

Detecting Overflow in Compressed Token Representations for Retrieval-Augmented Generation

Julia Belikova^{1,2}, Danila Rozhevskii¹, Dennis Svirin^{1,4},
Konstantin Polev², and Alexander Panchenko^{1,3}

¹Skoltech, ²Sber AI Lab, ³AIRI

⁴Institute for Information Transmission Problems of the Russian Academy of Sciences

Correspondence: {julia.belikova, a.panchenko}@skol.tech

Abstract

Efficient long-context processing remains a crucial challenge for contemporary large language models (LLMs), especially in resource-constrained environments. Soft compression architectures promise to extend effective context length by replacing long token sequences with smaller sets of learned *compressed tokens*. Yet, the limits of compressibility – and when compression begins to erase task-relevant content – remain underexplored. In this paper, we define *token overflow* as a regime in which compressed representations no longer contain sufficient information to answer a given query, and propose a methodology to characterize and detect it. In the xRAG soft-compression setting, we find that query-agnostic saturation statistics reliably separate compressed from uncompressed token representations, providing a practical tool for identifying compressed tokens but showing limited overflow detection capability. Lightweight probing classifiers over both query and context xRAG representations detect overflow with 0.72 AUC-ROC on average on HotpotQA, SQuADv2, and TriviaQA datasets, demonstrating that incorporating query information improves detection performance. These results advance from query-independent diagnostics to query-aware detectors, enabling low-cost pre-LLM gating to mitigate compression-induced errors.

1 Introduction

Large language models (LLMs) remain computationally constrained when processing long contexts, even as architectural advances and extended context windows become widely available (Vaswani et al., 2017; Liu et al., 2024). In retrieval-augmented generation (RAG), this limitation is particularly acute: retrieved evidence must be aggressively compressed or truncated, creating a tension between efficiency and faithfulness (Lewis et al., 2020; Aushev et al., 2025). Soft compression architectures address this by mapping long contexts into

dense vectors that can be directly consumed by the model, dramatically reducing token count while preserving global semantics (Liao et al., 2025).

However, the same mechanism that enables extreme compression also introduces a critical failure mode. As more information is packed into a fixed-dimensional *compressed token*, its representation can enter token overflow: it no longer carries sufficient task-relevant signal for the query and effectively behaves like noise, silently degrading downstream performance. Recent work on trainable tokens shows that individual embeddings have substantial theoretical capacity, but also that practical limits depend strongly on architecture, training, and input complexity (Kuratov et al., 2025). Yet, current compression systems are typically evaluated only via end-task metrics, offering little insight into *when* a single compressed token crosses from informative to overflowed states.

This paper investigates token overflow in soft compression architectures. We ask: **(RQ1)** How can we characterize overflow in compressed representations? **(RQ2)** Can overflow be detected efficiently, without full LLM inference, using lightweight diagnostics? **(RQ3)** Is overflow detectable from compressed tokens alone, or does it require modeling query-context interactions?

To address these questions, we:

- formalize **token overflow** and propose a methodology advancing from query-independent to query-aware detection approaches;
- demonstrate that **saturation statistics** reliably distinguish compressed tokens from standard tokens, providing a practical tool for identifying compressed representations, but show limited overflow detection capability;
- show that **attention patterns** during generation provide moderate overflow signal but

require LLM forward passes;

- develop **learned probing classifiers** operating on joint query-context representations that achieve strong overflow detection without LLM inference (Table 1), showing that incorporating query information improves detection performance;

Although our experiments focus on the xRAG architecture, the methodology is general and we expect it to yield similar results when applied to other setups. The source code is publicly available online¹.

2 Related Work

Long-context modeling and compression Efficient long-context processing has been tackled through architectural changes (Beltagy et al., 2020; Zaheer et al., 2020; Dai et al., 2019) and explicit compression. Context compression is systematized into *hard*, *soft*, and *hybrid* paradigms (Liao et al., 2025): hard compression selects token subsets with strict information bottlenecks; soft compression maps contexts into dense vectors accessible via attention; hybrid methods combine both approaches. Our work analyzes the *failure modes* of soft compression, asking when compressed vectors fail to carry useful task information.

Soft compression in RAG Retrieval-augmented generation (RAG) frameworks extend LLMs with external corpora (Lewis et al., 2020), motivating compression of retrieved passages. Several soft compression methods have been proposed: Auto-Compressors (Chevalier et al., 2023) learn summary vectors by training the model to reconstruct compressed context through attention; ICAE (Ge et al., 2024) employs in-context autoencoding to compress sequences into memory slots with combined reconstruction and language modeling objectives ($\sim 4\times$ compression, 1% parameters). We focus our experiments on xRAG (Cheng et al., 2024), utilizing it not merely as a baseline, but as a representative *projector-based* compression paradigm. Unlike autoencoder-based methods that compress context via complex recurrence or reconstruction objectives, xRAG treats dense retrieval embeddings as a distinct modality. It employs a lightweight projector to map these embeddings directly into the LLM’s input space. This architectural choice

isolates the compression mechanism from the complexities of extensive parameter fine-tuning ($<0.1\%$ parameters), allowing us to study the interactions between the projector and the frozen LLM in a controlled setting. By decoupling the retrieval representation from the generative process, xRAG provides clean access to pre- and post-projection states, making it an ideal testbed for analyzing signal degradation and token saturation without the confounders of end-to-end model adaptation.

Motivation for overflow detection Detecting *information overflow* – where input complexity exceeds compressed token capacity – is critical for optimizing RAG pipelines. It enables *adaptive chunking*, allowing systems to dynamically resize input segments based on semantic density rather than arbitrary fixed lengths. Furthermore, *early overflow detection* facilitates computational pruning: identifying and discarding saturated representations immediately after projection prevents wasteful LLM inference on degraded context. Despite progress across methods, most evaluations treat compressed vectors as black boxes and focus on downstream metrics (Ge et al., 2024; Cheng et al., 2024). Recent work shows single vectors can theoretically encode thousands of tokens, yet *practical* capacity depends on architecture and complexity (Kuratov et al., 2025). In contrast, we *operationalize* capacity limits through overflow detection in xRAG, advancing from query-independent saturation statistics to query-aware learned probing.

3 Methodology

Our goal is to characterize and detect *token overflow* in soft compression architectures across tasks and context regimes. We focus on xRAG-style compressors attached to frozen LLM backbones, studying how compressed token properties change as context complexity increases and downstream quality degrades. We employ a spectrum of detection approaches with increasing query-awareness: from query-agnostic saturation statistics, through query-conditioned attention patterns, to fully query-aware learned probing classifiers.

Our methodology is motivated by the observation that the same compressed representation may be sufficient for one query but overflowed for another. This motivates our approach, advancing from query-independent to query-aware detection:

1. **Context complexity and saturation statistics** (query-agnostic): Measure intrinsic prop-

¹<https://github.com/s-nlp/overflow-detection>

erties of compressed representations independent of any query – useful for identifying and characterizing compressed tokens.

2. **Attention features** (query-conditioned): Analyze how the LLM utilizes compressed tokens during generation for a specific query – captures behavioral signals but requires LLM forward passes.
3. **Learned probing** (query-aware): Train classifiers on joint query-context representations to detect overflow in embedding space – achieves strong detection without LLM inference.

This approach allows us to evaluate whether overflow detection improves as query information is incorporated, while identifying the most efficient deployment strategy.

3.1 Problem Setup

Let \mathcal{M} be a frozen LLM and \mathcal{C} a soft compression module (e.g., xRAG’s modality-fusion compressor) that maps an input sequence of n tokens with embeddings $\mathbf{X} \in \mathbb{R}^{n \times d}$ to $k \ll n$ compressed tokens $\mathbf{C} = \mathcal{C}(\mathbf{X}) \in \mathbb{R}^{k \times d}$. The compressed tokens are then injected into \mathcal{M} (e.g., as extra prefix tokens or interleaved context) and used to solve a downstream task such as extractive QA.

Given an input instance i with original context x_i , question q_i , and gold output y_i , we define task performance under compression, $\mathcal{T}_i(\mathbf{C}_i)$, as a scalar metric (e.g., F1, EM, or ROUGE), indicating whether the generated answer is judged correct. We compare it to a reference performance $\mathcal{T}_i^{\text{ref}}$ obtained from either (i) an uncompressed baseline (full context within the model’s window), or (ii) a lightly compressed setting where degradation is empirically negligible.

We define an *overflow* state for instance i as:

$$\mathcal{O}_i = \mathbf{1}(\mathcal{T}_i^{\text{ref}} = 1 \wedge \mathcal{T}_i(\mathbf{C}_i) = 0). \quad (1)$$

Our objective is to (a) understand how compressed representations differ between overflow and non-overflow regimes, and (b) learn detectors that can predict overflow from representations alone, without recomputing \mathcal{T}_i .

More generally, this formulation can be extended by defining overflow via a degradation threshold,

$$\mathcal{T}_i^{\text{ref}} - \mathcal{T}_i(\mathbf{C}_i) \geq \epsilon, \quad (2)$$

where ϵ is task-dependent. Exploring such threshold-based criteria with alternative evaluation functions is a promising direction for future work.

3.2 Context Complexity Measures

For each input context x_i (or aggregated retrieved context), we compute a set of *context complexity* measures intended to approximate how “hard” the context is to compress:

- **Context length** N_{ctx} : the number of tokens in the original, uncompressed context before any truncation. This is the simplest proxy for potential compression pressure and directly correlates with computational cost.
- **Language-model perplexity** PPL_i : the average per-token negative log-likelihood under the base LLM (without compression), which captures how predictable the context is given the model’s training distribution. Higher perplexity indicates linguistically or semantically atypical content.
- **Statistical compressibility** R_i : the compression ratio achieved by a standard lossless compressor (e.g., gzip or LZMA) on the raw text. We define $R_i = \frac{|x_i|_{\text{bytes}}}{|\text{zip}(x_i)|_{\text{bytes}}}$, where larger values indicate more redundancy and thus higher statistical compressibility.

These metrics allow us to analyze how overflow correlates with raw length, lexical predictability, and sequence-level redundancy (compressibility).

3.3 Token Saturation Statistics

We quantify saturation at the level of compressed tokens and their propagated hidden states. For each compressed token vector $\mathbf{c} \in \mathbb{R}^d$ and its corresponding hidden states $\mathbf{h}^{(\ell)}$ at layer ℓ , we compute the following statistics.

Hoyer’s sparsity Hoyer’s index (Hoyer, 2004) measures how concentrated a vector’s energy is in a few dimensions:

$$H(\mathbf{v}) = \frac{\sqrt{d} - \frac{\|\mathbf{v}\|_1}{\|\mathbf{v}\|_2}}{\sqrt{d} - 1}. \quad (3)$$

It ranges from 0 (all components equal) to 1 (only one non-zero component). Informative compressed tokens are hypothesized to exhibit higher sparsity (structured, selective activations), while overflowed tokens tend towards low sparsity (flat, noise-like patterns).

Spectral entropy We apply a discrete cosine transform (DCT) to \mathbf{v} and treat the normalized squared magnitudes as an energy distribution p over frequency components. The spectral entropy is defined as following:

$$S(\mathbf{v}) = - \sum_{i=1}^d p_i \log p_i, \quad p_i = \frac{|\text{DCT}(\mathbf{v})_i|^2}{\|\text{DCT}(\mathbf{v})\|_2^2}. \quad (4)$$

Low entropy corresponds to concentrated energy (structured signals), whereas near-maximum entropy indicates white-noise-like spectra.

Kurtosis We compute the excess kurtosis of the entries of \mathbf{v} :

$$K(\mathbf{v}) = \frac{\mathbb{E}[(v_j - \mu)^4]}{\sigma^4} - 3, \quad (5)$$

where μ and σ are the mean and standard deviation across dimensions. Heavy-tailed distributions (positive kurtosis) suggest a few large, informative coordinates, while overflowed tokens are expected to approach Gaussian-like behavior ($K \approx 0$).

3.4 Attention Features: Query-conditioned Overflow Signals

While saturation statistics measure intrinsic token properties, they ignore how the LLM actually *uses* compressed tokens during generation. To capture this behavioral dimension, we extract attention-based features that quantify the model’s reliance on xRAG tokens when processing a specific query.

For each instance, we perform a forward pass through the LLM with both the query and compressed context, extracting attention weights $\mathbf{A} \in \mathbb{R}^{L \times H \times T \times T}$ across all layers L , heads H , and sequence positions T . We compute:

Mean attention to xRAG tokens For each layer ℓ and head h , we measure the average attention mass directed to compressed token positions:

$$\bar{a}_{\text{xRAG}}^{(\ell,h)} = \frac{1}{|T_q|} \sum_{i \in T_q} \sum_{j \in T_{\text{xRAG}}} A_{i,j}^{(\ell,h)}, \quad (6)$$

where T_q denotes query token positions and T_{xRAG} denotes xRAG token positions. We aggregate across layers and heads to obtain instance-level statistics: mean, max, min, and standard deviation of attention to xRAG tokens.

Attention ratios To contextualize xRAG attention, we compute ratios comparing attention to compressed versus uncompressed tokens:

$$r_{\text{xRAG}/\text{non-xRAG}} = \frac{\bar{a}_{\text{xRAG}}}{\bar{a}_{\text{non-xRAG}}}. \quad (7)$$

This ratio isolates whether the model preferentially attends to compressed representations or relies more heavily on other context.

Attention entropy For each query position i , we compute the entropy of its attention distribution over all positions:

$$\text{Ent}_i = - \sum_{j=1}^T A_{i,j} \log A_{i,j}. \quad (8)$$

High entropy indicates diffuse attention (potentially signaling uncertainty or lack of relevant information), while low entropy indicates focused attention to specific tokens (Rykov et al., 2025).

3.5 Overflow Detection Methods

We evaluate overflow detection through two complementary approaches that span a spectrum from interpretable to representational methods. First, we test **feature-based classification** using explicit, hand-crafted features to determine whether overflow manifests in interpretable, low-dimensional signals. Second, we develop **learned probing classifiers** that operate directly on high-dimensional query and context embedding vectors.

Feature-based classification We aggregate the hand-crafted features described in §3.2–§3.4 (context complexity, saturation statistics, attention patterns) and train a **logistic regression classifier** implemented in scikit-learn². All hyperparameters are detailed in Appendix D.

Learned probing on vector representations

While feature-based methods offer interpretability, they may fail to capture complex interactions between query and context that manifest in the geometry of representation spaces. We therefore develop **learned probing classifiers** that operate directly on joint query-context representations. Our hypothesis is that overflow detection requires modeling alignment patterns in shared representation space.

²<https://scikit-learn.org>

Representation extraction For each instance i with query q_i and context x_i , we extract embeddings at multiple stages:

- *Query representations*: $q_i^{\text{preproj}} \in \mathbb{R}^{d_{\text{ret}}}$ (retriever embedding), $q_i^{\text{postproj}} \in \mathbb{R}^{d_{\text{LLM}}}$ (after projection), $q_i^{\text{mid}}, q_i^{\text{last}} \in \mathbb{R}^{d_{\text{LLM}}}$ (hidden states from intermediate and final LLM layers).
- *Context representations*: $x_i^{\text{preproj}} \in \mathbb{R}^{d_{\text{ret}}}$ (retriever embedding), $x_i^{\text{postproj}} \in \mathbb{R}^{d_{\text{LLM}}}$ (compressed token after projection), $x_i^{\text{mid}}, x_i^{\text{last}} \in \mathbb{R}^{d_{\text{LLM}}}$ (hidden states from intermediate and final layers).

We construct **joint feature vectors** by concatenating query and context representations at matched or complementary stages:

$$\phi_i = [x_i^{(s_c)}; q_i^{(s_q)}], \quad (9)$$

where $s_c, s_q \in \{\text{preproj}, \text{postproj}, \text{mid}, \text{last}\}$ denote the extraction stage. Our primary experiments use projection-stage representations (preproj, postproj) which are available immediately after encoding without requiring LLM forward passes. Following prior work demonstrating that intermediate transformer layers encode complementary information useful for interpretation tasks (CH-Wang et al., 2024; Belikova et al., 2025), we additionally evaluate multi-layer representations (including mid, last) to assess the efficiency-accuracy trade-off.

Classifier architectures To systematically assess the role of model capacity and training objectives in overflow detection, we evaluate three neural probe architectures:

- *Linear Probe*: A single linear transformation applied to the joint feature vector ϕ_i . This minimal architecture tests whether overflow is linearly separable in the concatenated representation space.
- *MLP Probe*: A two-layer feedforward network with one hidden layer, introducing nonlinear feature interactions while maintaining computational efficiency.
- *MLP Probe with Supervised Contrastive Learning (SCL)*: An MLP trained with a hybrid objective that combines standard binary cross-entropy with a supervised contrastive term (Khosla et al., 2020). This architecture

explicitly structures the representation space by encouraging same-class instances to cluster while pushing apart opposite-class instances.

For the SCL probe, we minimize the combined objective

$$\mathcal{L} = \mathcal{L}_{\text{BCE}} + \lambda \mathcal{L}_{\text{SCL}}, \quad (10)$$

where \mathcal{L}_{BCE} provides direct classification supervision, while \mathcal{L}_{SCL} imposes metric constraints by maximizing cosine similarity between same-label pairs and minimizing it between different-label pairs in the learned representation space.

4 Results

4.1 Experimental Setup

As a preliminary study, we use the xRAG-7B model³ as the base LLM and SFR-Embedding-Mistral⁴ as the retriever embedding model for all experiments. We focus on three extractive question answering datasets: **SQuADv2** (Rajpurkar et al., 2018), a context-based QA benchmark over Wikipedia passages; **TriviaQA** (Joshi et al., 2017), a large-scale reading comprehension dataset with independently collected evidence documents; and **HotpotQA** (Yang et al., 2018), a multi-hop reasoning dataset requiring information synthesis across multiple paragraphs. All three datasets were part of the xRAG compression module’s training data, providing a realistic testbed for studying overflow in deployed systems. We use test set examples that were *correctly* answered with uncompressed context, filtering out instances where the model fails regardless of compression. This ensures overflow detection focuses on compression-induced failures rather than inherent task difficulty.

Evaluation protocol Answer correctness is evaluated using GPT-4o-mini for SQuADv2, which assesses semantic equivalence between generated and ground-truth answers. For TriviaQA and HotpotQA, we apply a substring-based exact-match criterion: predictions are marked correct if they contain any reference answer as a substring. All classifiers are evaluated using 5-fold stratified cross-validation; training and hyperparameter details are provided in §3.5 and Appendix D.

Stage	Features	TriviaQA	SQuADv2	HotpotQA
Pre-compression	Context	0.589 ± 0.019	0.605 ± 0.025	0.541 ± 0.017
Pre-inference	Representation	0.687 ± 0.015	0.662 ± 0.015	0.653 ± 0.023
	Representation-joint	0.725 ± 0.021	<u>0.703 ± 0.019</u>	<u>0.720 ± 0.007</u>
	Saturation	0.568 ± 0.022	0.529 ± 0.024	0.533 ± 0.010
Post-inference	Attention	0.627 ± 0.012	0.608 ± 0.020	0.623 ± 0.013
	Representation	0.684 ± 0.016	0.665 ± 0.016	0.655 ± 0.024
	Representation-joint	0.719 ± 0.015	0.713 ± 0.016	0.733 ± 0.011
	Saturation	<u>0.583 ± 0.019</u>	0.546 ± 0.011	0.549 ± 0.008
	Saturation-joint	0.600 ± 0.014	0.585 ± 0.016	0.633 ± 0.015

Table 1: Overflow prediction performance (ROC-AUC) across different pipeline stages. **Pre-compression**: a priori context features. **Pre-inference**: pre-LLM inference combining preprojection and postprojection features. **Post-inference**: post-LLM inference combining features from middle and last layers. Representation-joint combines query and context representations. **Bold**: best performance per dataset; underlined: second-best.

4.2 Main Results

We organize our findings around the three research questions posed in §1, advancing from characterizing overflow (**RQ1**) through efficient detection methods (**RQ2**) to understanding the role of query-context interactions (**RQ3**). Table 1 presents our main results across three datasets, comparing detection performance at different pipeline stages. We distinguish between two detection stages: **pre-inference** (before LLM processing) uses concatenated preprojection and postprojection representations for probing, while **post-inference** (requiring full LLM forward pass) uses concatenated middle and last layer hidden states. For saturation statistics, the same stage distinction applies, with features extracted from corresponding layer representations. Table 3 (Appendix B) provides a detailed ablation study examining feature extraction at different architectural layers.

4.2.1 RQ1: Characterizing overflow in compressed representations

To understand the nature of overflow, we first examined whether compressed tokens exhibit distinctive geometric properties that might correlate with information loss.

Saturation statistics distinguish token types but not overflow

We compared the defined saturation statistics across xRAG and non-xRAG tokens at multiple LLM layers across all three datasets. To avoid positional bias and control for contextual confounds, we compare xRAG token statistics against four baselines: (i) *mean of all non-xRAG tokens*

(when compression is applied), capturing the aggregate behavior of standard tokens in compressed sequences; (ii) *mean of original context tokens*, representing uncompressed context behavior; (iii) *first original context token*, isolating position-specific effects; and (iv) *first token in no-context scenarios*, establishing a baseline without any context information.

Tables 2, 4, and 5 (Appendix C) present percentage differences across these baselines. The results reveal consistent patterns across datasets: xRAG tokens show lower sparsity and kurtosis, and dramatically higher spectral entropy across all layers (all $p < 0.001$). Spectral entropy shows the largest differences (87% across all datasets and baselines), while excess kurtosis shows substantial differences ranging from 29–98% depending on layer and baseline. Hoyer’s sparsity demonstrates more modest but consistent differences of 7–33%.

Crucially, these patterns remain remarkably stable across datasets and all four baselines, validating that observed properties reflect genuine characteristics of xRAG tokens rather than measurement artifacts or positional biases. The differences persist from middle to final layers, suggesting that compression effects propagate through the network without being normalized away. To verify that these representational differences enable token-type identification, we tested linear separability between xRAG and non-xRAG tokens across all baseline configurations. Linear classifiers achieve near-perfect separation (> 0.95 AUC-ROC for all variants), confirming that saturation statistics reliably distinguish compressed from uncompressed representations in the model’s activation space. Notably, more complex classifier architectures (MLP, MLP-SCL) provide no improvement over linear

³<https://hf.co/Hannibal046/xrag-7b>

⁴<https://hf.co/Salesforce/SFR-Embedding-Mistral>

models for this token-type classification task (see Appendix A), further confirming that compressed and uncompressed tokens occupy distinctly separable regions.

However, while these metrics successfully *characterize* compressed tokens, they fail to *predict overflow*. Despite the substantial magnitude of differences and near-perfect linear separability of token types, saturation statistics achieve only near-random predictive performance for overflow detection across datasets (Table 1). Even when combined with query information (Saturation-joint), performance remains limited (0.55–0.63 AUC-ROC).

Context complexity provides minimal signal

Context-level features (shown in Table 1) also achieve near-random performance, only marginally exceeding saturation statistics. This indicates that overflow is not strongly predicted by general context properties alone (perplexity, length, statistical compressibility) in our experimental setting. While our current datasets involve relatively short passages compressed into single tokens, we suggest that context complexity features may become more informative in settings with substantially longer contexts or more extreme compression ratios.

Summary for RQ1: Saturation statistics provide a reliable method to separate compressed tokens from uncompressed tokens, achieving near-perfect linear separability and revealing distinct activation-space statistics with 7–87% relative differences across multiple metrics, layers, and tokens. However, these query-agnostic properties do not predict task-relevant information loss, indicating that while compressed tokens are distinct in representation space, overflow detection requires modeling query-context interactions beyond intrinsic token characteristics.

4.2.2 RQ2: Efficient overflow detection without full LLM inference

While saturation statistics and context complexity features show limited predictive power, we investigated whether learned classifiers can effectively detect overflow, and critically, *at which stage in the compression pipeline* degradation becomes detectable. Tables 1 and 3 compare detection performance at two stages: **pre-inference** (projection stage, before LLM processing) and **post-inference** (LLM hidden states, after forward pass).

Overflow is detectable immediately after compression

Learned probing classifiers achieve 0.72 AUC-ROC on average at the post-projection stage (Table 1), substantially outperforming context-only models and query-agnostic baselines. Crucially, *compression-induced information loss manifests in the representation space immediately after projection*, before any LLM processing. The overflow signal is already present in query-context alignment patterns, revealing that degradation is determined by the compression step itself rather than emerging during generation.

LLM processing provides no additional signal

Post-inference detection using middle-layer hidden states achieves identical performance, confirming that overflow established at compression time merely propagates through the network without amplification or masking (Table 3). Attention patterns (0.62 AUC-ROC on average) and saturation statistics (even when query-conditioned) provide no meaningful improvement over projection-stage features. Last-layer features show slightly degraded performance, suggesting earlier layers better preserve overflow-relevant signals.

Summary for RQ2: Overflow detection without LLM inference matches post-inference performance. This reveals that compression degradation manifests immediately after projection and is determined during compression rather than during generation, enabling both efficient detection and deeper understanding of compression capacity limits.

4.2.3 RQ3: The necessity of modeling query-context interactions

Our final question addresses whether overflow can be detected from compressed tokens alone or whether incorporating query information improves detection performance.

Joint representations substantially outperform single-source models

Tables 1 and 3 compare detection using context-only representations (Representation), joint query-context representations (Representation-joint), and query-agnostic saturation statistics. Given the poor performance of saturation statistics alone, we explored whether incorporating contextual information could improve detection by aggregating statistics (Hoyer’s sparsity, spectral entropy, and excess kurtosis) from all non-xRAG tokens in the compressed sequence, computing their mean, maximum, minimum, and

Stage	Statistic	Non-xRAG	Context (first)	Context (mean)	No Context
Middle Layer	Excess Kurtosis	92.0	29.1	29.6	-23.1
	Hoyer’s index	24.4	23.0	19.4	20.1
	Spectral Entropy	0.1	87.1	87.1	87.1
Last Layer	Excess Kurtosis	98.5	98.8	94.1	80.4
	Hoyer’s index	31.4	32.1	19.2	7.2
	Spectral Entropy	0.1	87.1	87.1	87.0

Table 2: Relative differences (%) in saturation statistics between xRAG and baseline tokens on SQuADv2, computed as $\frac{\text{baseline}-\text{xRAG}}{\text{baseline}} \times 100\%$. **Middle Layer:** LLM intermediate layer features. **Last Layer:** LLM final layer features. Positive values indicate xRAG tokens have lower saturation (more structured representations). Large differences ($\geq 50\%$) in Excess Kurtosis and Spectral Entropy demonstrate consistent xRAG-specific properties across multiple baselines.

standard deviation (Saturation-joint).

The results reveal a clear hierarchy: representation-joint models achieve 0.70–0.73 AUC-ROC across datasets and stages, substantially outperforming context-only models (0.64–0.69 AUC-ROC). Saturation-joint yields only modest improvements over saturation-only features (0.58–0.63 vs. 0.52–0.58 AUC-ROC), remaining substantially below representation-based methods. This confirms that token-level activation statistics alone are insufficient for overflow detection, regardless of aggregation strategy.

This reveals that overflow is not an intrinsic property of compressed representations but emerges from the *mismatch* between what information the compressed token contains and what the query requires. Joint representation models capture this alignment directly in representation space, enabling accurate overflow prediction. Notably, saturation statistics maintain consistent low performance across all pipeline stages, confirming their utility for *identifying* compressed tokens but not for *predicting* query-specific overflow. Similar to token-type classification (RQ1), linear classifiers prove sufficient for overflow detection, with more complex architectures providing minimal improvement (Appendix A), suggesting that overflow manifests as relatively simple (approximately linearly separable) structure in joint representation space.

Summary for RQ3: Overflow detection fundamentally requires modeling query-context interactions. Joint representation models yield the strongest performance, outperforming context-only models by 5–8 percentage points (Table 1).

5 Conclusion

We investigated token overflow in soft compression architectures and proposed a methodology advancing from query-independent to query-aware

detection. Our findings show that saturation statistics reliably separate compressed from uncompressed tokens (7–87% relative differences), while learned probing on joint query-context representations achieves efficient pre-inference overflow detection (0.72 AUC-ROC on average) without LLM forward passes. Post-inference detection achieves comparable performance, confirming that overflow can be detected efficiently before expensive LLM processing. These results enable safer deployment of compression modules through low-cost pre-LLM gating and adaptive chunking strategies.

Limitations

This work focuses on the xRAG architecture as an initial controlled study. Future work should extend the methodology to longer contexts, diverse tasks (summarization, multi-hop reasoning), and other compression architectures to validate generalizability. Exploring richer overflow definitions beyond task performance degradation could capture subtle information loss patterns. Our detection performance establishes a strong baseline, with promising directions including multi-task learning across different compression ratios, incorporating architectural features of the compressor, and developing adaptive systems that dynamically adjust compression based on predicted overflow risk. The methodology’s architecture-agnostic design facilitates such extensions to emerging compression techniques.

Ethical Considerations

Generation of text with LLMs using compressed and overflowed tokens can lead to hallucinations and untrustworthy output. This way, the created technology may be considered helpful for minimizing such effects. At the same time, as the absolute accuracy numbers of the developed classifier are rela-

tively low, and eventual false positive predictions could lead to overconfidence in trustworthiness of generated texts. Therefore, we suggest that more research is needed to raise the absolute values of the developed classifiers to ensure their safe use in various text generation workflows and applications.

Acknowledgements

The work of Alexander Panchenko was supported by the RSF project № 25-71-30008 “Laboratory for reliable, adaptive, and trustworthy Artificial Intelligence”.

References

- Islam Aushev, Egor Kratkov, Evgenii Nikolaev, Andrei Glinskii, Vasilii Krikunov, Alexander Panchenko, Vasily Kononov, and Julia Belikova. 2025. [RAGulator: Effective RAG for regulatory question answering](#). In *Proceedings of the 1st Regulatory NLP Workshop (RegNLP 2025)*, pages 114–120, Abu Dhabi, UAE. Association for Computational Linguistics.
- Julia Belikova, Konstantin Polev, Rauf Parchiev, and Dmitry Simakov. 2025. [Data-efficient meta-models for evaluation of context-based questions and answers in llms](#). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2025, Padua, Italy, July 13-18, 2025*, pages 4385–4389. ACM.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *CoRR*, abs/2004.05150.
- Sky CH-Wang, Benjamin Van Durme, Jason Eisner, and Chris Kedzie. 2024. [Do androids know they’re only dreaming of electric sheep?](#) In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 4401–4420. Association for Computational Linguistics.
- Xin Cheng, Xun Wang, Xingxing Zhang, Tao Ge, Si-Qing Chen, Furu Wei, Huishuai Zhang, and Dongyan Zhao. 2024. [xrag: extreme context compression for retrieval-augmented generation with one token](#). In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS ’24*, Red Hook, NY, USA. Curran Associates Inc.
- Alexis Chevalier, Alexander Wettig, Anirudh Ajith, and Danqi Chen. 2023. [Adapting language models to compress contexts](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 3829–3846. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. 2019. [Transformer-xl: Attentive language models beyond a fixed-length context](#). *CoRR*, abs/1901.02860.
- Tao Ge, Jing Hu, Lei Wang, Xun Wang, Si-Qing Chen, and Furu Wei. 2024. [In-context autoencoder for context compression in a large language model](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Patrik O. Hoyer. 2004. [Non-negative matrix factorization with sparseness constraints](#). *J. Mach. Learn. Res.*, 5:1457–1469.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. [Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension](#). *Preprint*, arXiv:1705.03551.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. [Supervised contrastive learning](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Yuri Kuratov, Mikhail Arkhipov, Aydar Bulatov, and Mikhail Burtsev. 2025. [Cramming 1568 tokens into a single vector and back again: Exploring the limits of embedding space capacity](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 19323–19339. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Huanxuan Liao, Wen Hu, Yao Xu, Shizhu He, Jun Zhao, and Kang Liu. 2025. [Beyond hard and soft: Hybrid context compression for balancing local and global information retention](#). *Preprint*, arXiv:2505.15774.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Trans. Assoc. Comput. Linguistics*, 12:157–173.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.

- Elisei Rykov, Valerii Olisov, Maksim Savkin, Artem Vazhentsev, Kseniia Titova, Alexander Panchenko, Vasily Konovalov, and Julia Belikova. 2025. [Smur-fCat at SemEval-2025 task 3: Bridging external knowledge and model uncertainty for enhanced hallucination detection](#). In *Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025)*, pages 1034–1045, Vienna, Austria. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [Hotpotqa: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2369–2380. Association for Computational Linguistics.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. [Big bird: Transformers for longer sequences](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

A Classifiers Ablation Study

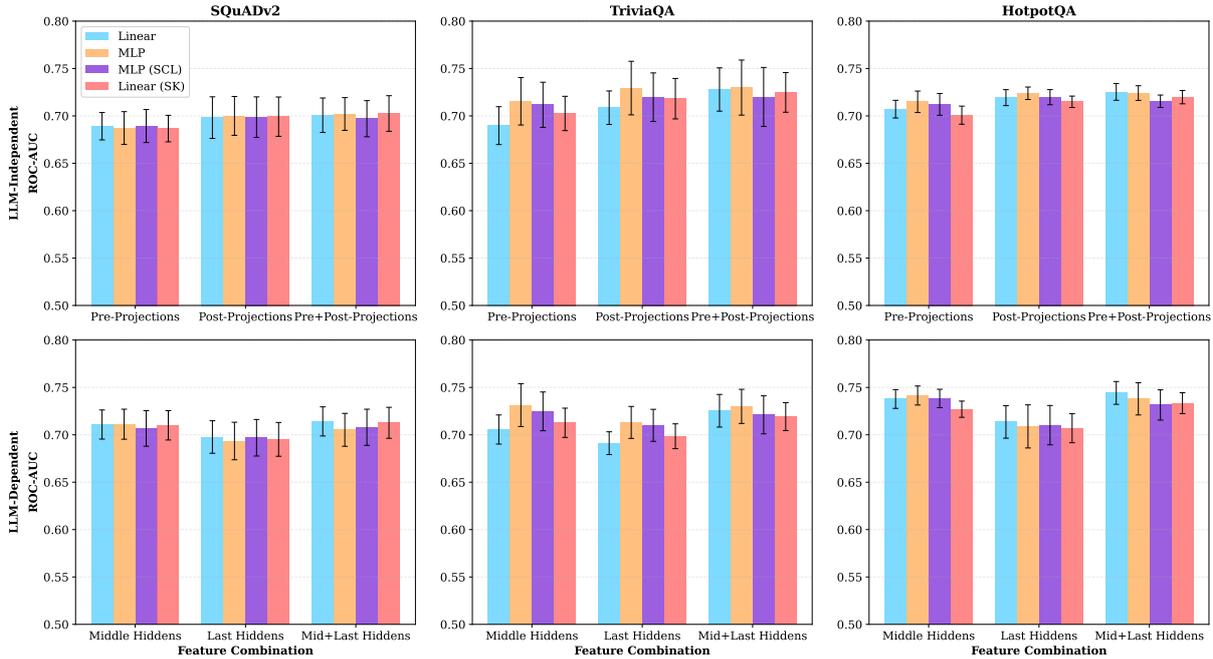


Figure 1: Comparison of classifier architectures (Linear scikit-learn, Linear PyTorch, MLP, MLP with SCL) across datasets and feature combinations. All architectures achieve comparable performance, with differences typically <1 percentage point, demonstrating that overflow is largely linearly separable in joint representation space.

B Features Ablation Study

Stage	Features	TriviaQA	SQuADv2	HotpotQA
Pre-projection	Representation	0.679 ± 0.019	0.641 ± 0.013	0.635 ± 0.017
	Representation-joint	0.703 ± 0.018	0.687 ± 0.014	0.701 ± 0.010
	Saturation	0.534 ± 0.016	0.526 ± 0.024	0.522 ± 0.009
Post-projection	Representation	0.675 ± 0.016	0.660 ± 0.016	0.652 ± 0.025
	Representation-joint	0.718 ± 0.021	<u>0.699 ± 0.021</u>	<u>0.715 ± 0.006</u>
	Saturation	0.553 ± 0.017	0.530 ± 0.011	0.518 ± 0.011
Middle layer	Attention	0.620 ± 0.012	0.581 ± 0.008	0.603 ± 0.014
	Representation	0.675 ± 0.016	0.660 ± 0.016	0.652 ± 0.025
	Representation-joint	0.713 ± 0.015	0.710 ± 0.015	0.727 ± 0.009
	Saturation	0.553 ± 0.017	0.530 ± 0.011	0.518 ± 0.011
	Saturation-joint	0.557 ± 0.021	0.560 ± 0.026	0.597 ± 0.017
Last layer	Attention	0.584 ± 0.011	0.584 ± 0.027	0.615 ± 0.014
	Representation	0.675 ± 0.016	0.661 ± 0.016	0.653 ± 0.024
	Representation-joint	0.698 ± 0.013	0.695 ± 0.018	0.707 ± 0.015
	Saturation	0.554 ± 0.020	0.523 ± 0.015	0.530 ± 0.010
	Saturation-joint	0.582 ± 0.015	0.541 ± 0.011	0.618 ± 0.011

Table 3: Ablation study examining feature extraction at different architectural stages (ROC-AUC). **Pre-projection**: retriever embeddings before projection. **Post-projection**: representations after projection. **Middle layer**: intermediate LLM hidden states. **Last layer**: final LLM hidden states. Representation-joint combines query and context representations. **Bold**: best performance per dataset; underlined: second-best.

C Saturation Statistics

Stage	Statistic	Non-xRAG	Context (first)	Context (mean)	No Context
Middle Layer	Excess Kurtosis	92.4	41.3	40.5	-4.4
	Hoyer’s index	24.4	26.1	20.6	21.1
	Spectral Entropy	0.1	87.1	87.1	87.1
Last Layer	Excess Kurtosis	98.3	98.4	93.1	81.4
	Hoyer’s index	28.4	28.1	16.0	7.4
	Spectral Entropy	0.0	87.1	87.1	87.0

Table 4: Relative differences (%) in saturation statistics between xRAG and baseline tokens on TriviaQA, computed as $\frac{\text{baseline}-\text{xRAG}}{\text{baseline}} \times 100\%$. **Middle Layer:** LLM intermediate layer features. **Last Layer:** LLM final layer features. Positive values indicate xRAG tokens have lower saturation (more structured representations). Large differences ($\geq 50\%$) in Excess Kurtosis and Spectral Entropy demonstrate consistent xRAG-specific properties across multiple baselines.

Stage	Statistic	Non-xRAG	Context (first)	Context (mean)	No Context
Middle Layer	Excess Kurtosis	91.6	39.7	46.0	-10.6
	Hoyer’s index	24.4	25.1	22.0	20.8
	Spectral Entropy	0.1	87.1	87.1	87.1
Last Layer	Excess Kurtosis	98.1	98.9	92.3	80.8
	Hoyer’s index	28.3	33.5	14.7	7.0
	Spectral Entropy	0.0	87.1	87.1	87.0

Table 5: Relative differences (%) in saturation statistics between xRAG and baseline tokens on HotpotQA, computed as $\frac{\text{baseline}-\text{xRAG}}{\text{baseline}} \times 100\%$. **Middle Layer:** LLM intermediate layer features. **Last Layer:** LLM final layer features. Positive values indicate xRAG tokens have lower saturation (more structured representations). Large differences ($\geq 50\%$) in Excess Kurtosis and Spectral Entropy demonstrate consistent xRAG-specific properties across multiple baselines.

D Hyperparameters

Table 6 summarizes all hyperparameters used in our experiments. All hyperparameters were tuned on the SQuADv2 validation set and then fixed across TriviaQA and HotpotQA datasets.

Method	Parameter	Value
<i>Feature-based Classification (Logistic Regression)</i>		
	Solver	lbfgs
	Regularization	L2, $C = 10^{-5}$
	Max iterations	1000
	Preprocessing	StandardScaler
<i>Common Settings (All Neural Probes)</i>		
	Cross-validation	5-fold stratified, 80%/20% train/val
	Batch size	256
	Optimizer	Adam, learning rate = 10^{-4}
	Preprocessing	StandardScaler
	Regularization	$\mathcal{L}_{\text{reg}} = \frac{\lambda_2}{2N} \ \theta\ _2^2 + \frac{\lambda_1}{N} \ \theta\ _1$ with $(\lambda_2, \lambda_1) = (500, 100)$
	Early stopping	Patience = 20 epochs
<i>Linear Probe</i>		
	Architecture	Single linear layer
	Max epochs	150
<i>MLP Probe</i>		
	Architecture	Two-layer feedforward
	Hidden dimension	1024
	Activation	ReLU
	Max epochs	50
<i>MLP-SCL Probe</i>		
	Architecture	Two-layer feedforward
	Hidden dimension	1024
	Activation	SiLU
	Dropout	0.1 (before and after hidden layer)
	Normalization	BatchNorm1d after hidden layer
	Contrastive weight	$\lambda = 0.3$
	Temperature	$\tau = 0.07$
	Max epochs	50

Table 6: Hyperparameters for all overflow detection methods. The regularization term for neural probes combines L2 and L1 penalties scaled by the number of model parameters N (excluding biases).