-Alpes (grant ANR-19-P3IA-0003 and ANR-23-IACL-0006).

References

- AI@Meta. 2024. [Llama models card](#).
- Jie Bao, Prithwish Basu, Mike Dean, Craig Partridge, Ananthram Swami, Will Leland, and James A. Hendler. 2011. [Towards a theory of semantic communication](#). *2011 IEEE Network Science Workshop*, pages 110–117.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2023. [Discovering latent knowledge in language models without supervision](#). In *The Eleventh International Conference on Learning Representations*.
- Nick Cammarata, Shan Carter, Gabriel Goh, Chris Olah, Michael Petrov, Ludwig Schubert, Chelsea Voss, Ben Egan, and Swee Kiat Lim. 2020. [Thread: Circuits](#). *Distill*. <https://distill.pub/2020/circuits>.
- Rudolf Carnap and Yehoshua Bar-Hillel. 1952. [An outline of a theory of semantic information](#).
- Shan Carter, Zan Armstrong, Ludwig Schubert, Ian Johnson, and Chris Olah. 2019. [Activation atlas](#). *Distill*.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Zhichen Dong, Zhanhui Zhou, Zhixuan Liu, Chao Yang, and Chaochao Lu. 2025. [Emergent response planning in LLMs](#). In *Forty-second International Conference on Machine Learning*.
- Harold P Edmundson. 1969. New methods in automatic extracting. *Journal of the ACM (JACM)*, 16(2):264–285.
- Benoit Favre, Evgeny Stepanov, Jérémy Trione, Frédéric Béchet, and Giuseppe Riccardi. 2015. [Call centre conversation summarization: A pilot task at multiling 2015](#). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 232–236, Prague, Czech Republic. Association for Computational Linguistics.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. [Dissecting recall of factual associations in auto-regressive language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12216–12235, Singapore. Association for Computational Linguistics.
- Mohamed Imed Eddine Ghebriout, Gaël Guibon, Ivan Lerner, and Emmanuel Vincent. 2025. [QUARTZ: QA-based unsupervised abstractive refinement for task-oriented dialogue summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 14689–14706, Suzhou, China. Association for Computational Linguistics.
- Daniel Gillick, Benoit Favre, and Dilek Hakkani-Tür. 2008. The icsi summarization system at tac 2008. In *Tac*.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. [SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China. Association for Computational Linguistics.
- Tanya Goyal and Greg Durrett. 2021. [Annotating and modeling fine-grained factuality in summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1449–1462, Online. Association for Computational Linguistics.
- Tianjie Ju, Zhenyu Shao, Bowen Wang, Yujia Chen, Zhuosheng Zhang, Hao Fei, Mong-Li Lee, Wynne Hsu, Sufeng Duan, and Gongshen Liu. 2025. [Probing then editing response personality of large language models](#). In *Second Conference on Language Modeling*.
- Chris Kedzie, Kathleen McKeown, and Hal Daumé III. 2018. [Content selection in deep learning models of summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1818–1828, Brussels, Belgium. Association for Computational Linguistics.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. [Attention is not only a weight: Analyzing transformers with vector norms](#). In *Proceedings of the 2020 Conference on Empirical*

- Methods in Natural Language Processing (EMNLP)*, pages 7057–7075, Online. Association for Computational Linguistics.
- Abhinav Kumar, Chenhao Tan, and Amit Sharma. 2022. Probing classifiers are unreliable for concept removal and detection. *Advances in Neural Information Processing Systems*, 35:17994–18008.
- Chen Li, Xian Qian, and Yang Liu. 2013. [Using supervised bigram-based ILP for extractive summarization](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1004–1013, Sofia, Bulgaria. Association for Computational Linguistics.
- Raymond Li, Wen Xiao, Linzi Xing, Lanjun Wang, Gabriel Murray, and Giuseppe Carenini. 2022. [Human guided exploitation of interpretable attention patterns in summarization and topic segmentation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10189–10204, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Wei Li, Lei He, and Hai Zhuge. 2016. [Abstractive news summarization based on event semantic link network](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 236–246, Osaka, Japan. The COLING 2016 Organizing Committee.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Hans Peter Luhn. 1958. The automatic creation of literature abstracts. *IBM Journal of research and development*, 2(2):159–165.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Ryan McDonald. 2007. A study of global inference algorithms in multi-document summarization. In *Proceedings of the 29th European Conference on IR Research, ECIR’07*, page 557–564, Berlin, Heidelberg. Springer-Verlag.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in GPT. *Advances in Neural Information Processing Systems*, 36. ArXiv:2202.05262.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Giovanni Monea, Maxime Peyrard, Martin Josifoski, Vishrav Chaudhary, Jason Eisner, Emre Kiciman, Hamid Palangi, Barun Patra, and Robert West. 2024. [A glitch in the matrix? locating and detecting language model grounding with fakepedia](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6828–6844, Bangkok, Thailand. Association for Computational Linguistics.
- Maxime Méloux, Giada Dirupo, François Portet, and Maxime Peyrard. 2025. [The dead salmon of ai interpretability](#). *Preprint*, arXiv:2512.18792.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Ani Nenkova and Kathleen McKeown. 2012. *A Survey of Text Summarization Techniques*, pages 43–76. Springer US, Boston, MA.
- Ani Nenkova, Rebecca Passonneau, and Kathleen McKeown. 2007. [The pyramid method: Incorporating human content selection variation in summarization evaluation](#). *ACM Trans. Speech Lang. Process.*, 4(2):4–es.
- Ani Nenkova and Lucy Vanderwende. 2005. The impact of frequency on summarization. *Microsoft Research, Redmond, Washington, Tech. Rep. MSR-TR-2005*, 101.
- Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan Belinkov. 2025. [LLMs know more than they show: On the intrinsic representation of LLM hallucinations](#). In *The Thirteenth International Conference on Learning Representations*.
- Maxime Peyrard. 2019. [A simple theoretical model of importance for summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1059–1073, Florence, Italy. Association for Computational Linguistics.
- Maxime Peyrard and Judith Eckle-Kohler. 2016. [Optimizing an approximation of ROUGE - a problem-reduction approach to extractive multi-document summarization](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1825–1836, Berlin, Germany. Association for Computational Linguistics.
- Sanjana Ramprasad, Elisa Ferracane, and Zachary Lipton. 2024. [Analyzing LLM behavior in dialogue summarization: Unveiling circumstantial hallucination trends](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12549–12561, Bangkok, Thailand. Association for Computational Linguistics.

- Abhilasha Ravichander, Yonatan Belinkov, and Eduard Hovy. 2021. [Probing the probing paradigm: Does probing accuracy entail task relevance?](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3363–3377, Online. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Julius Steen and Katja Markert. 2024. [Bias in news summarization: Measures, pitfalls and corpora](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5962–5983, Bangkok, Thailand. Association for Computational Linguistics.
- Xu Sun, Lionel Delphin-Poulat, Christèle Tarnec, and Anastasia Shimorina. 2025. [PoSum-bench: Benchmarking position bias in LLM-based conversational summarization](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 7996–8020, Suzhou, China. Association for Computational Linguistics.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- Damien Teney, Maxime Peyrard, and Ehsan Abbasnejad. 2022. [Predicting is not understanding: Recognizing and addressing underspecification in machine learning](#). In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*, page 458–476, Berlin, Heidelberg. Springer-Verlag.
- Jan Trienes, Jörg Schlötterer, Junyi Jessy Li, and Christin Seifert. 2025. [Behavioral analysis of information salience in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 23428–23454, Vienna, Austria. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- David Wan, Jesse Vig, Mohit Bansal, and Shafiq Joty. 2025. [On positional bias of faithfulness for long-form summarization](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8791–8810, Albuquerque, New Mexico. Association for Computational Linguistics.
- Linzi Xing, Wen Xiao, and Giuseppe Carenini. 2021. [Demoting the lead bias in news summarization via alternating adversarial learning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 948–954, Online. Association for Computational Linguistics.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2024. [Benchmarking large language models for news summarization](#). *Transactions of the Association for Computational Linguistics*, 12:39–57.
- Yixin Zhong. 2017. [A theory of semantic information](#). *Proceedings*, 1(3).
- Yongxin Zhou, François Portet, and Fabien Ringeval. 2022. [Effectiveness of French language models on abstractive dialogue summarization task](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3571–3581, Marseille, France. European Language Resources Association.
- Yongxin Zhou, Fabien Ringeval, and François Portet. 2024. [PSentScore: Evaluating sentiment polarity in dialogue summarization](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13290–13302, Torino, Italia. ELRA and ICCL.
- Yongxin Zhou, Fabien Ringeval, and François Portet. 2025. [Can GPT models follow human summarization guidelines? a study for targeted communication goals](#). In *Proceedings of the 18th International Natural Language Generation Conference*, pages 249–273, Hanoi, Vietnam. Association for Computational Linguistics.
- Markus Zopf, Eneldo Loza Mencía, and Johannes Fürnkranz. 2016. [Beyond centrality and structural features: Learning information importance for text summarization](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 84–94, Berlin, Germany. Association for Computational Linguistics.

A Multilingual Study

To assess the cross-lingual generalizability of our findings, we extend our evaluation to a French dialogue summarization dataset, complementing the English datasets described in Section 3.1.

We employ the **DECODA** corpus (Favre et al., 2015), a French call center dialogue collection created for the Multiling 2015 CCCS abstractive sum-

marization task. The corpus contains 1,000 unannotated dialogues and a test set of 100 dialogues with human-written synopses. We use the preprocessed samples from Zhou et al. (2022), which retain speaker identifiers while removing extraneous labels.

All reported results, including behavioral analysis, attention alignment, and article-level probing (Scenario 3), are based on the 100-dialogue test set. These results are presented in the following sections alongside supplementary results for the CNN/DailyMail and SAMSum datasets.

B Metrics for Word Importance Evaluation

Evaluating word importance attribution is non-trivial due to the lack of a standard metric for comparing *importance distribution* $I_M(D)$, unlike for generated text (e.g., ROUGE, BERTScore). A core challenge is the absence of a clear theoretical baseline, which complicates both metric selection and the interpretation of results.

B.1 Candidate Metrics

To rigorously assess the alignment between model-generated *importance distribution* $I_M(D)$ and a human gold standard (derived from reference summaries), we evaluate a comprehensive suite of metrics spanning four categories:

- **Correlation:** Spearman’s ρ and Kendall’s τ , measuring global rank alignment.
- **Ranking:** Normalized Discounted Cumulative Gain (NDCG@ k), Precision@ k , and Recall@ k , focusing on the retrieval of top-important words.
- **Distributional Divergence:** Kullback-Leibler (KL) Divergence, Jensen-Shannon Divergence (JSD), and Rényi Divergence, measuring the distance between probability distributions.
- **Set Overlap:** Jaccard Similarity, measuring the intersection of selected important word sets.

B.2 Meta-Evaluation Methodology

Given the diversity of available metrics, we perform a meta-evaluation to identify the most robust indicators of quality. We evaluate each metric based on two key criteria:

1. **Discrimination Power (D):** The standard deviation of the metric scores across different models. A higher D indicates the metric effectively distinguishes between models of varying quality.
2. **Sensitivity (S):** The ratio of the observed range of scores to the theoretical range of the metric. An ideal sensitivity (approaching 1.0) implies the metric utilizes its full scale and is not saturated.

We compute a **Composite Score** for each metric, defined as the average of its normalized *Discrimination* and *Sensitivity* scores, to rank their overall empirical utility.

B.3 Empirical Results

Our meta-evaluation on the CNN/DailyMail (Long-form news, 3,000 samples) and SAMSum (Short-form dialogue, 819 samples) datasets yielded the following insights:

CNN/DailyMail Results The divergence-based metrics demonstrated high discrimination power. **Rényi Divergence** ($\alpha = 2.0$) achieved the highest composite score (0.864), followed by **KL Divergence** (0.748). Among bounded metrics, **Spearman’s ρ** ranked third (0.746) with high sensitivity (0.86), and **NDCG@10** ranked fifth (0.704) with perfect sensitivity (1.00).

SAMSum Results For the dialogue dataset, overlap-based metrics performed exceptionally well. **Jaccard@15** achieved the highest composite score (0.965), likely due to the shorter, more keyword-centric nature of the dialogues. However, **Rényi Divergence** ($\alpha = 2.0$) remained robust (3rd, 0.816), and **NDCG@10** (4th, 0.754) and **Spearman’s ρ** (8th, 0.706) continued to show strong performance with perfect sensitivity (1.00).

B.4 Choice of Metrics

Based on the empirical results and theoretical properties, we select **Spearman’s Rank Correlation** (ρ) and **NDCG@10** as our primary universal metrics for the following reasons:

1. **Complementarity:** Spearman captures the *global* alignment of the entire importance distribution, while NDCG@10 focuses specifically on the *local* quality of the most important words (top-10), which are most critical for summarization.

2. **Robustness:** Both metrics demonstrated high sensitivity (> 0.85) across both datasets, ensuring they remain meaningful regardless of the data domain (news vs. dialogue).

Although Rényi Divergence ($\alpha = 2.0$) demonstrated high discriminative power, we reserve it as a secondary diagnostic tool due to its unbounded range, prioritizing the bounded and more interpretable Spearman ρ and NDCG@10 for our primary analysis.

C Details on the Experimental Setup

C.1 Prompts for Data Generation

The prompts used for generating length-variant summaries across datasets are presented in Table 1, where $N \in \{10, 20, 30, \dots, 100\}$.

Prompt for Data Generation
<p>CNN/DailyMail: “Summarize the following text in exactly {N} words: {text}”</p> <p>SAMSum: “Summarize the following dialogue in exactly {N} words: {dialogue}”</p> <p>DECODA: “Résumez le dialogue suivant en exactement {N} mots: {dialogue}”</p>

Table 1: Prompts used for generating length-variant summaries across datasets.

C.2 Model Specification

In Table 2, we list the models and their corresponding links.

C.3 Analysis of Importance Score Distributions

Figure 6 presents the distributions of importance scores across all models for the three datasets. The distributions reveal that the majority of words receive low importance scores, with approximately 50% of all word-importance pairs assigned a score of 0.1 across all three datasets. Only a small fraction ($\sim 6\text{-}8\%$) of words receive scores ≥ 0.8 .

CNN/DailyMail For the CNN/DailyMail dataset, the derived importance scores show a hierarchy: the highest scores (≥ 0.9) are assigned to **central content words**, including named entities, main events. Medium scores correspond to **supporting information**, while the lowest scores (≤ 0.2) are

given to function words, stop words, and peripheral background details.

All evaluated models demonstrate strong agreement on the most important words, consistently identifying main subjects, verbs, and unique entities. Substantial overlap occurs in the top 10–20% of important words across models, with greater variation appearing in the middle and low importance ranges. For example:

- High Importance (≥ 0.9): “international”, “authority”, “court”, “criminal”, “palestinian”, “state”, “stray”, “washington”, “dog”, “in” (central entities and concepts).
- Medium Importance (0.4–0.6): “peace”, “marking”, “into”, “investigate”, “over”, “jurisdiction”, “on”, “as”, “field” (supporting details and context).
- Low Importance (≤ 0.2): “joins”, “formal”, “with”, “giving”, “face”, “but”, “if”, “be”, “its”, “whether” (function words, connectors, and background terms).

SAMSum The SAMSum dataset shows similar patterns, with **named entities and main actors** receiving the highest importance scores, while function words and background terms are consistently assigned low importance. For example:

- High Importance (≥ 0.9): “amanda”, “find”, “eric”, “rob”, “lenny”, “bob” (named entities and key actors).
- Medium Importance (0.4–0.6): “for”, “if”, “ask”, “phone”, “routine” (supporting details and context).
- Low Importance (≤ 0.2): “tries”, “texting”, “exchange”, “so”, “suggesting”, “do”, “of”, “goodbye” (function words, connectors and supplementary verbs).

DECODA The DECODA dataset (French customer service dialogues) exhibits a distinct pattern where the model prioritizes both **domain-specific content** and essential **grammatical markers**. Unlike the English datasets, determiners and prepositions frequently receive high importance scores, likely due to their critical role in French syntax and coreference resolution.

- High Importance (≥ 0.9): “numéro”, “métro”, “bus”, “client”, “remboursement” (domain entities), along with “le”, “un”, “pour”, “de” (determiners/prepositions).
- Medium Importance (0.4–0.6): “l’agent”, “rappeler”, “après”, “donc”, “faire” (procedural verbs and connectors).

Models	Links
Llama-3.2-1B-Instruct (AI@Meta, 2024)	https://huggingface.co/meta-llama/Llama-3.2-1B-Instruct
Llama-3.1-8B-Instruct (AI@Meta, 2024)	https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct
Qwen2.5-1.5B-Instruct (Team, 2024)	https://huggingface.co/Qwen/Qwen2.5-1.5B-Instruct
Qwen2.5-3B-Instruct (Team, 2024)	https://huggingface.co/Qwen/Qwen2.5-3B-Instruct
Qwen2.5-7B-Instruct (Team, 2024)	https://huggingface.co/Qwen/Qwen2.5-7B-Instruct
Qwen2.5-14B-Instruct (Team, 2024)	https://huggingface.co/Qwen/Qwen2.5-14B-Instruct
deepseek-chat [†] (DeepSeek-AI et al., 2025)	https://api-docs.deepseek.com/

[†]Experiments done in September 2025, the model points to non-thinking mode of DeepSeek-V3.1-Terminus.

Table 2: Models used in the experiments and their corresponding links.

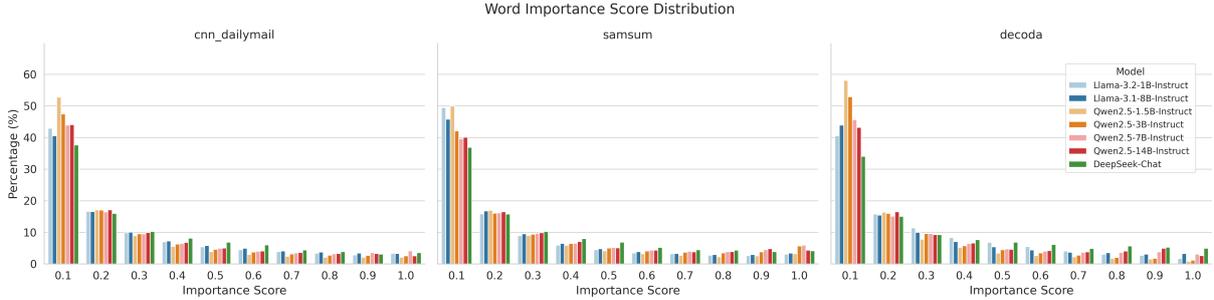


Figure 6: Distribution of word importance scores across models and datasets.

- Low Importance (≤ 0.2): “conversation”, “échange”, “précise”, “bonjour”, “est”, “on”, “il” (meta-dialogue descriptors, greetings, and generic auxiliary verbs).

These patterns demonstrate that while models consistently prioritize named entities and main actions, the treatment of function words varies by language, with French models retaining more grammatical structure in the high-importance tier.

C.4 Hidden States Extraction

Model & Dataset	Zero (%)	Annot. (%)	Words/Art.
Llama-3.2-1B-Instruct			
CNN-DailyMail	53.5	46.5	298.1
SAMSum	22.3	77.7	82.2
Llama-3.1-8B-Instruct			
CNN-DailyMail	59.8	40.2	295.6
SAMSum	36.8	63.2	81.8
Qwen2.5-1.5B-Instruct			
CNN-DailyMail	58.5	41.5	318.0
SAMSum	45.8	54.2	81.4
Qwen2.5-3B-Instruct			
CNN-DailyMail	60.7	39.3	308.2
SAMSum	44.2	55.8	77.6
Qwen2.5-7B-Instruct			
CNN-DailyMail	64.3	35.7	301.6
SAMSum	52.2	47.8	73.3
Qwen2.5-14B-Instruct			
CNN-DailyMail	65.3	34.7	301.2
SAMSum	49.4	50.6	75.9

Table 3: Statistics on hidden states extraction in different models on both datasets.

Table 3 presents detailed statistics on the ratio

of zero-score words to annotated words. The annotated words are identified based on their frequency in the generated summaries and their presence in the corresponding article or dialogue.

The analysis encompasses 300 articles from the CNN/DailyMail dataset, with approximately 88K-95K total words per model, as well as 819 dialogues from the SAMSum dataset, containing approximately 60K-67K total words per model.

The statistics reveal a distinct pattern in hidden-state extraction between datasets. For the conversational SAMSum dataset, models identified a majority of source words as annotated (47.8%–77.7%), with a minority as zero-score words (22.3%–52.2%). The opposite trend holds for the news-based CNN/DailyMail dataset, where zero-score words constitute the majority (53.5%–65.3%) and annotated words are the minority (34.7%–46.5%).

This suggests that summarization operates differently by genre. In long-form news, the task is inherently more abstractive, requiring the distillation of many source words into a concise summary, which results in a high proportion of zero-score words. In shorter dialogues, summaries are more extractive, preserving a higher density of the original words. Consequently, the informative features captured by hidden states are distributed differently, with conversational data exhibiting a higher concentration of annotated, salient tokens.

D Behavioral Analysis

D.1 Model Similarity

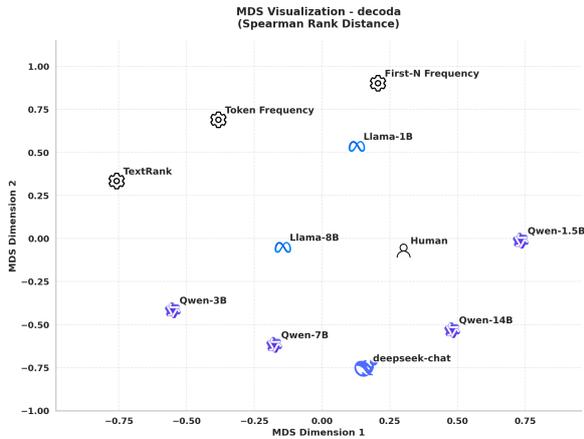


Figure 7: Pairwise model similarity based on Spearman rank correlation distance for importance distributions, visualized via two-dimensional Multidimensional Scaling (MDS). Results are shown for the DECODA dataset.

We extend the model similarity analysis from Section 4.2 by including results from the French DECODA dataset (Figure 7) and by evaluating similarity using NDCG@10 (Figure 8). The latter figure visualizes pairwise model distances based on NDCG@10 across all three datasets.

The key observations from Section 4.2 are reinforced and sharpened using NDCG@10: the behavioral distinction between LLMs and pre-LLM baselines (excluding *First-N-Frequency*) is more pronounced; clustering by model family becomes more visually distinct; and the *Human (Frequency)* consistently occupies an intermediate position between the pre-LLM baselines and the LLM cluster. These consistent patterns across two different metrics and three datasets strengthen the conclusion that LLMs share a common, family-influenced approach to attributing *importance distribution* $I_M(D)$ that differs from classical methods.

D.2 Quantifying Positional Bias

Since human tend to favor early information (e.g., historically strong baselines used the first sentences), we quantify this positional bias using a weighted positional importance score. For each model, the bias for a document D is defined as:

$$\text{Bias}(D) = \frac{1}{S_I} \sum_{t=1}^L I(t) \cdot p(t), \quad \text{where } p(t) = \frac{t}{L}.$$

Here, $I(t)$ is the importance score of token t , L is the document length in tokens, $S_I = \sum_t I(t)$

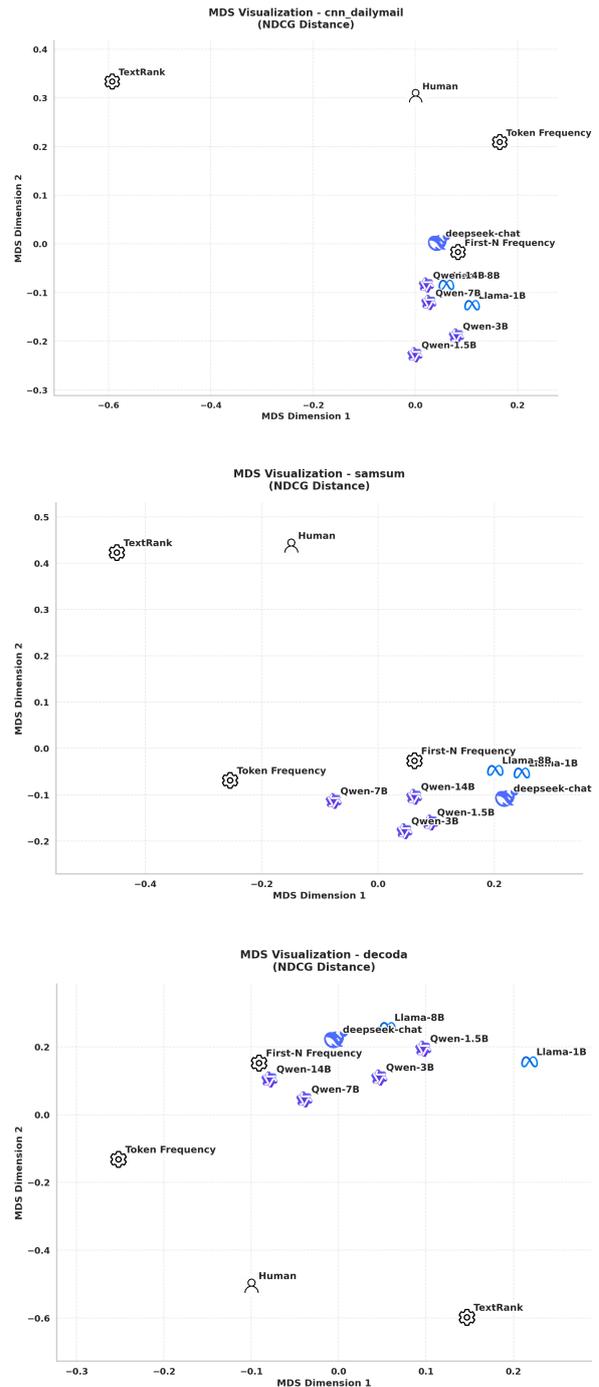


Figure 8: Model similarity in importance distribution rankings across the CNN/DailyMail, SAMSUM, and DECODA datasets, visualized using pairwise NDCG@10 distances.

is the total importance, and $p(t)$ is the normalized position (ranging from 0 at the start to 1 at the end). For words with multiple occurrences, we use the average positional index. A *lower* $\text{Bias}(D)$ value indicates a stronger preference for earlier tokens.

Figure 9 shows the bias distributions for each model on the three datasets. Consistent patterns

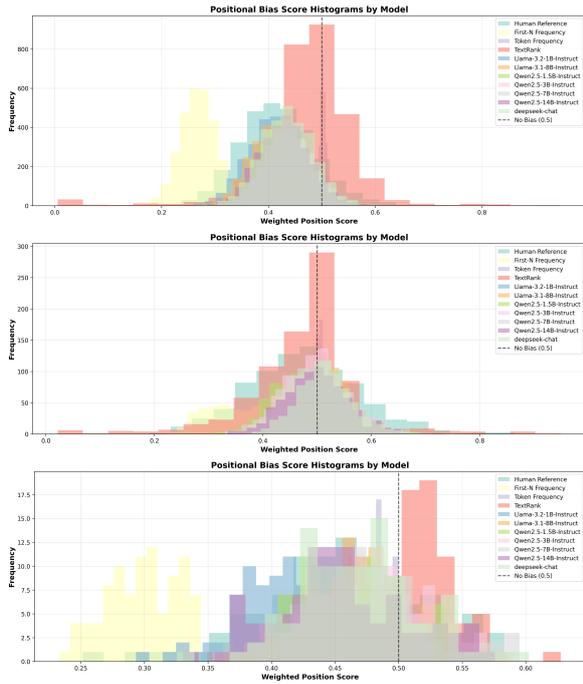


Figure 9: Distribution of positional bias scores for all models across the CNN/DailyMail (top), SAMSUM (middle), and DECODA (bottom) datasets. Lower scores indicate a stronger bias towards earlier document positions.

emerged: the baseline *First-N Frequency* exhibited the strongest early bias (lowest scores: 0.228 on CNN/DailyMail, 0.079 on SAMSUM, 0.196 on DECODA), while other baselines (*TextRank*, *Token Frequency*) showed more balanced or late-leaning distributions.

LLMs demonstrated moderate early biases. Qwen2.5 and Llama variants typically favored early positions more than the DeepSeek model. The weakest early bias (i.e., strongest late bias) varied by dataset: *Token Frequency* on CNN/DailyMail, *DeepSeek-Chat* on SAMSUM, and *TextRank* on DECODA. These findings suggest that while baseline methods rely heavily on document structure, LLMs exhibit more nuanced positional preferences that likely reflect learned content understanding rather than simple heuristics.

A comparison across datasets shows that CNN/DailyMail exhibits the strongest early-position bias, aligning with the standard inverted pyramid structure of news articles. In contrast, the two dialogue datasets display more moderate positional biases: SAMSUM shows the most balanced distribution, whereas DECODA displays a comparatively stronger, yet still moderate, early bias. This

difference likely reflects their distinct discourse structures: casual multi-turn conversations in SAMSUM versus the goal-oriented, problem-solving sequences characteristic of customer service dialogues in DECODA.

D.3 Model Entropy Comparison

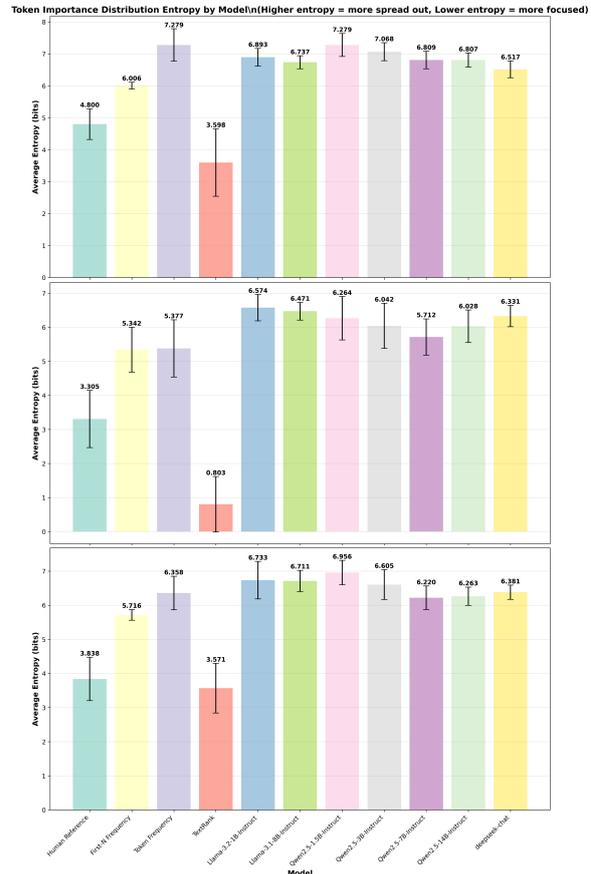


Figure 10: Entropy comparison by models on different datasets (top: CNN/DailyMail, middle: SAMSUM, bottom: DECODA).

We quantify the concentration of model-dependent *importance distribution* $I_M(D)$ using Shannon entropy. Figure 10 shows these entropy distributions across models, where lower entropy indicates importance is concentrated on fewer words, and higher entropy reflects more distributed attention.

We compute entropy as $H(I) = -\sum_i p_i \log_2(p_i)$ using the normalized importance scores $p_i = I_i / \sum_j I_j$. We observe systematically higher entropy for models on the CNN/DailyMail dataset compared to SAMSUM and DECODA, consistent with the greater length and informational density of news articles relative to conversational dialogues.

The results reveal distinct patterns across datasets. On CNN/DailyMail, DeepSeek-Chat produces the most focused distributions (lowest entropy: 6.52 ± 0.26 bits), while Qwen2.5-1.5B-Instruct yields the most dispersed (highest entropy: 7.28 ± 0.36 bits). This pattern differs on SAMSum, where Qwen2.5-7B-Instruct shows the most concentrated assignments (5.71 ± 0.53 bits) and Llama-3.2-1B-Instruct the most distributed (6.57 ± 0.39 bits). On DECODA, models maintain higher overall entropy, with Qwen2.5-1.5B-Instruct reaching 6.96 ± 0.36 bits and DeepSeek-Chat maintaining relatively selective attention at 6.38 ± 0.21 bits. This stability across languages suggests robust, language-invariant mechanisms for encoding importance.

E Attention-Importance Alignment

E.1 Heatmap Visualization

NDCG@10 To compare the top- k ranking consistency of attention heads across datasets, we visualize the average NDCG@10 per head in Figs. 11, 12, and 13 for CNN/DailyMail, SAMSum, and DECODA, respectively.

Spearman rank correlation We further measure attention-importance alignment using Spearman rank correlation. Figures 14, 15 and 16 visualize the average Spearman ρ per attention head for different LLMs on the three datasets, respectively, with the top three heads per model indicated by red ranked markers.

Our multi-head attention analysis shows that attention is not a monolithic proxy for importance. Early layers show negligible correlation, while specific heads in middle-to-late layers emerge as strong predictors (Spearman $\rho \approx 0.4$). Architecturally, Qwen2.5 models align importance in final layers, whereas Llama models align in middle layers. This suggests that *information importance* is captured by specialized attention heads, not by attention as a whole.

E.2 NDCG@10 by Layer

In Figure 17, we present the evolution of average NDCG@10 scores across model layers for the CNN/DailyMail (a), SAMSum (b), and DECODA (c) datasets. Each point represents the mean NDCG@10 across all attention heads within a layer, illustrating how the alignment between attention and the model-dependent *importance distribution* $I_M(D)$ varies with network depth.

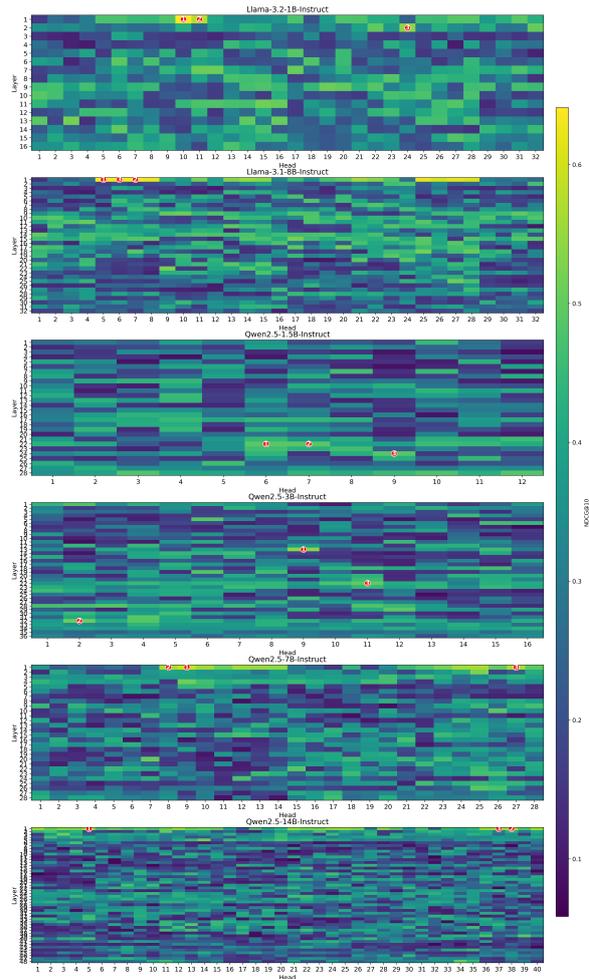


Figure 11: Average NDCG@10 per attention head on CNN/DailyMail. Red markers with ranks 1-3 indicate the top three heads per model.

E.3 Metrics Results for the Best Head

The MDS projection in Section 5.2 identified the best-performing attention head for Llama-3.2-1B-Instruct on SAMSum as Layer 13, Head 14 (NDCG@10 = 0.77). The results of other evaluation metrics for this head are reported in Table 4.

This head (Layer 13, Head 14) exhibits the strongest alignment with *importance distribution* $I_M(D)$ among all 512 attention heads. The NDCG@10 score of 0.769 indicates strong consistency in the ranking of the top-10 words, while the Spearman correlation of 0.352 reflects a moderate overall rank correlation. In addition, high precision scores (Precision@5 = 79.4%, Precision@10 = 68.4%) suggest that when this head assigns high attention to words, they are likely to be important. Conversely, the low recall scores (e.g., Recall@10 = 36.3%) indicate that it captures only a subset of

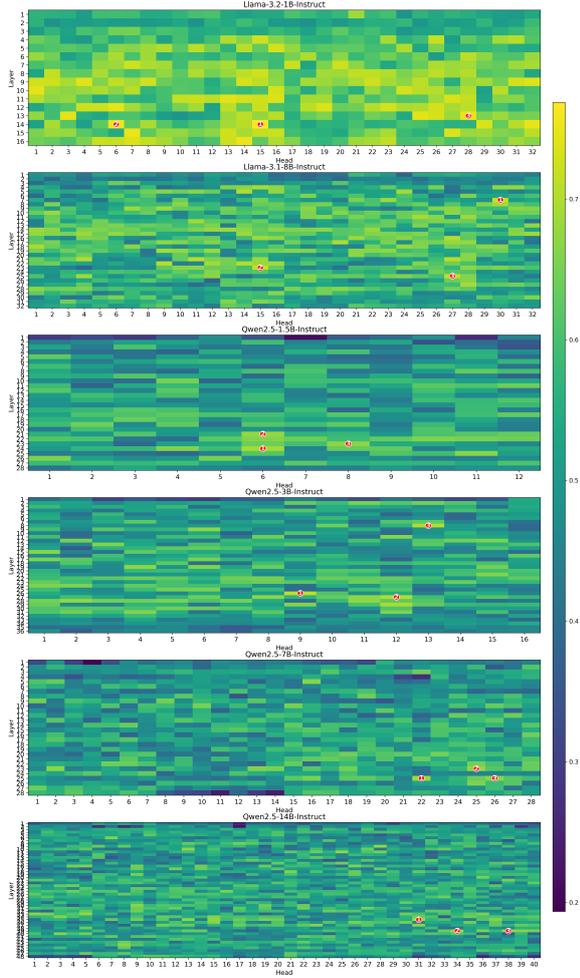


Figure 12: Average NDCG@10 per attention head on SAMSsum. Red markers with ranks 1–3 indicate the top three heads per model.

all important words, signifying a specialized rather than comprehensive focus.

The standard deviations across samples quantify the consistency of these metrics. For this head, the Spearman correlation shows moderate variability (std = 0.161). The NDCG scores exhibit high consistency (std \approx 0.11), indicating reliable top- k ranking quality. The Rényi divergence metrics show greater variability (std \approx 0.55), suggesting more sample-dependent distribution similarity. Overall, the low standard deviations, particularly for NDCG, confirm that the head’s performance is stable.

F Probing Results for the Three Scenarios

F.1 Baselines for Comparison

TextRank baseline values for each model and dataset are reported in Tables 5 (Spearman correlation) and 6 (NDCG@10). These baselines quan-

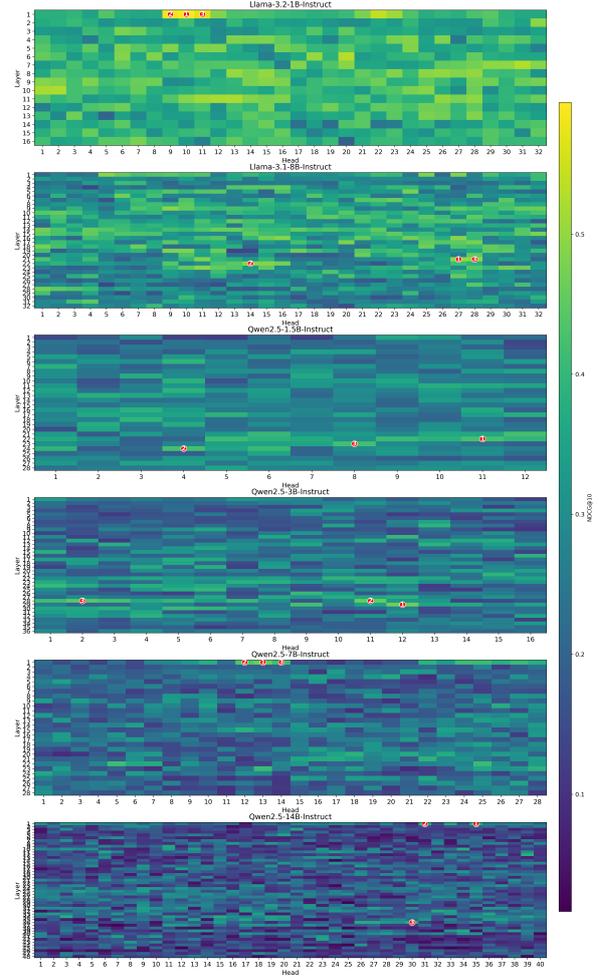


Figure 13: Average NDCG@10 per attention head on DECODA. Red markers with ranks 1–3 indicate the top three heads per model.

tify the alignment between unsupervised TextRank scores and the model-dependent *importance distribution* $I_M(D)$, establishing a content-aware performance lower bound for evaluating the probes.

F.2 Scenario 3: Article-level Probing

This section provides complementary probing analyses. In addition to the main results, we report NDCG@10 on the French summarization dataset DECODA (Figure 19) to examine cross-lingual robustness. On DECODA, the performance gap between probes trained on learned and randomized weights is narrower than on the English datasets, with peak performance (\approx 0.81 for Llama-3.1-8B-Instruct) occurring in middle-to-late layers. This suggests that importance ranking on DECODA may rely less on learned patterns, possibly due to more directly accessible base representations or distinct structural properties of the task.

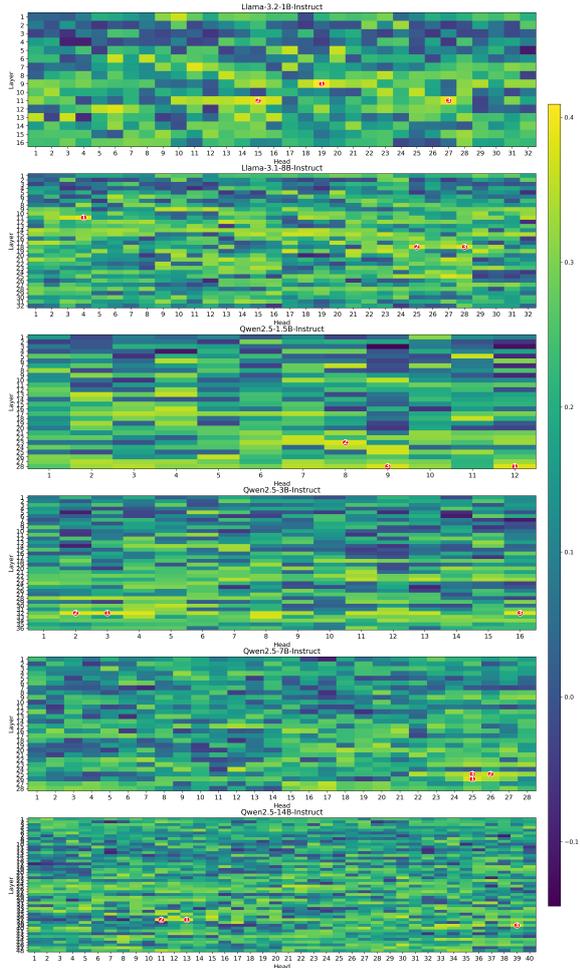


Figure 14: Average Spearman ρ per attention head on CNN/DailyMail. Red markers with ranks 1–3 indicate the top three heads per model.

Complementary Spearman correlation results are shown in Figure 18.

Overall effectiveness of information importance probing. The figure presents layer-wise Spearman correlations between probe-predicted information importance rankings and the empirical *importance distribution* $I_M(D)$. Across all evaluated models and datasets, Spearman values range from moderate to high, indicating that hidden states encode substantial information relevant to word-level importance in summarization. Notably, even the worst-performing layers consistently outperform the TextRank baseline (see Tables 5), which is omitted from the figure for clarity.

Layer-wise localization of information importance. A consistent pattern emerges across models: peak probing performance is typically attained in the middle-to-late layers rather than at the embedding or final layers. For instance, on

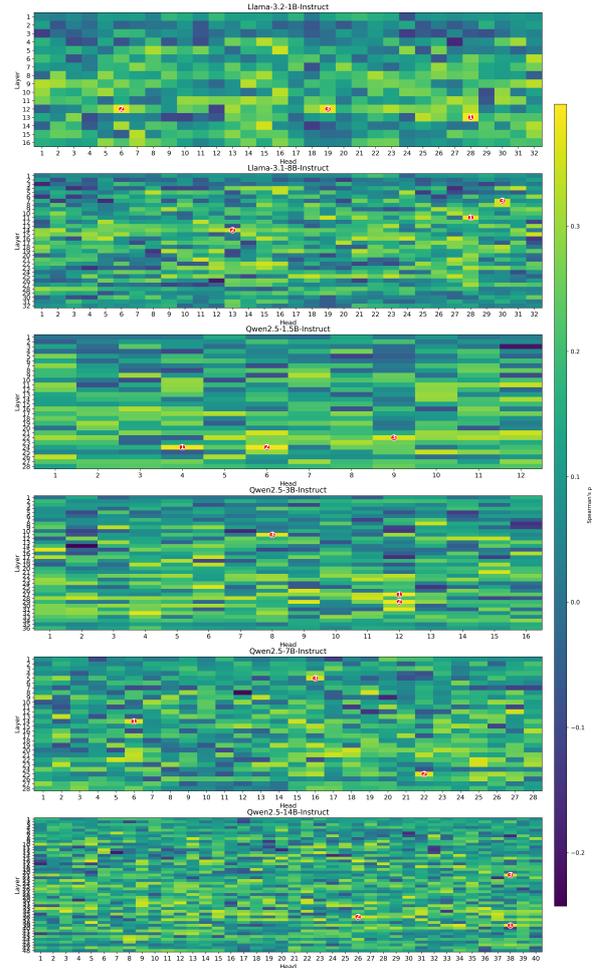


Figure 15: Average Spearman ρ per attention head on SAMSum. Red markers labeled with rank 1–3 indicate the top three heads per model.

CNN/DailyMail, LLaMA-3.2-1B-Instruct reaches its maximum correlation at layer 9 (out of 16 layers), while larger models such as Qwen-2.5-14B-Instruct peak substantially deeper (e.g., around layer 32 out of 48 layers). A similar trend is observed on SAMSum, where most models achieve their highest correlations in the upper-middle layers. This layer-wise behavior suggests that *information importance* is most explicitly represented after initial lexical processing but before the final layers, which are more specialized toward generation and output distribution modeling.

Cross-task differences in layer sensitivity. While the overall trend is consistent across datasets, the degree of layer-wise variation differs markedly between tasks. On CNN/DailyMail, which involves long, information-dense articles, probing performance varies substantially across layers, indicating a more pronounced redistribution of importance-

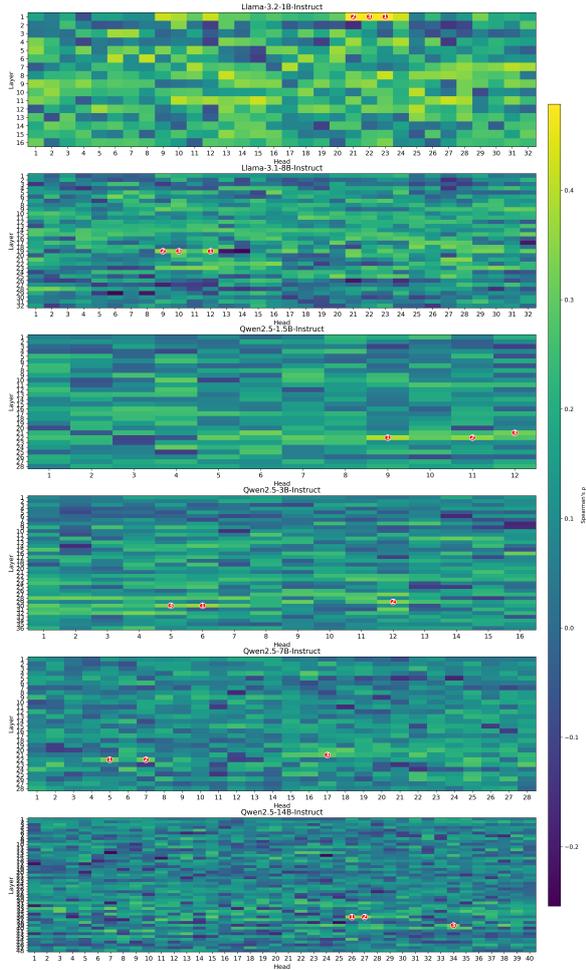
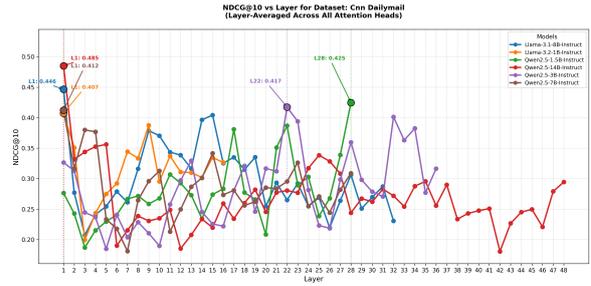


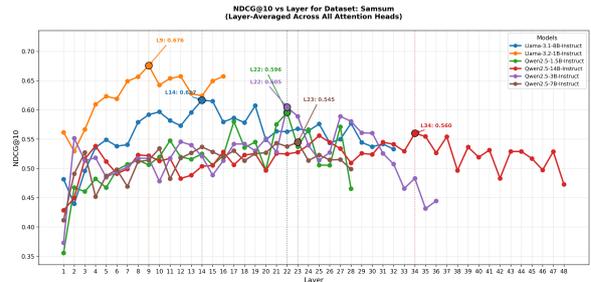
Figure 16: Average Spearman ρ per attention head on DECODA. Red markers with ranks 1–3 indicate the top three heads per model.

related information throughout the network depth. In contrast, on SAMSUm, a dialogue summarization task with shorter inputs and more localized salient content, Spearman correlations remain relatively stable across a broad range of layers. This reduced layer sensitivity suggests that importance cues in conversational data may be encoded more uniformly across representations, possibly due to lower structural complexity and weaker long-range dependencies.

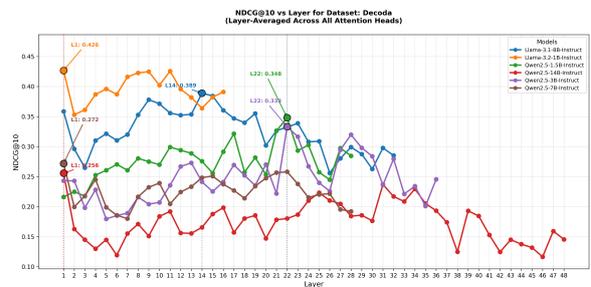
Spearman correlations on the DECODA dataset remain moderate to high across models, indicating that *information importance* is reliably captured by hidden representations in this French summarization setting. Similar to the trends observed on CNN/DailyMail and SAMSUm, peak correlations are typically achieved in the middle-to-late layers rather than at the embedding or final layers. These results provide additional evidence that the



(a) CNN/DailyMail



(b) SAMSUm



(c) DECODA

Figure 17: Evolution of average NDCG@10 scores across transformer layers for the CNN/DailyMail (a), SAMSUm (b), and DECODA (c) datasets. Each point represents the mean NDCG@10 across all attention heads within a layer.

proposed probing approach captures importance-related representations in a manner that is robust across languages.

F.3 Scenario 1 (Layer-wise Probing)

For resource efficiency and because preliminary results showed consistent trends between layer-wise and all-layers probing (with middle-to-late layers generally performing best), we conducted layer-wise probing only for Llama-3.2-1B-Instruct and Qwen2.5-1.5B-Instruct on the SAMSUm. The corresponding Spearman ρ and NDCG@10 results are presented in Figures 20 and 21.

The results reveal trends consistent with article-level probing, with best performance achieved in middle-to-late layers. Specifically, for Llama-3.2-1B-Instruct, Layer 10 performs best, while for Qwen2.5-1.5B-Instruct, Layer 25 is optimal, as

Metric	Value (Mean \pm Std)
<i>Ranking Correlation</i>	
Spearman Correlation	0.352 \pm 0.161
<i>Distribution Similarity</i>	
KL Score	-4.038
Power-Weighted KL Score	2.378
Rényi Divergence ($\alpha = 1.5$)	3.080 \pm 0.496
Rényi Divergence ($\alpha = 2.0$)	3.496 \pm 0.543
Rényi Divergence ($\alpha = 3.0$)	4.033 \pm 0.626
<i>Ranking Quality</i>	
NDCG@5	0.822 \pm 0.115
NDCG@10	0.769 \pm 0.114
<i>Information Retrieval</i>	
Precision@1	0.986
Recall@1	0.060
Precision@3	0.868
Recall@3	0.150
Precision@5	0.794
Recall@5	0.222
Precision@10	0.684
Recall@10	0.363

Table 4: Complete evaluation metrics for the best-performing attention head (Layer 13, Head 14) of Llama-3.2-1B-Instruct on the SAMSum dataset. Reported as mean \pm standard deviation where applicable; values are rounded to three decimal places.

Model	CNN/DailyMail	SAMSum	DECODA
<i>Llama Family</i>			
Llama-3.2-1B-Instruct	-0.133 \pm 0.179	0.130 \pm 0.112	-0.057 \pm 0.146
Llama-3.1-8B-Instruct	-0.162 \pm 0.194	0.102 \pm 0.124	-0.165 \pm 0.131
<i>Qwen Family</i>			
Qwen2.5-1.5B-Instruct	-0.145 \pm 0.175	0.092 \pm 0.137	-0.261 \pm 0.094
Qwen2.5-3B-Instruct	-0.158 \pm 0.184	0.106 \pm 0.139	-0.290 \pm 0.128
Qwen2.5-7B-Instruct	-0.188 \pm 0.182	0.099 \pm 0.158	-0.352 \pm 0.120
Qwen2.5-14B-Instruct	-0.194 \pm 0.192	0.095 \pm 0.146	-0.351 \pm 0.107

Table 5: TextRank Spearman correlation baselines (mean \pm std) per model and dataset.

Model	CNN/DailyMail	SAMSum	DECODA
<i>Llama Family</i>			
Llama-3.2-1B-Instruct	0.446 \pm 0.156	0.536 \pm 0.132	0.311 \pm 0.099
Llama-3.1-8B-Instruct	0.443 \pm 0.162	0.531 \pm 0.118	0.236 \pm 0.091
<i>Qwen Family</i>			
Qwen2.5-1.5B-Instruct	0.424 \pm 0.166	0.496 \pm 0.121	0.171 \pm 0.074
Qwen2.5-3B-Instruct	0.428 \pm 0.163	0.553 \pm 0.121	0.151 \pm 0.073
Qwen2.5-7B-Instruct	0.427 \pm 0.163	0.545 \pm 0.126	0.155 \pm 0.069
Qwen2.5-14B-Instruct	0.420 \pm 0.161	0.550 \pm 0.126	0.155 \pm 0.076

Table 6: TextRank NDCG@10 baselines (mean \pm std) per model and dataset.

indicated by both Spearman ρ and NDCG@10 visualizations in Figures 20 and 21. Compared to

all-layers probing, the best single layers yield similar but slightly lower Spearman correlations on SAMSum: peak ρ values are 0.56 for Llama-3.2-1B-Instruct and 0.70 for Qwen2.5-1.5B-Instruct, compared to 0.580 and 0.715, respectively, for all-layer probes. In terms of NDCG@10, the best single layers perform comparably: Llama-3.2-1B-Instruct achieves 0.81 versus 0.797 with all-layer probing, while Qwen2.5-1.5B-Instruct reaches 0.78 versus 0.795.

F.4 Scenario 2 (All-layers Probing)

In addition to the article-level and layer-wise probing experiments, we also conducted an all-layer probing by concatenating the hidden states from all layers of the model. This approach aggregates information across the network hierarchy, providing a richer representation than any single layer. As shown in Table 7, the all-layer concatenation achieves performance comparable to the best-performing individual layer: it is slightly better in some settings, while the optimal single layer performs better in others, with overall differences remaining small.

F.5 Computational Experiments

All experiments were conducted using PyTorch 2.6.0 with CUDA 12.4 on NVIDIA GPUs (RTX A6000, RTX 8000, H100). To ensure reproducibility, we fixed the random seed to 42 and enabled PyTorch’s deterministic mode.

G Cross Dataset Transfer Capability

We investigate the cross-dataset transfer capability of article-level probes by comparing their importance predictions on a shared CNN/DailyMail sample from Qwen2.5-3B-Instruct. As shown in Figure 22, the heatmap visualizes layer-wise predictions across the first 50 tokens, comparing probes trained separately on CNN/DailyMail (top) and SAMSum (bottom).

The bottom heatmap (from a probe trained on SAMSum) captures most importance patterns identified in the top heatmap (from a probe trained on CNN/DailyMail), demonstrating partial cross-dataset generalization. Both probes consistently identify key entities such as “Palestinian”, “Authority”, “criminal”, and “court” as highly important.

However, the probe trained on SAMSum exhibits a tendency to assign higher overall importance scores compared to the CNN/DailyMail-

Model	Dataset	Spearman	NDCG@10	Input Dim	Dataset Size	Epochs	Early Stop	Time (min)
Llama-3.2-1B-Instruct	CNN/DailyMail	0.664	0.835	34,816	89,421	5	Yes	191.6
	SAMSum	0.580	0.797	34,816	67,333	5	Yes	48.7
Llama-3.1-8B-Instruct	CNN/DailyMail	0.668	0.828	135,168	88,674	5	Yes	255.8
	SAMSum	0.733	0.821	135,168	66,962	6	Yes	67.4
Qwen2.5-1.5B-Instruct	CNN/DailyMail	0.636	0.798	44,544	95,396	5	Yes	232.3
	SAMSum	0.715	0.795	44,544	66,641	6	Yes	62.0
Qwen2.5-3B-Instruct	CNN/DailyMail	0.643	0.796	75,776	92,475	5	Yes	232.9
	SAMSum	0.736	0.819	75,776	63,574	5	Yes	54.8
Qwen2.5-7B-Instruct	CNN/DailyMail	0.648	0.818	103,936	90,476	5	Yes	236.9
	SAMSum	0.739	0.827	103,936	60,045	5	Yes	54.3
Qwen2.5-14B-Instruct	CNN/DailyMail	0.639	0.798	250,880	90,351	5	Yes	153.2
	SAMSum	0.735	0.822	250,880	62,200	5	Yes	73.9

Table 7: All-layers probing results and training details across models and datasets. Performance metrics (Spearman ρ , NDCG@10) were computed for each sample and then averaged.

trained probe. This suggests that the internal representation of importance may be calibrated differently when learned from the more extractive, dialogue-based data of SAMSum versus the more abstractive, long-form data of CNN/DailyMail.

H Licensing of Artifacts

All datasets and models used in this work are released under the following licenses.

- **CNN/DailyMail dataset** (v1.0.0): Apache-2.0 License
- **SAMSum dataset**⁵: CC BY-NC-ND 4.0 (non-commercial)
- **DECODA dataset**: Available from the official website⁶ upon acceptance of usage terms
- **Qwen models**: Apache 2.0 License
- **Llama models**: Llama Community License
- **DeepSeek-V3**⁷: MIT License

I AI Assistants In Research Or Writing

This research was conducted with the assistance of AI tools for text refinement and with coding support from Copilot.

⁵<https://huggingface.co/datasets/knkarthick/samsum>

⁶<https://pageperso.lis-lab.fr/benoit.favre/ccs/>

⁷<https://huggingface.co/deepseek-ai/DeepSeek-V3.1-Terminus>

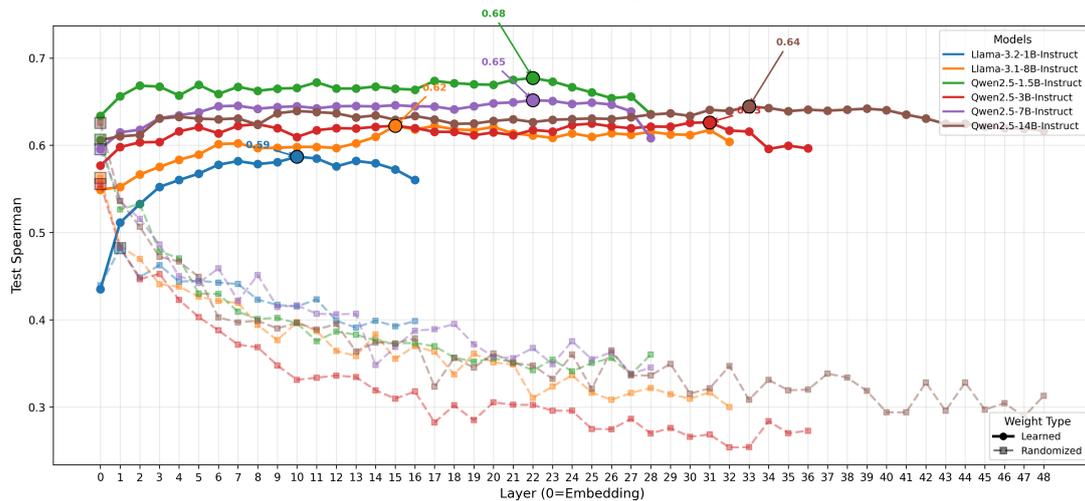
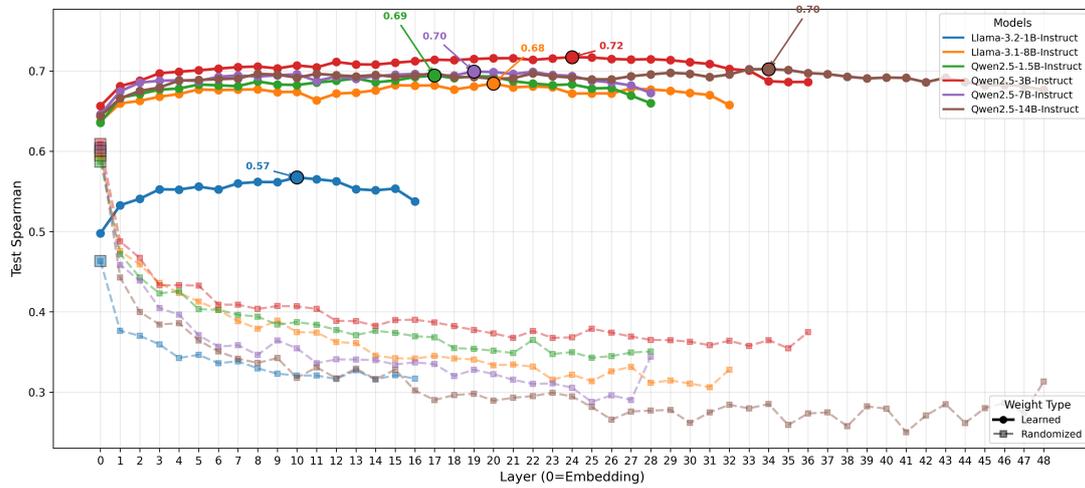
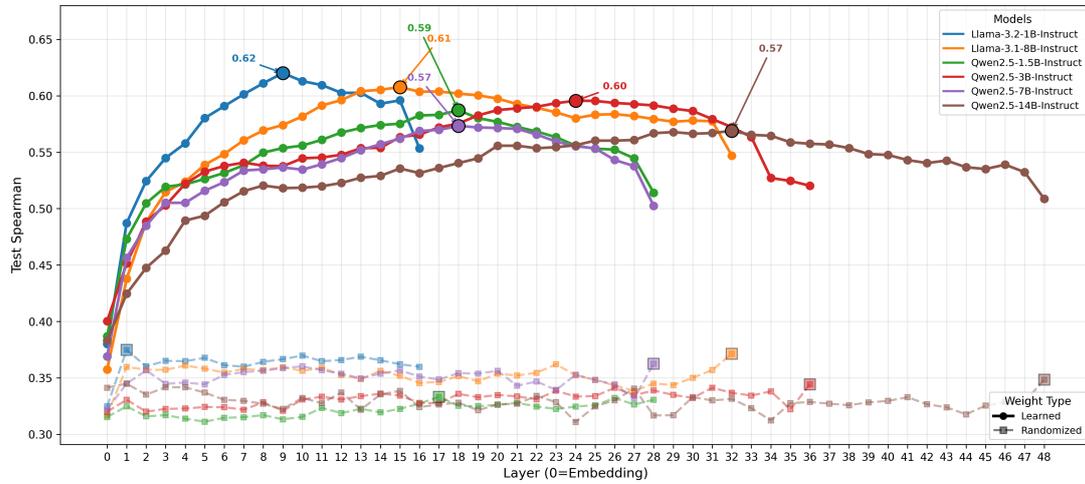


Figure 18: Spearman rank correlation (ρ) for article-level probing across layers on the CNN/DailyMail (top), SAMSum (middle), and DECODA (bottom) datasets. Performance is indicated by round markers, while square markers denote the *Randomized Weights Baseline*. The best-performing layer for each dataset is highlighted and annotated. TextRank baseline results (approximately zero) are omitted for clarity.

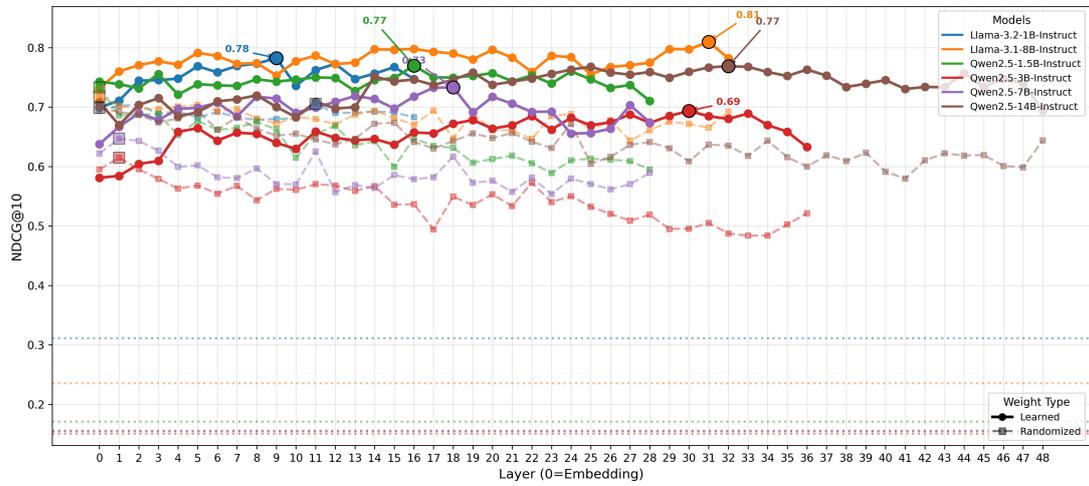


Figure 19: Article-level probing NDCG@10 across layers for DECODA. Round dots show learned model performance; square dots show the *Randomized Weights Baseline*. The best-performing layers are annotated. Horizontal dashed lines show the TextRank baseline for each model.

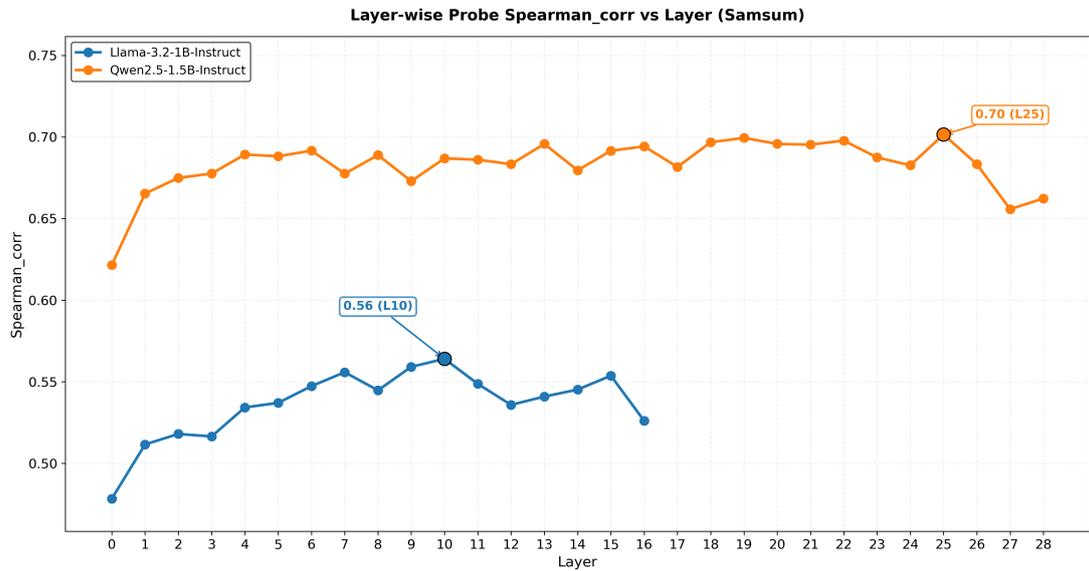


Figure 20: Layer-wise probing performance (Spearman ρ) across layers for Llama-3.2-1B-Instruct and Qwen2.5-1.5B-Instruct on the SAMSsum dataset. Best-performing layers are annotated with their index and score. Values represent per-sample averages.

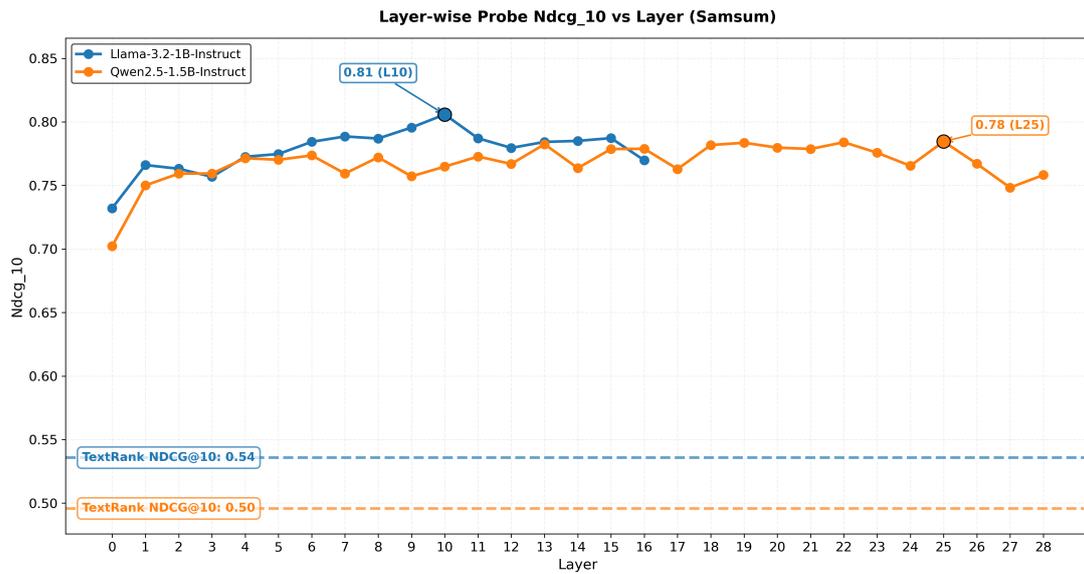


Figure 21: Layer-wise probing performance (NDCG@10) across layers for Llama-3.2-1B-Instruct and Qwen2.5-1.5B-Instruct on the SAMSum dataset. Best-performing layers are annotated with their index and score. Values represent per-sample averages.

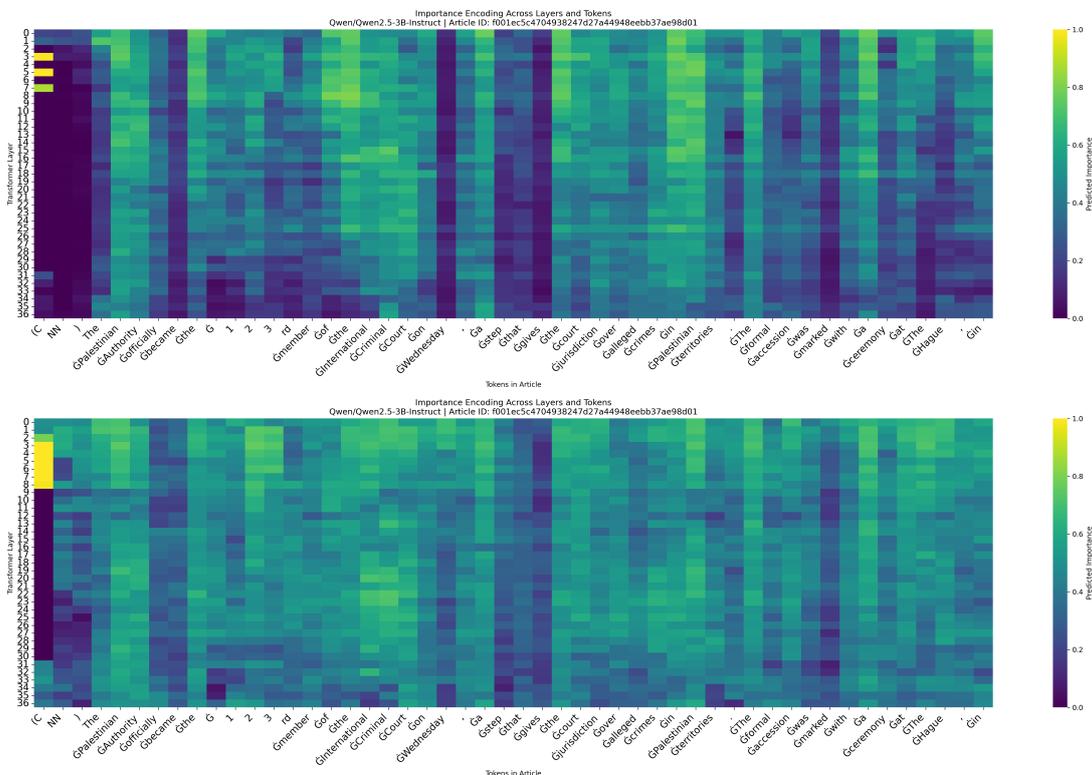


Figure 22: Layer-wise importance predictions for Qwen2.5-3B-Instruct on a CNN/DailyMail sample, comparing article-level probes trained on CNN/DailyMail (top) versus SAMSum (bottom). The heatmap displays probe outputs across all layers and the first 50 tokens.