

# 000 NEURAL MESSAGE-PASSING ON ATTENTION GRAPHS 001 FOR HALLUCINATION DETECTION 002

003 **Anonymous authors**

004 Paper under double-blind review

## 005 ABSTRACT

006 Large Language Models (LLMs) often generate incorrect or unsupported content,  
007 known as hallucinations. Existing detection methods rely on heuristics or simple  
008 models over isolated computational traces such as activations, or attention maps.  
009 We unify these signals by representing them as *attributed graphs*, where tokens  
010 are nodes, edges follow attentional flows, and both carry features from attention  
011 scores and activations. Our approach, CHARM, casts hallucination detection as a  
012 graph learning task and tackles it by applying GNNs over the above attributed  
013 graphs. We show that CHARM provably subsumes prior attention-based heuristics  
014 and, experimentally, it consistently outperforms other leading approaches across  
015 diverse benchmarks. Our results shed light on the relevant role played by the graph  
016 structure and on the benefits of combining computational traces, whilst showing  
017 that CHARM exhibits promising zero-shot performance on cross-dataset transfer<sup>1</sup>.  
018

## 019 1 INTRODUCTION

020 Despite their impressive capabilities, LLMs frequently produce outputs that are factually inaccurate,  
021 logically inconsistent, or unsupported by the input context, broadly referred to as *hallucinations*  
022 (Pagnoni et al., 2021; Cao et al., 2022; Qiu et al., 2023). As LLMs are increasingly applied in diverse  
023 domains, detecting hallucinations becomes crucial for ensuring their safe and reliable use. This  
024 phenomenon is inherently complex and multi-faceted, and methods for *automated hallucination*  
025 *detection* (HD) have recently received significant attention (Yin et al., 2024; Bar-Shalom et al., 2025).

026 A straightforward approach for HD is to query LLMs multiple times, either by asking them to judge  
027 their own outputs (Kadavath et al., 2022) or by sampling alternative generations to measure semantic  
028 variability (Kuhn et al., 2023). While effective in some cases, this strategy requires repeated rollouts,  
029 making it both slow and computationally expensive, and thus unsuitable for real-time or large-scale  
030 use. A more scalable line of work leverages the internal signals produced by LLMs during decoding,  
031 which we refer to as *computational traces*. In particular, most works focus on *linearly* probing  
032 residual stream activations on selected layers and token positions (Orgad et al., 2024; Azaria &  
033 Mitchell, 2023; Belinkov, 2022). More recently, attention maps have shown to provide an additional  
034 perspective on model behaviour, e.g., by leveraging prompt-response attention ratios (Chuang et al.,  
035 2024). Although providing meaningful, alternative cues on hallucinations existing attention-based  
036 techniques rely on simple models or handcrafted *heuristics* (Sriramanan et al., 2024; Binkowski  
037 et al., 2025). Furthermore, all the above methods treat computational traces in isolation, despite  
038 capturing complementary aspects of hallucinations. To date, a systematic exploration of the interplay  
039 of computational traces is still lacking. More broadly, the field currently lacks a framework applying  
040 modern deep learning techniques to structured, holistic representations of computational traces,  
041 leaving the community to rely on heuristic, single-signal approaches.

042 In this paper, we propose a unified framework that represents LLM computational traces as *attributed*  
043 *graphs*, a natural, yet under-explored perspective in the HD literature. Similarly to recent works  
044 dealing with the analysis of learnt attention computational flows (Barbero et al., 2024; El et al.,  
045 2025), this formulation considers tokens as nodes and draws connections between them based on  
046 the structure of attention maps calculated during text generation. Crucially, both nodes and edges  
047 can be endowed with features derived from the values of computational traces across layers (and

048 049 050 051 052 053 <sup>1</sup>Code will be released upon acceptance.

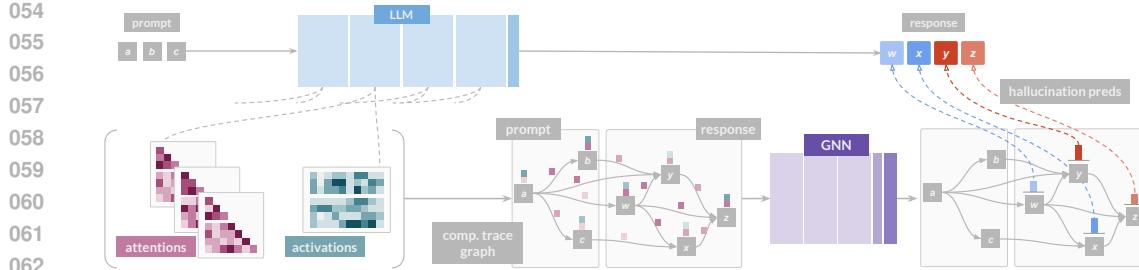


Figure 1: **Overview of CHARM.** We extract attention and activation matrices from LLM computations and build an attributed graph from them: edges and their features are derived from off-diagonal attention scores; node features are based on activations, and diagonal attention values. The resulting graph is processed by a GNN-based architecture, which outputs either token-level hallucination scores (as illustrated) or a global hallucination score for the entire sentence.

heads): node features capture token-wise signals such as activations and self-scores (the attention a token assigns to itself), while edge features encode pairwise interactions, most prominently the attention between distinct tokens. This perspective casts HD as a *graph learning* problem, which has recently obtained successes in broad-ranging domains (Monti et al., 2019; Gonzalez et al., 2021; Liu et al., 2023) and, we argue, is well suited to this task. First, representing computational traces as attributed graphs allows to naturally integrate heterogeneous signals, which may hold varying predictive value across generation tasks. Second, the framework accommodates different levels of detection granularity, with the standard setup of graph classification corresponding to response-level detection, and that of node classification to the token-level one. Finally, this formulation directly leverages the rich body of work on Graph Neural Networks (GNNs) and their code-libraries (Fey & Lenssen, 2019), providing principled and well-studied tools for tailored HD models.

Motivated by these advantages, we introduce CHARM (see Figure 1), an HD approach based on a Graph Neural Network operating on computational trace graphs (Gilmer et al., 2017). Our framework can jointly process different computational traces, and subsumes known detection heuristics: we prove that it can express recent attention-based methods (Chuang et al., 2024; Sriramanan et al., 2024) either at the token or response granularity levels. Experimentally, CHARM *consistently outperforms* these heuristics, as well as other leading methods across benchmarks and detection resolutions. Our analyses further reveal that incorporating activations into computational trace graphs alongside attention-features may improve detection of non-contextual hallucinations. Beyond state-of-the-art comparisons, our ablation studies demonstrate the importance of the graph structure, the main principle driving CHARM. Finally, we report promising zero-shot cross-dataset transfer results and observe robustness to graph sparsifications, indicating viable trade-offs between accuracy and efficiency.

**Contributions** are summarised as follows. (1) We introduce a unified view of LLM *computational traces* as *attributed graphs*, where tokens are nodes connected by attention-induced edges, and both nodes and edges are enriched with features such as activations and attention scores. (2) We introduce CHARM, which casts HD as graph learning on computational trace graphs. It uses a GNN that provably subsumes attention-based heuristics, opening new application frontiers for machine learning on graphs. (3) We show that CHARM consistently outperforms leading HD methods across diverse benchmarks and granularities, while exhibiting promising zero-shot transfer capabilities.

## 2 RELATED WORK

**Hallucinations and their detection in LLMs.** The term “hallucinations” in LLMs broadly refers to errors in text generation where outputs are unfaithful to the input or external facts (Orgad et al., 2024). These include knowledge inaccuracies, flawed reasoning, biases, and references to non-existing (Liu et al., 2021; Huang et al., 2023a; Ji et al., 2023; Rawie et al., 2023). Hallucinations can involve complex failures and manifest in subtle ways, including at the granularity of single tokens (Orgad et al., 2024). Early detection approaches leverage uncertainty measures in next-token prediction or semantic consistency of responses Kadavath et al. (2022); Varshney et al. (2023); Kuhn et al. (2023); Manakul et al. (2023). Alternatively, recent work propose detectors on LLM computational

traces; prominent examples include (hidden) activations Kadavath et al. (2022); Snyder et al. (2024); Yuksekgonul et al. (2023); Zou et al. (2023); Yin et al. (2024); Chen et al. (2024); Simhi et al. (2024); Li et al. (2024); Marks & Tegmark (2023); Burns et al. (2022); Rateike et al. (2023) and attention matrices Sriramanan et al. (2024); Chuang et al. (2024); Binkowski et al. (2025); Bazarova et al. (2025); Zhang et al. (2023). Different traces may be more or less informative for different types of hallucinations: e.g., attention-based heuristics have been evidenced to be predictive in contextual hallucination settings Chuang et al. (2024). Existing methods mostly rely on heuristics or simple classifiers on specific traces; some are also constrained to coarse detection levels (e.g., whole responses) (Sriramanan et al., 2024; Binkowski et al., 2025; Kuhn et al., 2023).

**Attention-based HD.** Irregular or skewed attention behaviours have been observed to often signal pathological text generation (Xu et al., 2023; Chuang et al., 2024; Binkowski et al., 2025; Bazarova et al., 2025). E.g., hallucinated translations may exhibit localised scores on narrow context windows (Xu et al., 2023); hallucinations in contextual question answering may correlate with excessive focus on response tokens w.r.t. context ones (Chuang et al., 2024). The recent *Lookback Lens* proposes a detection feature based on this intuition. Other works extract spectral or structural features from attention matrices, e.g., via graph Laplacians combined with logistic regression (Binkowski et al., 2025). These approaches, however, remain limited by fixed heuristics and shallow classifiers. Our work generalises this line by employing attention matrices to construct attributed graphs, a formulation that supports predictions at multiple levels of granularity, integrates additional computational signals and unlocks the application of modern (graph-based) deep learning techniques.

**Graphs of LLM computation and Graph Neural Networks.** Recent works have applied graph-theoretic perspectives to neural computations (Vitytksyi et al., 2025), often graphs induced by attention matrices (Barbero et al., 2024; El et al., 2025). Notably, Barbero et al. (2024) analyse signal propagation on attention graphs, uncovering phenomena such as representational collapse and *oversquashing* (Alon & Yahav, 2021; Topping et al., 2022). These studies highlight the value of attention graphs, but are limited to descriptive and structural analyses. In contrast, we extend attention graphs to more general *attributed graphs* to integrate other computational traces and, importantly, we propose to directly *learn* on these graphs for the task of HD. To this end, we leverage Graph Neural Networks (Kipf & Welling, 2017; Gilmer et al., 2017; Battaglia et al., 2018), a family of architectures which have recently achieved remarkable results in relevant structured domains (Qasim et al., 2019; Monti et al., 2019; Stokes et al., 2020; Gonzalez et al., 2021; Liu et al., 2023).

### 3 LLM COMPUTATIONAL TRACES AS ATTRIBUTED GRAPHS

**Preliminaries.** Throughout this paper, we focus on attention-based, decoder-only LLMs. Abstracting away architectural specifics, we treat them as sharing a common backbone: a stack of transformer-decoder blocks. Let  $\mathcal{L}$  denote a reference LLM consisting of  $L$  decoder-block layers of  $H$  heads each<sup>2</sup>,  $\vec{p}$  refer to a prompt in input, and  $\vec{r}$  to the response  $\mathcal{L}$  generates. We consider  $\vec{p}, \vec{r}$  to be sequences of tokens of size, resp.,  $n_p, n_r$  ( $n := n_p + n_r$ ). We use  $T_i$  to refer to token at position  $i$  in the concatenation  $\vec{p} \mid \vec{r}$ . Within each transformer block, multi-head attention produces scores, which we collect in attention matrices  $A^{l,h} \in [0, 1]^{n \times n}$  for layer  $l$  and head  $h$ . Due to the causal structure of decoder-only transformers, these matrices are lower-triangular. For convenience, we define  $\alpha_{i,j} \in [0, 1]^{L \cdot H}$  as the vector of attention scores between  $T_i$  and  $T_j$  across all layers and heads. In addition to multi-head attention values, residual stream activations constitute another key source of information about the computation performed by  $\mathcal{L}$ . For each token  $T_i$ , we denote by  $\mathbf{a}_i^l \in \mathbb{R}^d$  its  $d$ -dimensional activation vector at layer  $l$ ; this captures the computational state of the model at such token position and processing stage. Together, attention values and activations form the primary signals we use to describe the computational traces of LLMs. While our framework focuses on these two, it can also naturally accommodate additional sources of information, such as logits.

**From computational traces to attributed graphs.** The attention values calculated along the way are, in fact, *pairwise scores* that induce a (non-symmetric) binary relation between tokens. In fact, they define a *directed graph*  $G = (V, E)$  on any sequence of tokens  $\vec{s} = \vec{p} \mid \vec{r}$ , where:

<sup>2</sup>One can also consider, without loss of generality, a different number of heads for each layer.

- 162     •  $V$ , the node (vertex) set, is the set of all tokens in  $\vec{s}$ , namely  $\{T_i\}_{i=0}^{n-1}$ ;  
 163     •  $E$ , the edge set, is the set of ordered pairs  $(T_i, T_j)$ ,  $i > j$ , signifying  $T_i$  attends to  $T_j$  in the  
 164        generation of next tokens:  $\alpha_{i,j}^{l,h} > 0$  for some of  $\mathcal{L}$ 's layers and corresponding heads.

166     We consider these graphs as *attributed*, in the sense that nodes and edges can host *features* representing  
 167         $\mathcal{L}$ 's computational traces. Edge features are given by the set of attention scores between *distinct*  
 168        tokens, i.e.,  $x_{E,(i,j)} = \alpha_{i,j}$ 's, with  $i \neq j$ . Node features are given by the attention scores "paid"  
 169        by a token to itself, i.e.,  $x_{V,i} = \alpha_{i,i}$ . Node features can also host other token-wise computational  
 170        traces; we consider residual stream activations  $\mathbf{a}_i^l$  at any layer  $l$ , so that node / token  $T_i$  is endowed  
 171        with feature vector  $x_{V,i} = (\alpha_{i,i} \mid \mathbf{a}_i^l)$ . Formally, we gather node and edge features in matrices  
 172         $X_V \in \mathbb{R}^{n \times (L \cdot H + d)}$  (or  $\mathbb{R}^{n \times (L \cdot H)}$  should activations be neglected), and  $X_E \in \mathbb{R}^{n_E \times L \cdot H}$ . The  
 173        resulting graph is  $G = (V, E, X_V, X_E)$ . This representation captures both token interactions and  
 174        per-token computational states, and can be extended to incorporate other traces, e.g., activations from  
 175        multiple layers or output logits. We leave investigating these aspects to future research endeavours.

177     **Sparsifying computational trace graphs.** Very small attention scores convey noisy and weak  
 178        contribution to updating the representation of a token. To reduce computational overhead, we  
 179        threshold attention scores at  $\tau$ , zeroing values below it and dropping edges unsupported by any head  
 180        or layer after this process. In formulae, new graph is defined as  $G = (V, E, X_V, X_E^\tau)$ , with:

$$181 \quad (X_E^\tau)_{(i,j),(l,h)} = \begin{cases} 0 & \text{if } \alpha_{i,j}^{l,h} \leq \tau, \\ \alpha_{i,j}^{l,h} & \text{otherwise.} \end{cases} \quad E = \{(T_i, T_j) \mid i > j \text{ and } \exists d \text{ s.t. } (X_E^\tau)_{(i,j),d} > 0\}. \quad (1)$$

184     As we experimentally show in Section 5.3, sparsifying the graph in this way may significantly improve  
 185        the efficiency of our yet-to-be-described model, while retaining information about the most relevant  
 186        token interactions. In the next section we illustrate how, starting from the above formalism, we can  
 187        instantiate problems such as automated HD as graph learning tasks.

## 189     4 NEURAL MESSAGE PASSING FOR HALLUCINATION DETECTION

### 191     4.1 HALLUCINATION DETECTION IS A GRAPH MACHINE LEARNING TASK

193     **Problem formulation.** Computational trace graphs can be naturally associated with *labels* one  
 194        seeks to predict for the underlying text generation process. In our specific use-case of HD, these  
 195        can indeed reflect hallucination annotations at the level of response tokens or the overall response.  
 196        Concretely, our reference graph  $G$  — encoding the computation of  $\mathcal{L}$  on prompt  $\vec{p}$  and response  $\vec{r}$  as  
 197        per the above Section 3 — can be *annotated* as  $(G, y)$ , where:

$$198 \quad \begin{aligned} \text{(i)} \quad y \in \{0, 1\}, y &= \begin{cases} 1, & \text{if } \vec{r} \text{, contains } \textit{hallucinating} \text{ passages} \\ 0, & \text{otherwise} \end{cases}, \quad \text{or} \\ 200 \quad \text{(ii)} \quad y \in \{0, 1\}^{n_r}, y_i &= \begin{cases} 1, & \text{if token } T_i, i > n_p \text{ is part of a } \textit{hallucinating} \text{ passage within } \vec{r}, \\ 0, & \text{otherwise} \end{cases}. \end{aligned}$$

203     Here, (i) stands for a *graph-wise* label, while (ii) represents labels at the granularity of single  
 204        (response) tokens. With these premises, we formalise HD as learning a parametric function  $f(G) = \hat{y}$   
 205        mapping a computational trace graph  $G$  to predictions  $\hat{y}$  of the corresponding labels  $y$ . Depending on  
 206        the task,  $\hat{y} \in [0, 1]^k$  with  $k = 1$  for graph-wise or  $k = n_r$  for token-wise detection.

207     **Our CHARM architecture.** We parameterise  $f$  in the family of message-passing Graph Neural  
 208        Network (GNN) (Kipf & Welling, 2017; Gilmer et al., 2017). These networks are, to date, the  
 209        de-facto standard for learning on attributed graphs, while possessing an architectural pattern which,  
 210        as we show next, well aligns them to generalise known approaches. Message-passing networks, in  
 211        particular, implement local computations reflecting the structure of the input graph, hence benefitting  
 212        from the aforementioned sparsification and offering a compelling advantage in terms of computational  
 213        complexity. In particular, we structure  $f$  as:  $f = f_{\text{pred}} \circ f_{\text{pool}} \circ f_{\text{mp}}$  (see Figure 2), where:

- 215     •  $f_{\text{mp}}$  stacks learnable message-passing layers to compute updated token representations from  
 the input graph  $G$ ;

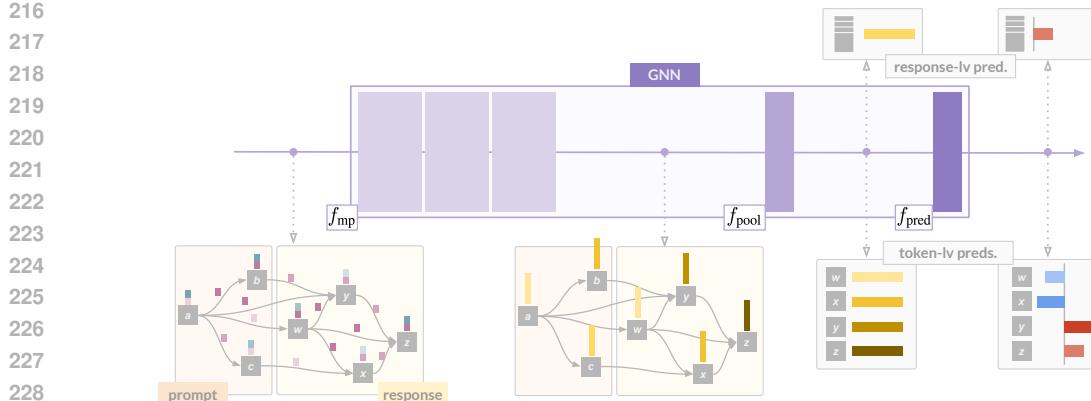


Figure 2: **HD with CHARM.** The input is an attributed graph (shown in the bottom left). First  $f_{mp}$  obtains refined node / token representations via msg-passing. Next,  $f_{pool}$  aggregates these if response level predictions are required. Finally, a projection head,  $f_{pred}$ , outputs the detection score.

- $f_{pool}$  aggregates token representations into a single graph representation (e.g., via averaging or summation), or acts as the identity for token-wise detection;
- $f_{pred}$  applies dense layers to compute graph- or token-wise hallucination predictions.

Starting from the original node (token) features  $X_V$ , each layer in  $f_{mp}$  calculates and updates hidden node representations by aggregating them (possibly, non-linearly) along the connectivity defined by the attention scores as per Equation (1). That is, the  $t$ -th layer updates token  $i$ 's embedding as:

$$h_i^{(t+1)} = \text{up}_t \left( h_i^{(t)}, \bigcup_{j: (i,j) \in E} \text{msg}_t(h_i^{(t)}, h_j^{(t)}, x_{E,(i,j)}^\tau, p_{i,j}) \right) \quad (2)$$

where:  $\bigcup$  is a permutation invariant<sup>3</sup> aggregator such as sum, average, or max;  $x_{E,(i,j)}^\tau$  corresponds to the features of edge  $(i,j)$  in  $X_E^\tau$ ;  $h_i^{(0)} = x_{V,i}$  are the initial node features;  $p_{i,j}$  is a one-hot vector indicating whether edge  $(i,j)$  connects prompt to response or response to response tokens. Functions  $\text{up}_t$ ,  $\text{msg}_t$  are parameterised as Multi-Layer Perceptrons (MLPs) running on the concatenation of their arguments. We refer to our overall approach as CHARM, a mnemonic formula for “Catching Hallucinated Responses via (learnable) Message-passing”.

This formulation offers two main advantages. First, it provides a unified framework capable of handling both token-level and response-level HD, in contrast to prior approaches that target a single granularity. Second, by leveraging message-passing over computational trace graphs, CHARM can flexibly integrate multiple signals and naturally subsume heuristic-based detectors. As we exemplify next, rather than discarding prior meaningful intuitions, our approach can generalise them.

## 4.2 EXPRESSIVENESS

Here, we demonstrate how hand-crafted heuristics emerge as special cases of our CHARM. To illustrate this, we focus on two representative methods: (i) Lookback Lens (Chuang et al., 2024), which produces tokenwise hallucination scores, and (ii) LLM-Check (Sriramanan et al., 2024), which outputs global sentence-level (graphwise) scores. We show that both heuristics can be approximated, to arbitrary precision and under mild assumptions, by CHARM, highlighting its expressiveness.

**Lookback Lens** extracts, for each *response* token  $i = n_p, \dots, n_r + n_p - 1$ , layer  $l$  and head  $h$ , a heuristic feature  $\ell_i^{l,h}$  quantifying the average proportion of attention paid to prompt w.r.t. previously

<sup>3</sup>Permutation invariance ensures that calculated message does not depend on the ordering of nodes in the neighbourhood.

270 generated response tokens (Chuang et al., 2024). Precisely, Lookback Lens outputs scores:  
 271

$$272 \quad P_i^{l,h} = \frac{1}{n_p} \sum_{j=0}^{n_p-1} \alpha_{i,j}^{l,h}, \quad R_i^{l,h} = \frac{1}{i - n_p} \sum_{j=n_p}^{i-1} \alpha_{i,j}^{l,h}; \quad \ell_i^{l,h} = \frac{P_i^{l,h}}{P_i^{l,h} + R_i^{l,h}}. \quad (3)$$

$$273$$

$$274$$

275 where  $P_i^{l,h}$ ,  $R_i^{l,h}$  correspond to the average attention paid by response tokens to, resp., the prompt  
 276 and the (previously) generated response. We argue that this heuristic can, in fact, be interpreted in the  
 277 form of message-passing on the (non-thresholded) attention graph described in Section 3, and can  
 278 thus be captured by our approach:  
 279

280 **Proposition (informal) 1.** *Equipped with a single-layer message-passing stack  $f_{mp}$ , CHARM running*  
 281 *on a non-thresholded computational trace graph ( $\tau = 0$ ) can arbitrarily well approximate token-wise*  
 282 *Lookback Lens features  $\ell_i$  for bounded prompt and response lengths.*

283 *Proof idea.* The ability to perform such approximation relies on the following: (1) aggregation over  
 284 previous tokens, as required by  $P_i^{l,h}$ 's and  $R_i^{l,h}$ 's, is naturally captured by propagating (and aggre-  
 285 gating) attention features on the neighbourhoods of the directed attention graph as per Equation (2);  
 286 (2) conditioned on mark input  $p_{i,j}$ , MLP msg can differently route attention features from prompt  
 287 vs. response tokens in *separate* subspaces of the internal representations; (3) message summation  
 288 can separately accumulate attention to prompt versus response tokens (*non-normalised*  $P_i^{l,h}$ ,  $R_i^{l,h}$ );  
 289 (4) MLP up can normalise and combine these aggregated scores as required, calculating the ratio  
 290 in Equation (3). A formal statement of informal proposition 1 is, along with its proof, in Appendix A.  
 291

292 **LLM-Check** proposes to detect hallucinations at the level of *entire responses*: for a chosen LLM  
 293 layer  $l$ , LLM-Chk- $l$  obtains an “Attention Score”  $c^l$  by averaging the log-determinants of attention  
 294 matrices across the set of corresponding heads (Sriramanan et al., 2024). Given the peculiar lower-  
 295 triangular structure of these matrices, such scores can be calculated as:  
 296

$$297 \quad c^l = \frac{1}{H} \sum_{h=0}^{H-1} \sum_{i=0}^{n-1} \log(\alpha_{i,i}^{l,h}), \quad (4)$$

$$298$$

299 i.e., by summing the log-transformed attentions paid by each token to itself, averaged across heads.  
 300 In practice, the inner summation is replaced with averaging, more robust to prompt and response  
 301 lengths. We note, in our CHARM, these “self-scores” are gathered and processed as node features,  
 302 which can be transformed and then later aggregated by our architecture to reproduce scores  $c^l$ 's.  
 303

304 **Proposition (informal) 2.** *With a single-layer message-passing stack  $f_{mp}$ , CHARM can arbitrarily*  
 305 *well approximate global LLM-Chk- $l$  features  $c^l$ , provided attentions are clipped away from zero<sup>4</sup>.*

306 *Proof idea.* Intuitively: (1) up MLPs can calculate the initial log-transform on the features of the  
 307 receiving node/token  $i$  — that is, on each token’s “self-scores” — while discarding information from  
 308 neighbours; (2)  $f_{pool}$ , set to summation — or averaging, if required — can aggregate these values  
 309 across tokens; (3) last,  $f_{pred}$  can average these values across heads and selectively for the desired  
 310 layer  $l$  to implement the outer averaging in Equation (4). Again, a formal statement of informal  
 311 proposition 2 is reported and proved in Appendix A.

312 Both the two above propositions guarantee that, while general and learnable, our approach can  
 313 also provably default to known, hand-crafted heuristics (under mild, reasonable assumptions). This  
 314 showcases the expressiveness of the CHARM framework, further evidencing how graph representations  
 315 and message-passing networks can offer a valid and compelling perspective into the task of HD.  
 316

## 317 5 EXPERIMENTS

$$318$$

319 We evaluate different aspects of learning with CHARM through the following research questions: **(Q1)**  
 320 Is our GNN-based formulation effective in practice? **(Q2)** Is CHARM effective at detecting different  
 321 types of hallucinations and across different granularities (e.g., token-level and full-response)? **(Q3)**  
 322

323 <sup>4</sup>This assumption ensures the log is continuous on an appropriate compact set, rendering its approximation  
 amenable (see Appendix A); in practice we also did observe it was necessary to ensure numerical stability.

324 Does CHARM exhibit any zero-shot transferability across datasets? **(Q4)** How crucial is the underlying  
 325 graph structure for CHARM’s performance? **(Q5)** Can CHARM handle large/dense graphs? **(Q6)** Can  
 326 the combination of attention and activations be effective in CHARM?

327  
 328 We initialise training of CHARM with 3 different random seeds, and, in the following, report the  
 329 mean test performance along with std. All results are obtained by models optimising validation  
 330 performance (AUPR, see below). Additional information, including dataset details, hyperparameter  
 331 searches, implementation notes, and extended results, are available in Appendices B to D.

### 332 5.1 CONTEXTUAL TOKEN-LEVEL HALLUCINATION DETECTION

333  
 334 **Datasets.** We first evaluate our approach at a token-level granularity on the NQ (Kwiatkowski  
 335 et al., 2019) and CNN (See et al., 2017) datasets. These consists of prompt-response pairs with  
 336 hallucination annotations available at the level of single response tokens (see Section 4.1). These pairs  
 337 are obtained by prompting a target LLM to perform either document-based question answering (NQ)  
 338 or text summarisation (CNN). These datasets contain instances of *contextual hallucinations*: although  
 339 the relevant and correct facts are provided in the input context, the target LLM is still observed to  
 340 generate incorrect responses (Chuang et al., 2024). Original generations and annotations for this  
 341 dataset are derived from (Chuang et al., 2024)<sup>5</sup>; coherently with the setup in the same work, we take  
 342 LLaMa-2-7B-chat as the reference LLM on both datasets. More details are in Appendix B.1.

343  
 344 **Method comparisons.** We compare against a set of representative  
 345 baselines. *Probability-based detectors* (Probas) (Guerreiro et al.,  
 346 2022; Kadavath et al., 2022; Varshney et al., 2023; Huang et al., 2023b)  
 347 leverage the next-token probabilities to estimate LLM uncertainty  
 348 and predict hallucinations; *Activation probes* (Orgad et al., 2024;  
 349 Azaria & Mitchell, 2023; Belinkov, 2022) (Act-\*) train a logistic classifier  
 350 on *activations* at specific layers; the *attention-based*, Lookback  
 351 Lens heuristic (Chuang et al., 2024) fits the same model on hand-  
 352 crafted token-wise features calculated over all layers and heads (see  
 353 Equation (3)). We run Act-\*'s on the common choice of  
 354 LLM layers 24, 28, 32, motivated by the findings in (Chuang et al., 2024; Azaria & Mitchell, 2023).  
 355 As for CHARM, we instantiate it in two configurations: one only employing attention features (att),  
 356 another also utilising activations from a specific layer, which we set to 24 (att+act-24) due to its  
 357 consistently superior performance in Act-\*'s. We run CHARM on graphs sparsified with a fixed  
 358  $\tau = 0.05$  (see Equation (1) and related ablations in Section 5.3). The HD task is at the level of single  
 359 nodes/tokens, so  $f_{pool}$  is set to *identity*. We additionally consider two ablated versions of CHARM  
 360 (att): Neigh-Avg (N), Neigh-Avg (E). These extract token-wise features through a single, non-  
 361 learnable msg-passing step, aggregating, resp., either node or edge features across neighbourhoods in  
 362 the same computational trace graphs considered in CHARM (details are in Appendix D.1.1). Comparing  
 363 with these allows us to evaluate the relevance of our multi-layer, learnable procedure. *All methods*  
 364 in comparison have their hyperparameters tuned on the Val. set. Performance is measured in terms of  
 365 Test AUROC and AUPR.

366  
 367 **Results,** reported in Table 1, show our approach consistently outperforms all baselines across  
 368 both datasets and metrics in the att.-only configuration. The significant margins over Lookback  
 369 Lens, Neigh-Avg (N) and Neigh-Avg (E) underscore the benefits of an expressive, learnable  
 370 graph-based method over attention-based heuristics. This pure attention configuration also surpasses

371 Table 1: Test AUROC and AUPR (%) for NQ and CNN (token-  
 372 wise, higher is better). **Bold**: best, Underlined: runner-up. <sup>†</sup>:  
 373 we also tune the regular. strength, diff. than the original.

Method	NQ		CNN	
	AUROC	AUPR	AUROC	AUPR
Probas	49.8	16.2	54.4	8.2
Act-24	<u>73.0</u>	<u>36.2</u>	71.3	<u>20.3</u>
Act-28	71.6	34.6	70.1	18.4
Act-32	67.4	28.6	67.7	15.4
Lookback Lens	70.8	31.0	71.9	17.4
Lookback Lens <sup>†</sup>	71.9	34.3	<u>74.4</u>	19.7
Neigh-Avg (N)	66.0	24.5	70.1	14.9
Neigh-Avg (E)	66.8	30.4	70.5	18.6
ours	<b>CHARM (att)</b>	<b><u>74.8</u><sub>±0.6</sub></b>	<b><u>40.3</u><sub>±1.7</sub></b>	<b><u>75.4</u><sub>±0.2</sub></b>
	CHARM (att+act-24)	72.2 <sub>±1.2</sub>	35.5 <sub>±1.6</sub>	70.9 <sub>±0.2</sub>
				19.8 <sub>±0.5</sub>

374  
 375 <sup>5</sup>Differently than (Chuang et al., 2024), however, we construct and experiment with a different split ensuring  
 376 full textual disjointness between train, test, and validation, see more in Appendix B.1.

378 *all* activation-based Act- $^*$  probes, this contributing to answer positively to **Q1**. We interestingly  
 379 report that including activations from an intermediate layer into CHARM reveals detrimental. We  
 380 hypothesise that, on these contextual benchmarks, attention features alone carry most of the relevant  
 381 predictive signal — an expressive enough model like ours can leverage it at best and struggle to find  
 382 additional complementary signals on activations, which could, instead lead to fit spurious correlations.  
 383

## 384 5.2 RESPONSE-LEVEL HALLUCINATION DETECTION

385 **Datasets.** We next evaluate CHARM at the coarser response-level HD on three datasets: Movies (Orgad et al., 2024), WinoBias (Zhao et al., 2018), and Math (Sun et al., 2024). Unlike NQ and CNN,  
 386 these benchmarks address failure modes different than contextual grounding, namely: factual knowl-  
 387 edge recall (Movies), intrinsic bias in coreference resolution (WinoBias), and arithmetic reasoning  
 388 (Math). This allows to assess generalisation across fundamentally different hallucination types. The  
 389 relative role of attention thereon is not obviously clear, but we hypothesise they can provide infor-  
 390 mative signals, e.g. by capturing systematic biases in attention to demographic cues or by reflecting  
 391 unusual patterns in intermediate calculations steps. Exploring their interplay with other computational  
 392 traces is thus a insightful direction. For these experiments we derive text generations and hallucination  
 393 annotations following the procedure in (Orgad et al., 2024) and, consistently with this work, we target  
 394 a different LLM, Mistral-7B-instruct. More dataset details are in Appendix B.2.  
 395

396 **Method comparisons.** As for CHARM and its non-learnable counterparts (iv), we set component  
 397  $f_{pool}$  to *average*. We compare our method to Act- $^*$ ’s probing activations in notoriously relevant  
 398 token positions, e.g., the last token of the prompt, or the last of the response (Orgad et al., 2024)  
 399 (see Appendix C.4). Here, we compare against the response-level attention-based LLM-Check (Sri-  
 400 ramanan et al., 2024) (LLM-Chk- $^*$ ) and the spectral-method proposed in (Binkowski et al., 2025)  
 401 (LapEig). We also run an enhanced counterpart of (LLM-Chk++- $^*$ ) whereby per-head scores are  
 402 considered as inputs to logistic regression, rather than being averaged (Equation (4)).  
 403

404 **Results** are re-  
 405 ported in Table 2.  
 406 Overall, CHARM  
 407 attains the best per-  
 408 formance across these  
 409 benchmarks, with  
 410 particularly notable  
 411 gaps in Math. To-  
 412 gether with the above,  
 413 these results answer  
 414 positively to **Q1** and  
 415 **Q2**. Interestingly, we  
 416 observe a markedly  
 417 different behaviour  
 418 than in the previously  
 419 considered context-  
 420 ual HD datasets.  
 421 Other than Movies  
 422 — where both our  
 423 configurations work  
 424 equally well — on both Winobias and Math, activation-based features work in synergy with  
 425 attention-based ones. Instead of leading to fit spurious correlations as *hypothesised* for NQ and CNN  
 426 (Section 5.1), they contribute to strongly boost performance over the att.-only CHARM, and over  
 427 all considered att.- and act.-based methods. This is particularly pronounced on Winobias — there,  
 428 CHARM (att.) is surpassed by Act- $^*$ ’s, but the inclusion of activations leads it (att. & act.-24) to  
 429 significantly outperform them both. Overall these datasets provide cases leading to a positive answer  
 430 to **Q6**. We last note that Math is the only dataset where Act – 24 outperformed by other variants,  
 431 namely Act-32. We thus also ran CHARM (att+act-32), which scored Test AUROC of  $81.7 \pm 0.2$ ,  
 and AUPR of  $83.8 \pm 0.3$ .

Table 2: Test AUROC and AUPR (%) for Movies, Winobias, Math (response-  
 lv, higher is better). **Bold**: best, Underlined: runner-up.

Method	Movies		Winobias		Math	
	AUROC	AUPR	AUROC	AUPR	AUROC	AUPR
Probas	58.6	81.6	64.5	20.0	54.5	57.4
Act-24	77.0	90.4	<u>76.6</u>	<u>37.8</u>	77.7	77.5
Act-28	77.0	90.4	73.9	34.3	<u>78.1</u>	77.8
Act-32	76.3	90.2	72.7	35.3	76.6	77.9
LLM-Chk-24	47.5	74.6	38.9	10.9	64.5	68.2
LLM-Chk-28	51.1	76.7	41.6	11.4	65.5	69.2
LLM-Chk-32	61.5	82.1	41.6	11.3	64.0	67.6
LLM-Chk++-24	66.3	84.5	64.6	20.5	67.3	69.1
LLM-Chk++-28	67.8	86.3	64.8	21.0	67.0	70.6
LLM-Chk++-32	73.0	88.8	67.2	24.1	68.6	72.5
LapEig	72.9	88.4	74.1	33.3	73.6	76.3
Neigh-Avg (N)	78.6	91.2	63.8	23.0	77.4	79.2
Neigh-Avg (E)	54.9	78.5	65.8	21.9	76.7	78.3
ours	CHARM (att)	<b><u>80.3</u></b> $\pm$ <b><u>0.2</u></b>	<b><u>92.0</u></b> $\pm$ <b><u>0.1</u></b>	70.4 $\pm$ 0.7	29.1 $\pm$ 1.0	76.5 $\pm$ 1.1
ours	CHARM (att+act-24)	79.7 $\pm$ 0.3	91.8 $\pm$ 0.2	<b><u>77.8</u></b> $\pm$ <b><u>0.4</u></b>	<b><u>39.8</u></b> $\pm$ <b><u>1.3</u></b>	<b><u>80.8</u></b> $\pm$ <b><u>0.7</u></b>
ours						<b><u>79.7</u></b> $\pm$ <b><u>0.5</u></b>
ours						<b><u>83.1</u></b> $\pm$ <b><u>0.7</u></b>

432 5.3 ADDITIONAL ANALYSES  
433

434 **The role of the graph.** To what extent  
435 does message-passing on the attention-  
436 induced graph contribute to the perfor-  
437 mance of CHARM? To answer this, we ab-  
438 late the connectivity in our input samples  
439 and train CHARM on two representative  
440 datasets: CNN and Math. In this setup, CHARM  
441 defaults to a set model which, instead of message-  
442 passing, applies a stack of dense layers over node features. We train and extensively tune this baseline,  
443 denoted “CHARM (no g.)”, considering both att.-only and att. & act. configurations. We report its  
444 best results in Table 3 compared to the best corresponding CHARM. Results clearly show the positive  
445 contribution of message-passing on the constructed topology, answering positively to **Q4**.  
446

447 **Efficiency and robustness.** We exper-  
448 iment with different values of the attention  
449 threshold  $\tau$  (Equation (1)), studying how  
450 graph sparsity and (inference) memory con-  
451 sumption vary in relation to performance.  
452 We run this study on NQ, with results in Ta-  
453 ble 4. Test AUPR is reported along with  
454 the number of edges, sparsity and infer-  
455 ence memory footprint averaged over test  
456 graphs. We observe CHARM’s performance is robust to various levels of sparsifications, whilst this  
457 can provide dramatic reduction in resource consumption. Performance drops more notably only for  
458  $\tau = 0.5$ , which, we note, still outperforms the best competitor, i.e., Act-24 (see Table 1). Overall  
459 our default  $\tau = 0.05$  attains a good trade-off, whilst we note it maximises val. AUPR. These results  
460 answers positively to **Q5**. We finally measure a distinctly contained inference latency of  $\approx 1e^{-3}$  secs.  
461 Refer to Appendix C.1 for run-time and performance comparisons with other popular HD methods  
462 relying on multiple prompting, which incur significantly higher latency.  
463

464 **Zero-shot transfer.** CHARM is a learn-  
465 able, expressive multi-layer approach —  
466 this raises a natural question: *To what ex-  
467 tent can it generalise cross-datasets zero-  
468 shot?* To investigate this, we follow the  
469 setup in (Chuang et al., 2024): we train on  
470 NQ and evaluate on CNN, and vice-versa.  
471 Results are in in Table 5. Overall *no single method consistently outperforms the others* in this  
472 challenging setup. In fact, despite its larger expressiveness, CHARM demonstrates promising generali-  
473 sation: it outperforms activation-based probes, ranks best in CNN  $\rightarrow$  NQ, and places second in NQ  
474  $\rightarrow$  CNN (behind Lookback Lens, which conversely performs the *worst* in CNN  $\rightarrow$  NQ). These  
475 results suggest that zero-shot transfer remains an open challenge, but our graph-based formulation is  
476 competitive and can capture generalisable signals, answering positively to **Q3**.  
477

478 6 CONCLUSIONS  
479

480 In this work, we proposed attributed graphs as a principled formulation of LLM computational traces,  
481 showing how diverse signals can be unified in this framework and how neural message passing can  
482 be applied thereon for diverse HD tasks. We showed our approach, CHARM, can provably generalise  
483 prior methods and that it achieves strong empirical performance, consistently outperforming existing  
484 methods. Additional analyses underscored the importance of graph structure and demonstrated  
485 promising zero-shot generalization across datasets.  
486

487 Future endeavours will consider integrating other computational traces (e.g., logits), as well as  
488 extensions to new tasks such as detecting data contamination, identifying LLM-generated text, or  
489 flagging jailbreak attempts. Future work may also explore alternative message-passing architectures,  
490 including positional and structural encodings tailored to these attributed graphs.  
491

492 Table 3: Results from ablating the graph structure.  
493

Method	CNN		Math	
	AUROC	AUPR	AUROC	AUPR
CHARM (no g.)	70.8 $\pm$ 0.5	19.2 $\pm$ 0.5	80.6 $\pm$ 0.7	82.7 $\pm$ 0.1
CHARM	<b>75.4</b> $\pm$ 0.2	<b>22.7</b> $\pm$ 0.4	<b>81.7</b> $\pm$ 0.2	<b>83.8</b> $\pm$ 0.3

494 Table 4: Avg. graph stats. on NQ at different thresholds,  
495 along with GPU memory footprint and Test AUPR.  
496

$\tau$	Num. Edges	Sparsity	Mem. (MB)	AUPR
0.5	1,118.80	0.993	22.99 $\pm$ 3.71	38.4 $\pm$ 0.4
0.1	7,458.67	0.952	60.44 $\pm$ 11.79	41.0 $\pm$ 1.2
0.05	14,884.44	0.906	104.15 $\pm$ 23.05	40.3 $\pm$ 1.7
0.01	58,998.88	0.645	363.02 $\pm$ 98.39	40.3 $\pm$ 0.9
0.001	19,7784.82	0.026	1177.20 $\pm$ 523.61	40.1 $\pm$ 0.0

497 Table 5: Cross-dataset zero-shot transf. NQ  $\leftrightarrow$  CNN.  
498

Method	NQ $\rightarrow$ CNN		CNN $\rightarrow$ NQ	
	AUROC	AUPR	AUROC	AUPR
LkbLens <sup>†</sup>	<b>68.6</b>	<b>14.9</b>	62.0	26.5
Act-24	63.4	11.3	<b>63.8</b>	<b>29.7</b>
CHARM	<b>64.1</b> $\pm$ 1.1	<b>12.0</b> $\pm$ 0.98	<b>65.5</b> $\pm$ 0.14	<b>31.6</b> $\pm$ 0.10

486 REPRODUCIBILITY STATEMENT  
487488 Our code will be released upon acceptance, along with all training and evaluation scripts. Section 5,  
489 as well as Appendices B to D provide the required implementation details to reproduce our results.  
490491 REFERENCES  
492493 Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications.  
494 *International Conference on Learning Representations*, 2021.495 Amos Azaria and Tom Mitchell. The internal state of an llm knows when it's lying. In *Findings of*  
496 *the Association for Computational Linguistics: EMNLP 2023*, pp. 967–976, 2023.497 Guy Bar-Shalom, Fabrizio Frasca, Derek Lim, Yoav Gelberg, Yftah Ziser, Ran El-Yaniv, Gal  
498 Chechik, and Haggai Maron. Learning on llm output signatures for gray-box behavior analysis.  
499 *arXiv:2503.14043*, 2025.500 Federico Barbero, Andrea Banino, Steven Kapturowski, Dharshan Kumaran, João G.M. Araújo, Alex  
501 Vitvitskyi, Razvan Pascanu, and Petar Veličković. Transformers need glasses! information over-  
502 squashing in language tasks. In *Advances in Neural Information Processing Systems*, volume 37,  
503 pp. 98111–98142, 2024.504 Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi,  
505 Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al.  
506 Relational inductive biases, deep learning, and graph networks. *arXiv:1806.01261*, 2018.507 Alexandra Bazarova, Aleksandr Yugay, Andrey Shulga, Alina Ermilova, Andrei Volodichev, Kon-  
508 stantin Polev, Julia Belikova, Rauf Parchiev, Dmitry Simakov, Maxim Savchenko, Andrey  
509 Savchenko, Serguei Barannikov, and Alexey Zaytsev. Hallucination detection in llms via topologi-  
510 cal divergence on attention graphs. *arXiv:2504.10063*, 2025.511 Yonatan Belinkov. Probing classifiers: Promises, shortcomings, and advances. *Computational*  
512 *Linguistics*, 48(1):207–219, 2022.513 Lukas Biewald. Experiment tracking with weights and biases, 2020. Software available from  
514 wandb.com.515 Jakub Binkowski, Denis Janiak, Albert Sawczyn, Bogdan Gabrys, and Tomasz Kajdanowicz. Hallu-  
516 cination detection in llms using spectral features of attention maps. *arXiv:2502.17598*, 2025.517 Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in language  
518 models without supervision. *arXiv:2212.03827*, 2022.519 Meng Cao, Yue Dong, and Jackie Cheung. Hallucinated but factual! inspecting the factuality of  
520 hallucinations in abstractive summarization. In Smaranda Muresan, Preslav Nakov, and Aline  
521 Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational*  
522 *Linguistics (Volume 1: Long Papers)*, pp. 3340–3354, Dublin, Ireland, May 2022. Association for  
523 Computational Linguistics.524 Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. Inside:  
525 Llms' internal states retain the power of hallucination detection. *arXiv:2402.03744*, 2024.526 Yung-Sung Chuang, Linlu Qiu, Cheng-Yu Hsieh, Ranjay Krishna, Yoon Kim, and James R. Glass.  
527 Lookback lens: Detecting and mitigating contextual hallucinations in large language models using  
528 only attention maps. In *Proceedings of the 2024 Conference on Empirical Methods in Natural*  
529 *Language Processing*, pp. 1419–1436. Association for Computational Linguistics, 2024.530 Batu El, Deepro Choudhury, Pietro Liò, and Chaitanya K. Joshi. Towards mechanistic interpretability  
531 of graph transformers via attention graphs. In *ICLR 2025 Workshop on Explainable AI for Science*  
532 (*XAI4Science*), 2025.533 Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with PyTorch Geometric. In  
534 *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

- 540 Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural  
 541 message passing for quantum chemistry. In *International Conference on Machine Learning*,  
 542 volume 70 of *Proceedings of Machine Learning Research*, pp. 1263–1272. PMLR, 2017.
- 543
- 544 Guadalupe Gonzalez, Shunwang Gong, Ivan Laponogov, Michael M. Bronstein, and Kirill Veselkov.  
 545 Predicting anticancer hyperfoods with graph convolutional networks. *Human Genomics*, 15(33),  
 546 2021.
- 547
- 548 Nuno M Guerreiro, Elena Voita, and André FT Martins. Looking for a needle in a haystack: A  
 549 comprehensive study of hallucinations in neural machine translation. *arXiv:2208.05309*, 2022.
- 550
- 551 Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong  
 552 Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language  
 553 models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information  
 Systems*, 2023a.
- 554
- 555 Yuheng Huang, Jiayang Song, Zhijie Wang, Shengming Zhao, Huaming Chen, Felix Juefei-Xu,  
 556 and Lei Ma. Look before you leap: An exploratory study of uncertainty measurement for large  
 557 language models. *arXiv:2307.10236*, 2023b.
- 558
- 559 Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang,  
 560 Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM  
 Computing Surveys*, 55(12):1–38, 2023.
- 561
- 562 Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot,  
 563 Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al.  
 564 Mistral 7b. *arXiv:2310.06825*, 2023.
- 565
- 566 Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas  
 567 Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly)  
 568 know what they know. *arXiv:2207.05221*, 2022.
- 569
- 570 Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks.  
 In *International Conference on Learning Representations (ICLR)*, 2017.
- 571
- 572 Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for  
 573 uncertainty estimation in natural language generation. *arXiv:2302.09664*, 2023.
- 574
- 575 Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris  
 576 Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion  
 577 Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav  
 578 Petrov. Natural questions: A benchmark for question answering research. *Transactions of the  
 Association for Computational Linguistics*, 7:452–466, 2019.
- 579
- 580 Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time  
 581 intervention: Eliciting truthful answers from a language model. *Advances in Neural Information  
 Processing Systems*, 36, 2024.
- 582
- 583 Gary Liu, Denise Catacutan, Khushi Rathod, Kyle Swanson, Wengong Jin, Jody Mohammed, Anush  
 584 Chiappino-Pepe, Saad Syed, Meghan Fragis, Kenneth Rachwalski, Jakob Magolan, Michael Surette,  
 585 Brian Coombes, Tommi Jaakkola, Regina Barzilay, James Collins, and Jonathan Stokes. Deep  
 586 learning-guided discovery of an antibiotic targeting *acinetobacter baumannii*. *Nature Chemical  
 Biology*, pp. 1–9, 2023.
- 587
- 588 Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and Bill Dolan.  
 589 A token-level reference-free hallucination detection benchmark for free-form text generation.  
 590 *arXiv:2104.08704*, 2021.
- 591
- 592 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv:1711.05101*, 2017.
- 593
- Potsawee Manakul, Adian Liusie, and Mark JF Gales. Selfcheckgpt: Zero-resource black-box  
 594 hallucination detection for generative large language models. *arXiv:2303.08896*, 2023.

- 594 Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language  
 595 model representations of true/false datasets. *arXiv:2310.06824*, 2023.  
 596
- 597 Federico Monti, Fabrizio Frasca, Davide Eynard, Damon Mannion, and Michael M. Bronstein.  
 598 Fake news detection on social media using geometric deep learning. In *ICLR 2019 Workshop on*  
 599 *Representation Learning on Graphs and Manifolds*, 2019.
- 600 Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan  
 601 Belinkov. Llms know more than they show: On the intrinsic representation of llm hallucinations.  
 602 *arXiv:2410.02707*, 2024.  
 603
- 604 Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. Understanding factuality in abstrac-  
 605 tive summarization with FRANK: A benchmark for factuality metrics. In Kristina Toutanova,  
 606 Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cot-  
 607 terell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the*  
 608 *North American Chapter of the Association for Computational Linguistics: Human Language*  
 609 *Technologies*, pp. 4812–4829, Online, June 2021. Association for Computational Linguistics.
- 610 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor  
 611 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward  
 612 Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,  
 613 Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance  
 614 deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and  
 615 R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran  
 616 Associates, Inc., 2019.  
 617
- 618 Allan Pinkus. Approximation theory of the mlp model in neural networks. *Acta Numerica*, 8:143–195,  
 619 1999.
- 620 Shah Rukh Qasim, Jan Kieseler, Yutaro Iiyama, and Maurizio Pierini. Learning representations  
 621 of irregular particle-detector geometry with distance-weighted graph networks. *The European*  
 622 *Physical Journal C*, 79(7), 2019.
- 623 Yifu Qiu, Yftah Ziser, Anna Korhonen, Edoardo Ponti, and Shay Cohen. Detecting and mitigating  
 624 hallucinations in multilingual summarisation. In Houda Bouamor, Juan Pino, and Kalika Bali  
 625 (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*,  
 626 pp. 8914–8932, Singapore, December 2023. Association for Computational Linguistics.
- 627 Miriam Rateike, Celia Cintas, John Wamburu, Tanya Akumu, and Skyler Speakman. Weakly  
 628 supervised detection of hallucinations in llm activations. *arXiv:2312.02798*, 2023.
- 629 Vipula Rawte, Swagata Chakraborty, Agnibh Pathak, Anubhav Sarkar, SM Tommoy, Aman Chadha,  
 630 Amit P Sheth, and Amitava Das. The troubling emergence of hallucination in large language  
 631 models—an extensive definition, quantification, and prescriptive remediations. *arXiv:2310.04988*,  
 632 2023.
- 633 Abigail See, Peter J. Liu, and Christopher D. Manning. Get to the point: Summarization with  
 634 pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for*  
 635 *Computational Linguistics (Volume 1: Long Papers)*, pp. 1073–1083, Vancouver, Canada, 2017.
- 636 Adi Simhi, Jonathan Herzig, Idan Szpektor, and Yonatan Belinkov. Constructing benchmarks and  
 637 interventions for combating hallucinations in llms. *arXiv:2404.09971*, 2024.
- 638 Ben Snyder, Marius Moisescu, and Muhammad Bilal Zafar. On early detection of hallucinations in  
 639 factual question answering. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge*  
 640 *Discovery and Data Mining*, pp. 2721–2732, 2024.
- 641 Gaurang Sriramanan, Siddhant Bharti, Vinu Sankar Sadasivan, Shoumik Saha, Priyatham Kattakinda,  
 642 and Soheil Feizi. Llm-check: Investigating detection of hallucinations in large language models.  
 643 In *Advances in Neural Information Processing Systems*, volume 37, pp. 34188–34216, 2024.

- 648 Jonathan M. Stokes, Kevin Yang, Kyle Swanson, Wengong Jin, Andres Cubillos-Ruiz, Nina M.  
 649 Donghia, Craig R. MacNair, Shawn French, Lindsey A. Carfrae, Zohar Bloom-Ackermann,  
 650 Victoria M. Tran, Anush Chiappino-Pepe, Ahmed H. Badran, Ian W. Andrews, Emma J. Chory,  
 651 George M. Church, Eric D. Brown, Tommi S. Jaakkola, Regina Barzilay, and James J. Collins. A  
 652 deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702, 2020.
- 653 Yuhong Sun, Zhangyue Yin, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Hui Zhao. Benchmarking hal-  
 654 lucination in large language models based on unanswerable math word problem. *arXiv:2403.03558*,  
 655 2024.
- 656 Jake Topping, Francesco Di Giovanni, Benjamin P. Chamberlain, Xiaowen Dong, and Michael M.  
 657 Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature. In *International  
 658 Conference on Learning Representations*, 2022.
- 659 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay  
 660 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian  
 661 Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin  
 662 Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar  
 663 Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann,  
 664 Artem Korenev, Punit Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich,  
 665 Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin  
 666 Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan  
 667 Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor,  
 668 Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela  
 669 Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and  
 670 Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv:2307.09288*, 2023.
- 671 Neeraj Varshney, Wenlin Yao, Hongming Zhang, Jianshu Chen, and Dong Yu. A stitch in time saves  
 672 nine: Detecting and mitigating hallucinations of llms by validating low-confidence generation.  
 673 *arXiv:2307.03987*, 2023.
- 674 Alex Vitvitskyi, João G. M. Araújo, Marc Lackenby, and Petar Veličković. What makes a good  
 675 feedforward computational graph? In *Proceedings of the 42nd International Conference on  
 676 Machine Learning*, Proceedings of Machine Learning Research. PMLR, 2025.
- 677 Weijia Xu, Sweta Agrawal, Eleftheria Briakou, Marianna J. Martindale, and Marine Carpuat. Un-  
 678 derstanding and detecting hallucinations in neural machine translation via model introspection.  
*Transactions of the Association for Computational Linguistics*, 11:546–564, 2023.
- 679 Fan Yin, Jayanth Srinivasa, and Kai-Wei Chang. Characterizing truthfulness in large language model  
 680 generations with local intrinsic dimension. *arXiv:2402.18048*, 2024.
- 681 Mert Yuksekgonul, Varun Chandrasekaran, Erik Jones, Suriya Gunasekar, Ranjita Naik, Hamid  
 682 Palangi, Ece Kamar, and Besmira Nushi. Attention satisfies: A constraint-satisfaction lens on  
 683 factual errors of language models. *arXiv:2309.15098*, 2023.
- 684 Tianhang Zhang, Lin Qiu, Qipeng Guo, Cheng Deng, Yue Zhang, Zheng Zhang, Chenghu Zhou,  
 685 Xinbing Wang, and Luoyi Fu. Enhancing uncertainty-based hallucination detection with stronger  
 686 focus. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language  
 687 Processing*, pp. 915–932. Association for Computational Linguistics, 2023.
- 688 Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Gender bias in  
 689 coreference resolution: Evaluation and debiasing methods. In *Proceedings of the 2018 Conference  
 690 of the North American Chapter of the Association for Computational Linguistics: Human Language  
 691 Technologies, Volume 2 (Short Papers)*, pp. 15–20, New Orleans, Louisiana, 2018. Association for  
 692 Computational Linguistics.
- 693 Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan,  
 694 Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, et al. Representation engineering: A  
 695 top-down approach to ai transparency. *arXiv:2310.01405*, 2023.
- 696
- 697
- 698
- 699
- 700
- 701

## 702 A EXPRESSIVENESS: CLAIMS AND PROOFS 703

704 **Proposition (informal) 1.** *Equipped with a single-layer message-passing stack  $f_{mp}$ , CHARM running  
705 on a non-thresholded computational trace graph ( $\tau = 0$ ) can arbitrarily well approximate token-wise  
706 Lookback Lens features  $\ell_i$  for bounded prompt and response lengths.*  
707

708 **Proposition 1.** *Let  $\ell_i(\vec{s}; \mathcal{L})$  denote the Lookback Lens features calculated, for token  $i$  on string  
709  $\vec{s} = \vec{p} \mid \vec{r}$  for LLM  $\mathcal{L}$  (Equation (3)). Also, let  $h_i^{(t)}(G_{\vec{s}})$  denote the  $t$ -layer representation for the same  
710 token in output from  $t$  CHARM msg-passing layers (Equation (2)) on the corresponding computational  
711 trace graph  $G_{\vec{s}}$ . For any precision  $\epsilon > 0$ , there exists a 1-layer stack of CHARM’s layers which  
712 approximate Lookback Lens features up to precision  $\epsilon$ , for maximum allowed prompt-length  $\bar{n}_p$  and  
713 response length  $\bar{n}_r$ . That is, for any prompt  $\vec{p}$ , response  $\vec{r}$  of the aforementioned maximum lengths,  
714 and for any response token  $i$ ,  $\|h_i^{(1)}(G_{\vec{p} \mid \vec{r}}) - \ell_i(\vec{p} \mid \vec{r}; \mathcal{L})\|_{\infty} < \epsilon$ .*  
715

716 *Proof.* To prove the above we show that, by setting hyperparameters  $\square \equiv \sum$  (summation) and  $\tau = 0$   
717 (no thresholding), there exist a single-layer stack  $f_{mp}$  which compute the desired approximation. In  
718 the following we will consider the attention scores across the  $L$  layers and  $H$  heads to be arranged  
719 in our node and edge features in a flattened manner. We will conveniently denote with  $b(l, h)$  the  
720 function which evaluates the index of layer  $l$  and head  $h$  in this flattened representation.  
721

722 **(0) Setup and inputs.** As our stack is made up of one single message-passing layer, our specific  
723 interest is thus in showing the existence of appropriate MLPs  $\text{msg}_0, \text{up}_0$  enabling  $f_{mp}$  to realise the  
724 target approximation. Being these components of the *first* — and only — message-passing layer in  
725  $f_{mp}$ , its input node and edge representations effectively correspond to the original node and edge  
726 features  $X_V, X_E$ . We are thus focussing on the following update:  
727

$$728 h_i^{(1)} = \underbrace{\text{up}_0}_{(2)} \left( x_{V,i}, \underbrace{\sum_{j < i} \text{msg}_0}_{(1)} (x_{V,i}, x_{V,j}, x_{E,(i,j)}, p_{i,j}) \right) \quad (5)$$

730 where the neighbourhood aggregation sums across *all* previous token positions ( $j < i$ ) since no  
731 thresholding is enforced ( $\tau = 0$ ) and due to the fact that attention values cannot exactly evaluate to 0  
732 because of the application of softmax normalisation.  
733

734 **(1) Message function.** Let us first describe what we desire  $\text{msg}_0$  to calculate. We would like it to  
735 map the concatenation of its arguments, with dimensionality  $3d + 2$ ,  $d = L \cdot H$ , to a vector of  
736 dimensionality  $2d + 2$ , where:  
737

- 738 • The mark feature  $p_{i,j}$  is replicated in the first two dimensions (channels 0, 1);  
739
- 740 • Edge features  $x_{E,(i,j)}$  are replicated either in channels 2 through  $2 + d - 1$  if  $p_{i,j} = [1, 0]$   
741 (message from prompt token) or in channels  $2 + d$  through  $2d - 1$  otherwise (message from  
742 response token);  
743
- 744 • Node features  $x_{V,i}, x_{V,j}$  are discarded.  
745

746 Now, an MLP exactly implementing the above message function does exist; in fact, it can be *explicitly  
747 constructed*.  
748

749 *First layer:* Weight matrix  $W_1$  is in  $\mathbb{R}^{(3d+2) \times (2d+2)}$ . We will describe it in terms of columns slices.  
750

- 751 • A first slice gathers the first two columns (0, 1); these are all zero except for the bottom  $2 \times 2$   
752 block, set as identity  $I_2$ . This slice copies the  $p_{i,j}$  mark features in the first two channels of  
753 the hidden representation.  
754
- 755 • A second slice gathers columns 2 through  $2 + d - 1$ ; these are all zero except for rows  
756  $2d$  through  $3d - 1$ , set to identity  $I_d$ , and row  $3d$ , set to a  $\vec{1}_d$  one-only vector. This slice  
757 calculates the sum between edge features  $x_{E,(i,j)}$  and the first channel of the mark  $p_{i,j}$ ,  
758 indicating whether the message comes from a prompt token.  
759

- 756 • A third — and last — slice gathers columns  $2 + d$  through  $2 + 2d - 1$ ; these are all zero  
 757 except for rows  $2d$  through  $3d - 1$ , set to identity  $I_d$ . This slice copies edge features  $x_{E,(i,j)}$   
 758 in the last  $d$  channels of the hidden representation.

760 Bias vector  $b_1$  in  $\mathbb{R}^{2d+2}$  is all zero except for channels  $2$  through  $2 + d - 1$ , set to vector  $-\vec{1}_d$ .  
 761 Recapping, the hidden representation is a vector in  $\mathbb{R}^{2d+2}$  with the following structure:

$$762 \left[ p_{i,j} \mid \underbrace{\cdots (x_{E,(i,j)})_{b(l,h)} + (p_{i,j})_0 - 1 \cdots}_{(h_1)} \mid \underbrace{\cdots (x_{E,(i,j)})_{b(l,h)} \cdots}_{(h_2)} \right].$$

766 *Second layer:* Weight matrix  $W_2$  is in  $\mathbb{R}^{(2d+2) \times (2d+2)}$ . We will describe it in terms of columns slices  
 767 again.

- 769 • A first slice gathers the first two columns  $(0, 1)$ ; these are all zero except for the top  $2 \times 2$   
 770 block, set as identity  $I_2$ . This slice replicates, again, the  $\text{ReLU}(p_{i,j}) = p_{i,j}$  mark features in  
 771 the first two channels of the output.
- 773 • A second slice gathers columns  $2$  through  $2 + d - 1$ ; these are zero in the first two rows,  
 774 rows  $2$  through  $2 + d - 1$  are set to identity  $-\vec{1}_d$  and the last  $d$  rows are set to  $I_d$ . This slice  
 775 calculates, channel-by-channel,  $\text{ReLU}(h_2) - \text{ReLU}(h_1) = h_2 - \text{ReLU}(h_1)$ .
- 776 • A third — and last — slice gathers columns  $2 + d$  through  $2 + 2d - 1$ ; these are all zero  
 777 except for rows  $2$  through  $2 + d - 1$ , set to identity  $I_d$ . This slice copies  $\text{ReLU}(h_1)$  in the  
 778 last  $d$  channels of the output.

780 Bias vector  $b_2$  is set to zero.

781 As a result, the output is a vector in  $\mathbb{R}^{2d+2}$  with the following structure:

$$783 \left[ p_{i,j} \mid \underbrace{\cdots (x_{E,(i,j)})_{b(l,h)} - \text{ReLU}((x_{E,(i,j)})_{b(l,h)} + (p_{i,j})_0 - 1) \cdots}_{(o_1)} \mid \underbrace{\cdots \text{ReLU}((x_{E,(i,j)})_{b(l,h)} + (p_{i,j})_0 - 1) \cdots}_{(o_2)} \right].$$

786 Now, note that, as desired:

- 788 • if  $(p_{i,j})_0 = 1$  (message from response token):
- 790    -  $(o_1)_{b(l,h)} = (x_{E,(i,j)})_{b(l,h)} - \text{ReLU}((x_{E,(i,j)})_{b(l,h)} + 1 - 1) = (x_{E,(i,j)})_{b(l,h)} -$   
 791        $(x_{E,(i,j)})_{b(l,h)} = 0$   

793    -  $(o_2)_{b(l,h)} = \text{ReLU}((x_{E,(i,j)})_{b(l,h)} + 1 - 1) = (x_{E,(i,j)})_{b(l,h)},$

795 • if  $(p_{i,j})_0 = 0$  (message from prompt token):

797    -  $(o_1)_{b(l,h)} = (x_{E,(i,j)})_{b(l,h)} - \text{ReLU}((x_{E,(i,j)})_{b(l,h)} + 0 - 1) = (x_{E,(i,j)})_{b(l,h)} - 0 =$   
 798        $(x_{E,(i,j)})_{b(l,h)}$   

800    -  $(o_2)_{b(l,h)} = \text{ReLU}((x_{E,(i,j)})_{b(l,h)} + 0 - 1) = 0$

801 Now, when aggregating these calculated messages through summation, it becomes clear that the  
 802 aggregated message vector  $m_i$  will eventually hosts:

- 804 • The number of response tokens preceding token  $i$ , i.e.,  $i - n_p - 1$ , in channel  $0$ ;  
 805 • The length of the prompt, i.e.,  $n_p$ , in channel  $1$ ;  
 807 • Summation  $\sum_{j=0}^{n_p-1} \alpha_{l,h}^{i,j}$  in channel  $b(l, h) + 2$ , denoted  $\hat{P}_i^{l,h}$ ;  
 808 • Summation  $\sum_{j=n_p}^{i-2} \alpha_{l,h}^{i,j}$  in channel  $b(l, h) + 2 + d$ , denoted  $(\hat{R}_i^{l,h})^-$ .

(2) **Update function.** Next, we desire  $up_0$  to implement a function  $up^*$  mapping the concatenation  $[x_{V,i} \mid m_i]$  to a vector of dimension  $d$  corresponding to the Lookback Lens output scores in Equation (3). Showing that  $up_0$  can approximate  $up^*$  up to desired precision  $\epsilon$  will complete the proof.

We now describe  $up^*$ . We build it as a composition of two functions  $f_A \circ f_B$ :

- $f_A$  is such that  $f_A(x_{V,i}, m_i) = y_i = c_i + m_i$ , where  $c_i$  has the same dimensionality as  $m_i$  ( $2d + 2$ ), and  $c_i = [0 \mid 1 \mid \vec{0} \mid x_{V,i}]$  — note that vector  $y_i$  is an “updated” version of  $m_i$  whereby channel 1 now equals  $i - n_p$  and channels  $b(l, h) + 2 + d$ ’s are now such that:

$$(\hat{R}_i^{l,h})^- + x_{V,i} = \sum_{j=n_p}^{i-1} \alpha_{l,h}^{i,j} = \hat{R}_i^{l,h} \quad (6)$$

- $f_B$  is such that  $f_B(y_i) = z_i$  where  $z_i$  is of dimensionality  $d$  and:

$$(z_i)_{b(l,h)} = \frac{\binom{y_i}_{b(l,h)+2}}{\binom{y_i}_1} \quad (7)$$

$$(z_i)_{b(l,h)} = \frac{\binom{y_i}_{b(l,h)+2}}{\binom{y_i}_1} + \frac{\binom{y_i}_{b(l,h)+2+d}}{\binom{y_i}_0}.$$

Note that, importantly,  $(z_i)_{b(l,h)} = \frac{\hat{P}_i^{l,h}/n_p}{\hat{P}_i^{l,h}/n_p + \hat{R}_i^{l,h}/(i - n_p)} = \frac{P_i^{l,h}}{P_i^{l,h} + R_i^{l,h}} = \ell_i^{l,h}$ .

First,  $f_A$  can be realised by a single affine transformation. This has weight matrix  $W_A$  in  $\mathbb{R}^{(3d+2) \times (2d+2)}$ , described in terms of row slices as follows.

- A first slice gathers the first  $d$  rows; it is zero except for the last  $d$ -column block, set as identity  $I_d$ .
- A second slice gathers rows  $d$  through  $3d + 2 - 1$  and it set as an identity  $I_{(2+2d)}$ .

The above linear transformation has the effect of copying  $m_i$  in the output and of summing  $x_{V,i}$  in its last  $d$  entries — where the aggregated message from response tokens is stored. Now, bias vector  $b_A$  is in  $\mathbb{R}^{2+2d}$  and is zero everywhere except for its first element which is set as 1. Adding  $b_A$  has the effect of simply increasing the second entry by one, thus “updating” the count of response messages stored there.

Second, we note that  $f_B$  computes the same exact scalar-valued function  $f_B^s$  on each output component; also, this  $f_B^s$  is continuous on (the non-compact) domain  $\{1, \dots, \bar{n}_p\} \times \{1, \dots, \bar{n}_r\} \times (0, 1)^2$ . We note that  $f_B^s$  can be trivially, continuously extended to the compact  $\{1, \dots, \bar{n}_p\} \times \{1, \dots, \bar{n}_r\} \times [0, 1]^2$ : it suffices to see that its limits exist and are finite on the boundary of  $[0, 1]^2$ . This is indeed the case; we note that denominators in each individual normalisation ratios are always greater or equal than one; and that the two addenda in the denominator of the main ratio are always non-negative, can never evaluate simultaneously to zero and their sum is bounded away from zero (given the maximum allowed length for prompt and response).

Given this premise, term  $f_B^{s,\text{ext}}$  this continuous extension; we can invoke the MLP Universal Approximation Theorem (Pinkus, 1999) to claim the existence of an MLP  $M_B^s$  with one hidden layer which approximates the continuous  $f_B^{s,\text{ext}}$  on the compact domain  $\{1, \dots, \bar{n}_p\} \times \{1, \dots, \bar{n}_r\} \times [0, 1]^2$  up to precision  $\epsilon$ . This implies the original  $f_B^s$  is also  $\epsilon$ -approximated in its original domain. Now, it is possible to (easily,) appropriately *replicate* the weights of  $M_B^s$  to distribute its same exact computation for each of the output coordinates, thus obtaining an MLP  $M_B$  approximating the overall  $f_B$ :  $\forall y, \forall i |M_B^s(y) - (f_B(y))_i| < \epsilon$ , that is,  $\forall y, \forall i |(M_B(y))_i - (f_B(y))_i| < \epsilon$ . But, then:

$$\begin{aligned} \forall y, \forall i |(M_B(y))_i - (f_B(y))_i| &< \epsilon \\ \implies \forall y \max \left( |(M_B(y))_i - (f_B(y))_i| \mid i = 0, \dots, d-1 \right) &< \epsilon \\ \implies \forall y \|M_B(y) - f_B(y)\|_\infty &< \epsilon \end{aligned}$$

Denote  $(W_B^1, b_B^1), (W_B^2, b_B^2)$  the weight and biases of, respectively, the first and second layers of  $M_B$ . Now, the above affine transformation exactly implementing  $f_A$  can be “absorbed” into the *first* layer of  $M_B$ , by replacing the weight and bias as  $W^1 = W_B^1 \cdot W_A, b_1 = (W_B^1 \cdot b_A + b_B^1)$ . The resulting MLP composed by  $((W^1, b^1), (W_B^2, b_B^2))$  now  $\epsilon$ -approximates the overall  $u^*$ , concluding the proof.  $\square$

**Proposition (informal) 2.** *With a single-layer message-passing stack  $f_{mp}$ , CHARM can arbitrarily well approximate global LLM-Chk-l features  $c^l$ , provided attentions are clipped away from zero<sup>6</sup>.*

**Proposition 2.** *Let  $c^l(\vec{s}; \mathcal{L})$  denote the LLM-Check Attention Score (Equation (4)) calculated on string  $\vec{s} = \vec{p} \mid \vec{r}$  for LLM  $\mathcal{L}$  (Equation (4)) and its layer  $l$ . Also, let  $y(G_{\vec{s}})$  denote the prediction in output from an CHARM stacking components  $f_{msg}, f_{pool}, f_{pred}$ , run on corresponding graph  $G_{\vec{s}}$ . For any precision  $\epsilon > 0$ , when  $f_{pool} \equiv \sum$  and attention values are clipped from below to value  $\alpha_{min} > 0$ , there exists a 1-layered  $f_{msg}$ , and an MLP  $f_{pred}$  such that CHARM approximates LLM-Check Attention Scores up to precision  $\epsilon$ . That is, for any prompt-response pair  $\vec{s}$ ,  $|y(G_{\vec{s}}) - c^l(\vec{s}; \mathcal{L})| < \epsilon$ .*

*Proof.* To prove the above we show that, in the setting described in the proposition, there exist a single-layer stack  $f_{mp}$ , as well as an MLP  $f_{pred}$  which compute the desired approximation. Here we consider the LLM-Chk-l variant which averages across heads instead of performing a summation, but the proof below is easily extended to this alternative configuration.

**(0) Setup, inputs and proof strategy.** Our stack is made up, again, of one single message-passing layer in  $f_{mp}$ , followed by sum-based pooling and an MLP  $f_{pred}$ . Our specific interest is in showing the existence of appropriate MLPs  $msg_0, up_0, f_{pred}$  enabling the full stack to realise the target approximation. Again,  $msg_0, up_0$  are the components of the *first* — and only — message-passing layer in  $f_{mp}$ , so its input node and edge representations effectively correspond to the original node and edge features  $X_V, X_E$ . The whole computation then takes the form:

$$y = \underbrace{f_{pred}}_{(3)} \left( \sum_{i=0}^{n-1} \underbrace{up_0}_{(2)} \left( x_{V,i}, \sum_{(i,j) \in E} \underbrace{msg_0}_{(1)} (x_{V,i}, x_{V,j}, x_{E,(i,j)}, p_{i,j}) \right) \right) \quad (8)$$

**(1) Message function.** The LLM-Chk-l method does not perform any aggregation on the attention graph — for our purposes it suffices for MLP  $msg_0$  to simply implement a function which outputs any constant non-negative vector  $v$ . W.l.o.g., set  $v = \vec{0}$ ; the MLP implementing  $msg_0$  is trivially obtained by setting both its weights and biases to zero.

**(2) Update function.** Note that  $x_{V,i} = \alpha_{i,i}$ ; as it will be clearer next, it suffices for  $up_0$  to approximate the log function applied thereon component-wise. Now, log is, in fact, operating on domain  $[\alpha_{min}, 1]$  in view of the applied clipping; there, the function is continuous. Consider the compact set  $[\alpha_{min}, 1]$  obtained as the closure of the above domain. The considered log is trivially continuously extended to this new domain, since its limit exists finite as the argument approaches 1 from the left (it evaluates to 0). We therefore invoke the Universal Approximation Theorem (Pinkus, 1999), which guarantees the existence of an MLP  $M_{log}$  approximating the component-wise log function on the extended domain and, in turn, on the original input domain  $[\alpha_{min}, 1]$  up to arbitrary precision  $\epsilon$ . We construct  $up_0$  by appropriately replicating the weights of  $M_{log}$  to output the approximated log-transform on each of the first input  $d$  channels in parallel, whilst discarding the last  $d$  remaining channels in output from the message function (see above).

**(3) Prediction function.** At this point, the input of  $f_{pred}$  corresponds to:  $\hat{z} = \frac{1}{n} \sum_{i=0}^{n-1} \hat{y}_i$ , where  $\hat{y}_i$  is an  $\epsilon$ -approximation of the log function applied element-wise to  $\alpha_{i,i}$ ’s. If  $f_{pred}$  implemented the final averaging over  $l$ ’s  $H$  heads (channels), then its output would be an overall approximation of the desired quantity  $c^l$  (of a certain precision yet to be quantified). We note that it easy to explicitly construct such an MLP exactly implementing the outer averaging over the selected heads. We describe its two layers in the following.

<sup>6</sup>This assumption ensures the log is continuous on an appropriate compact set, rendering its approximation amenable (see Appendix A); in practice we also did observe it was necessary to ensure numerical stability.

- 918
- First layer: it features a weight matrix of the form  $[I_d \mid -I_d]$  and a zero-valued bias; this  
919 layer expands the  $d$ -dimensional input to a  $2d$  representation, where the the first  $d$  channels  
920 host input  $\hat{z}$ , the last  $d$  channels the negated input  $-\hat{z}$ .

921

  - Second layer: it features a weight matrix in  $\mathbb{R}^{2d \times 1}$ . The upper  $d \times 1$  block hosts a vector  
922 where entries  $b(l, h)$  equal to  $1/H$  if  $l = \bar{l}$ , 0 otherwise. The lower  $d \times 1$  block has the same  
923 exact structure, but non-zero entries are, instead, set to  $-1/H$ .

924

925 It is easy to see this construction exactly implements the required averaging, making the ReLU  
926 activation act neutrally.

927

928 Now, we ask to what precision does the final output approximate overall target  $c^l$ . Note that the only  
929 source of approximation is in the previously introduced  $u_0$ ; we are thus only required to quantify  
930 how it “propagates” to the rest of the following computation. We have, by the triangular inequality:

$$\begin{aligned} & \left| \frac{1}{H} \sum_{h=0}^{H-1} \frac{1}{n} \sum_{i=0}^{n-1} (\hat{y}_i)_{b(l,h)} - \frac{1}{H} \sum_{h=0}^{H-1} \frac{1}{n} \sum_{i=0}^{n-1} \log(\alpha_{i,i}^{l,h}) \right| \leq \\ & \frac{1}{H} \sum_{h=0}^{H-1} \left| \frac{1}{n} \sum_{i=0}^{n-1} (\hat{y}_i)_{b(l,h)} - \frac{1}{n} \sum_{i=0}^{n-1} \log(\alpha_{i,i}^{l,h}) \right| \leq \\ & \frac{1}{H} \sum_{h=0}^{H-1} \frac{1}{n} \sum_{i=0}^{n-1} \left| (\hat{y}_i)_{b(l,h)} - \log(\alpha_{i,i}^{l,h}) \right| \leq \\ & \frac{1}{H} \sum_{h=0}^{H-1} \frac{n\epsilon}{n} = \frac{Hn\epsilon}{Hn} = \epsilon \end{aligned}$$

943 which concludes the proof. □

944

## 945 B DATASET DETAILS

946

### 947 B.1 NQ AND CNN

948

949 **Dataset construction.** These datasets are constructed precisely following the implementation  
950 described in (Chuang et al., 2024) and provided as part as a supplementary codebase at <https://github.com/voidism/Lookback-Lens>. From this repository we derive both prompts  
951 and pre-computed, annotated generations, which we re-use via teacher-forcing to hook out the  
952 required computational traces, namely attention and activation matrices. We tested the fidelity of  
953 these generated scores in early experiments: we recalculated original Lookback Lens features  
954 using the generated data and managed to reproduce the original results in (Chuang et al., 2024).

955

956 **Dataset details**, including a description of the text generation and annotation process, are found  
957 in (Chuang et al., 2024, Appendix A) and (Chuang et al., 2024, Appendix C.2), to which we refer the  
958 interested reader.

959

960 **Splitting.** Chuang et al. (2024) originally split the data randomly and in a way that, potentially,  
961 passages from the same response could appear across training and evaluation splits. We argue this is  
962 an undesired side-effect and, in an effort to cast the HD in a more challenging setup, we instead split  
963 the data at the level whole prompt-response pairs (graphs according to our framework). Specifically,  
964 we fix the seed to 42 and randomly obtain a prompt-response level split in the proportion 60% / 20%  
965 / 20% (train / val / test).

966

### 967 B.2 MOVIES, WINOBIAS, MATH

968

969 **Dataset construction.** These datasets are constructed following the process described in (Orgad  
970 et al., 2024), and by leveraging the authors’ code open-sourced at <https://github.com/technion-cs-nlp/LLMsKnow> (MIT License). The prompts and ground truth labels of all the  
971 three considered dataset, in particular, are provided by the authors themselves in the above codebase.

972 As for hallucination labels, we run the annotation process whose code is provided therein. These  
 973 annotation routines are mostly based on string matching procedures.  
 974

975 **Dataset details**, including a description of the datasets and how prompts have been derived are  
 976 provided in (Orgad et al., 2024, Appendix A.3), to which we refer the interested reader.  
 977

978 **Splitting.** We use the same train/test splits provided by Orgad et al. (2024), and additionally carve  
 979 out a random sample of 20% of training data points, treated as our validation set. We perform this  
 980 sampling by setting random seed to 42.  
 981

## 982 C EXTENDED EXPERIMENTAL SECTION

### 984 C.1 COMPARISON WITH SELF-CHECK AND MULTIPLE-PROMPTING-BASED METHODS

986 In this section, we compare against baseline methods that rely on additional prompting, specifically  
 987 P(True) Kadavath et al. (2022) and Semantic Entropy (SE) Kuhn et al. (2023). Both approaches  
 988 operate over multiple LLM generations or prompts, which introduces a non-negligible computational  
 989 overhead and may hinder their applicability in real-time settings. Table 6 reports results on the Movies  
 990 dataset using `Mistral-7B-instruct`. For SE, we follow the original evaluation setup (Kuhn  
 991 et al., 2023), employing the DeBERTa entailment model as described in the referenced work.  
 992

993 Table 6: Comparison with methods relying on multiple prompting.

995 <b>Method</b>	996 <b>Mis-7B – Movies (AUC)</b>
996 P(True)	62.00
997 Semantic Entropy	70.06
998	
999 CHARM (att) (ours)	<b>80.3±0.2</b>
1000 CHARM (att+act-24) (ours)	<u>79.7±0.3</u>

1001 We observe that in all cases, CHARM variants substantially outperform the competing approaches.  
 1002 To quantify the computational burden of these baselines, we measured the average runtime of SE  
 1003 for producing a prediction. This process involves generating 10 additional responses and clustering  
 1004 them by computing mutual entailments with an auxiliary DeBERTa model. On average, SE required  
 1005  $5.9 \pm 1.7$  seconds per evaluation. We minimised the overhead of auxiliary generations by running  
 1006 them in parallel through batching. Nevertheless, the clustering step alone accounts for about 1.35  
 1007 seconds of runtime, which is not negligible. These findings highlight the advantage of our method,  
 1008 which not only achieves higher accuracy but also operates orders of magnitude faster, with detection  
 1009 runtimes on this dataset in the range of  $10^{-4}$  seconds.  
 1010

### 1011 C.2 IMPACT OF THE CHOICE OF ACTIVATION LAYER(S)

1012 Here we report results on the Math dataset in an effort to assess the sensitivity of CHARM across  
 1013 different choices of the activation layer(s). We train (and tune) CHARM with features from the layers  
 1014 considered for activation probes, namely 24, 28, 32. We also experiment with considering all these  
 1015 layers jointly, with activations concatenated together across layers. Results are in Table 7.  
 1016

1017 CHARM remains robust w.r.t. the chosen activation layer and we observe that concatenating multiple  
 1018 layers together may further improve performance.  
 1019

### 1020 C.3 PERFORMANCE BY LENGTH

1022 We run additional experiments on the CNN dataset to better study how model performance, inference  
 1023 run-time and memory footprint vary as a function of the processed sequence lengths.  
 1024

1025 **Runtime and memory consumption by length** have their trends reported in Figure 3. These  
 1026 results show extremely contained run-times and memory consumptions even for lengths in the order

Table 7: Performance of CHARM with different choices of activation layers.

Method	Math	
	AUROC	AUPR
Act-24	77.7	77.5
Act-28	78.1	77.8
Act-32	76.6	77.9
CHARM (att)	$76.5 \pm 1.1$	$79.7 \pm 0.5$
CHARM (att+act-24)	$80.8 \pm 0.7$	$83.1 \pm 0.7$
CHARM (att+act-28)	$81.2 \pm 1.0$	$83.4 \pm 1.3$
CHARM (att+act-32)	$81.7 \pm 0.2$	$83.8 \pm 0.3$
CHARM (att+act-(24,28,32))	$82.7 \pm 0.1$	$84.0 \pm 0.5$

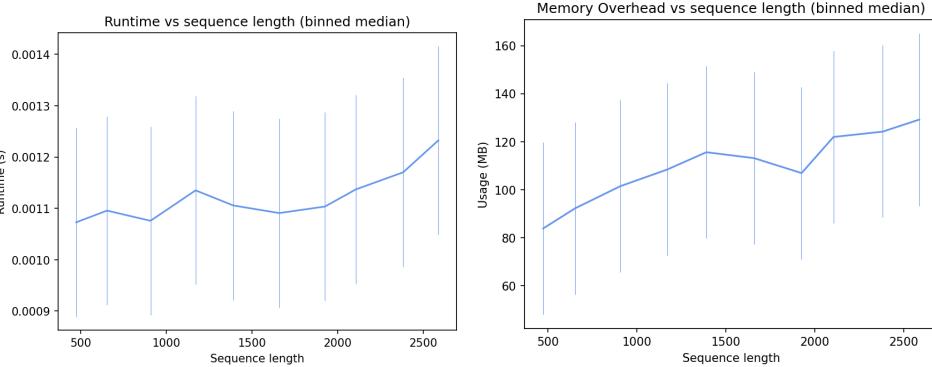


Figure 3: Inference runtime (left) and memory consumption (right) by sequence length on the CNN dataset. Values are grouped in 10 bins; x-ticks report median length per bin; the y-axis reports median measurements per bin, as well as the corresponding inter-quartile range.

of thousands of tokens. Overall, runtime and memory consumption scale very favourably in the considered range, underscoring the crucial computational advantage of running neural message-passing on *sparsified* attention graphs.

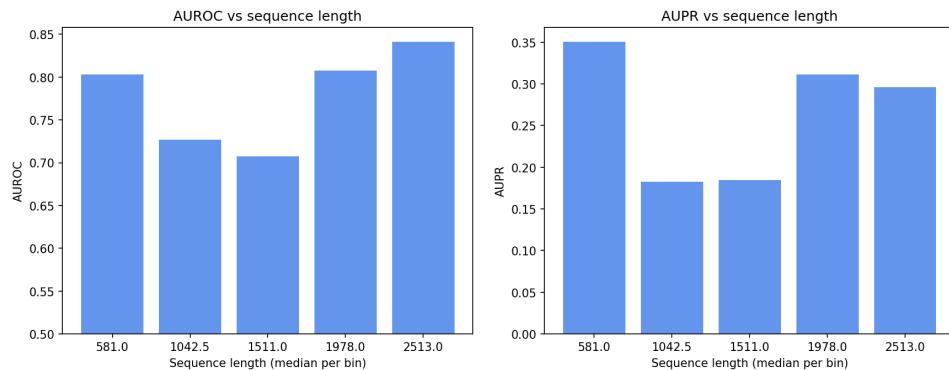


Figure 4: Test performance by sequence length on the CNN dataset: AUROC (left), AUPR (right). Values are grouped in 5 bins; x-ticks report median length per bin; the y-axis reports the performance within each bin.

**Performance by length** is reported in Figure 4 both in terms of AUROC and AUPR. We observe that CHARM’s performance does not degrade with increasing sequence length; in fact, the highest AUROC is achieved in the longest-sequence bin. Overall, these results suggest that CHARM is not particularly sensitive to sequence length.

1080  
 1081 **Generalising to longer sequences at test time.** We also experimented with training on shorter  
 1082 sequences and testing on longer ones, again on the CNN dataset. We produced a new non-uniform  
 1083 split whereby test samples are taken as the top 20% longest sequences, the rest considered for  
 1084 training and validation sets. This way, training and model selection is run on sequences only up to  
 1085  $\sim 1.2k$  tokens, while the model is tested on sequences of length ranging from this value to  $\sim 2.7k$   
 1086 tokens. Results are reported in Table 8, comparing CHARM with other Act-\* and Lookback Lens  
 1087 baselines.  
 1088

Table 8: Performance of CHARM on CNN, split by size (test set contains unseen longer sequences).

Method	CNN (longer seqs.)	
	AUROC	AUPR
Act-24	71.6	16.3
Act-28	70.3	16.0
Act-32	68.4	14.1
Lookback Lens	73.9	18.9
CHARM (att)	<b>74.5 <math>\pm 0.7</math></b>	<b>21.6 <math>\pm 1.2</math></b>

1098 We observe that CHARM still outperforms methods in comparisons on both metrics, and that its  
 1099 performance on unseen sequence lengths remains stable, only marginally lower than what obtained  
 1100 on uniform splits (Table 1).  
 1101

#### C.4 HYPERPARAMETER GRIDS

1104 We employed the same hyperparameter grid search across all datasets considered for CHARM, as  
 1105 summarized in Table 9. When incorporating activations into CHARM, we additionally searched over a  
 1106 separate weight decay parameter, applied only to the encoder of the activations, with candidate values  
 1107  $\{0.0, 0.05, 0.1\}$ .  
 1108

Table 9: Hyperparameter search space for CHARM.

Hyperparameter	Values
Learning Rate	$\{0.001, 0.0005\}$
Learning Rate Sched.	$\{\text{Reduce On Plateau, Cosine w/ Warmup}\}$
Batch Size	32
Dropout	$\{0.25, 0.5\}$
Hidden Dimension	$\{32, 64, 128\}$
Number of Layers	$\{1, 2, 3\}$
Weight Decay	$\{0.0, 0.001\}$
BatchNorm	$\{\text{yes, no}\}$
Residual Connections	$\{\text{yes, no}\}$

##### C.4.1 BASELINE HYPERPARAMETER SEARCHES

1124 All hyperparameters were selected based on validation performance and, in particular, in order to  
 1125 maximise the AUPR metric. The details are provided below.  
 1126

**Probas** We evaluated different readout functions — mean, max, and sum — applied to the next-  
 1127 token probabilities.  
 1128

**Act-\*** We experimented with the following regularisation parameters for logistic regression:  $C \in \{10^{-8}, 10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 100, 10^5\}$ . In addition, we probed token  
 1129 positions in:  $\{-3, -2, -1, 0, 1, 2\}$ .  
 1130

**LLM-Chk-\*** We were required to clamp the attention scores from below using  $\epsilon = 10^{-6}$  to avoid numerical errors. For LLM-Chk++-\*, we experimented with  $C \in \{10^{-8}, 10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 100, 10^5\}$ .  
 1131

1134 **LapEig** We experimented with the following values of  $k$ :  $\{4, 5, 6, 7, 8, 9, 10, 11, 12, 15, 20\}$ . For  
 1135 datasets where the minimum number of tokens in the test split was less than  $k$ , we restricted  
 1136 experiments to values of  $k$  below this threshold. As for logistic regression, we used  $C = 1$ , class  
 1137 balancing, and a maximum of 2,000 iterations, consistent with what prescribed in the original  
 1138 paper (Sriramanan et al., 2024).

1139 **Neigh-Avg (N) and Neigh-Avg (E)** We used mean readout exactly as in  
 1140 our model and tuned the logistic regression regularisation parameter over  $C \in$   
 1141  $\{10^{-8}, 10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 100, 10^5\}$ .  
 1142

1143 **Lookback Lens.** We implemented Lookback Lens exactly as described in the original paper,  
 1144 using logistic regression with a regularization parameter of  $C = 1$ . For Lookback Lens  $^\dagger$ , we per-  
 1145 formed a grid search over  $C \in \{10^{-8}, 10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 100, 10^5\}$ .  
 1146

1147 **CHARM (no g.)** We use the same exact grid as for CHARM (see Table 9), with the following  
 1148 exceptions: (1) no msg-passing layers; (2) a readout / prediction head that is either linear or a  
 1149 implemented as an MLP.  
 1150

## D IMPLEMENTATION DETAILS AND COMPUTATIONAL RESOURCES

### D.1 DETAILED ARCHITECTURAL FORMS AND TRAINING PARAMETERS

#### D.1.1 NEIGHBOURHOOD AVERAGING BASELINES

1156 Our baselines Neigh-Avg (N) and Neigh-Avg (E) calculate features in a non-learnable way as  
 1157 described below.  
 1158

##### 1159 **Neigh-Avg (N).**

$$\begin{aligned}
 1161 \quad h_i^{(1)} &= \frac{1}{1 + \deg_{\text{in}}(i)} \left( h_i^{(0)} + \sum_{j: (i,j) \in E} x_{E,(i,j)} \right) = \\
 1162 &= \frac{1}{1 + \deg_{\text{in}}(i)} \left( x_{V,i} + \sum_{j: (i,j) \in E} x_{E,(i,j)} \right) = \\
 1163 &= \frac{1}{1 + \deg_{\text{in}}(i)} \left( \alpha_{i,i} + \sum_{j: (i,j) \in E} \alpha_{j,j} \right) \tag{9}
 \end{aligned}$$

##### 1169 **Neigh-Avg (E).**

$$\begin{aligned}
 1171 \quad h_i^{(1)} &= \frac{1}{1 + \deg_{\text{in}}(i)} \left( h_i^{(0)} + \sum_{j: (i,j) \in E} h_j^{(0)} \right) = \\
 1172 &= \frac{1}{1 + \deg_{\text{in}}(i)} \left( x_{V,i} + \sum_{j: (i,j) \in E} x_{V,j} \right) = \\
 1173 &= \frac{1}{1 + \deg_{\text{in}}(i)} \left( \alpha_{i,i} + \sum_{j: (i,j) \in E} \alpha_{j,j} \right) \tag{10}
 \end{aligned}$$

1180 Outputs  $h_i^{(1)}$  are then fed in input to a logistic regression model, regularised as illustrated in Ap-  
 1181 pendix C.4.1. Before that, they are averaged-pooled in the case of response-wise predictions tasks.  
 1182

#### 1183 D.1.2 EXPERIMENTAL CHARM FORM

1184 Throughout all our experiments, CHARM implements the following msg-passing equation:

$$h_i^{(t+1)} = \text{up}_t \left( \left[ h_i^{(t)} \mid \frac{1}{\deg_{\text{in}}(i)} \sum_{j: (i,j) \in E} \text{msg}_t \left( [h_j^{(t)} \mid x_{E,(i,j)}^\tau \mid p_{i,j}] \right) \right] \right). \tag{11}$$

1188 Initial node features always include “reflexive attention”, i.e., that a token pays to itself. When  
 1189 additionally including activations (CHARM (att+act-\*)) we employed an additionally encoder to  
 1190 preprocess these. The output of this module is concatenated to the original attention features before  
 1191 message passing takes place. Note that, in all our experiments, for computational reasons and  
 1192 in alignment with the computational flow of Lookback Lens, we remove all connections from  
 1193 prompt to prompt tokens.

1194  
 1195 **D.1.3 OPTIMIZER AND SCHEDULERS**

1196 For all datasets and tasks, we use the AdamW optimizer Loshchilov & Hutter (2017). We experi-  
 1197 mented with two learning rate schedulers (see Table 6): “Reduce On Plateau” and “Cosine Annealing  
 1198 with Warmup”, where warmup spanned 10% of the total training steps. The scheduler yielding the  
 1199 best validation performance was selected.

1200  
 1201 **D.2 CODE IMPLEMENTATION**  
 1202

1203 Thee implementation of CHARM was realised by means of PyTorch (Paszke et al., 2019) and PyTorch  
 1204 Geometric (Fey & Lenssen, 2019) (available respectively under the BSD and MIT license). We  
 1205 performed hyperparameter tuning using the Weight and Biases framework (Biewald, 2020). For  
 1206 baselines and models running logistic regression, we resorted to the implementation exposed by  
 1207 the Sci-Kit Learn library (BSD license). LapEig required also running the PCA dimensionality  
 1208 reduction technique; we invoked the python implementation from the same library.

1209  
 1210 **D.3 EXPERIMENTAL RESOURCES AND ARTEFACTS**

1211 We ran our all our experiments on NVIDIA L40 GPUs. The two employed LLMs were both accessed  
 1212 via Hugging Face python API, in particular:

- 1214 • LLaMa-2-7b-chat (Touvron et al., 2023) (License: LLaMa 2 Community License). Ac-  
 1215 cessed at <https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>.
- 1216 • Mistral-7b-instruct (Jiang et al., 2023) (License: Apache-2.0). Accessed at  
 1217 <https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>.

1218  
 1219 **E VISUALISATIONS**  
 1220

1221 We provide here two sample visualisations of the constructed computational graphs in input to our  
 1222 CHARM architecture. These illustrate two *test* data points in the NQ dataset and are reported in  
 1223 Figures 5 and 6, where we show the full computational graph (left) along with a zoom-in on the  
 1224 response tokens (right).

1225 These visualisations arrange prompt tokens on the left, and response tokens on the right. Edges are  
 1226 drawn in a way that their thickness and transparency is proportional to the corresponding attention  
 1227 scores, averaged, for illustrative purposes, across layers and heads. Response tokens have their border  
 1228 coloured according to the ground truth — red: the token is in an hallucinated passage; blue: otherwise.  
 1229 Their interior is filled, instead, with a colour that conveys the model prediction — “more red”: the  
 1230 token is more likely to be in an hallucinated passage; “more blue”: more likely to be in a correct,  
 1231 non-hallucinated passage. Output scores for our model are matched to this colour map linearly after a  
 1232 min-max normalisation. Note the absence of prompt-to-prompt edges, as they are neglected in our  
 1233 experiments as explained at the end of Appendix D.1.2.

1234  
 1235 **F LARGE LANGUAGE MODEL (LLM) USAGE**  
 1236

1237 We employed large language models (LLMs) to support the writing process, specifically for improving  
 1238 clarity in technical explanations, refining grammar and style, and enhancing overall readability. LLMs  
 1239 were also used to a limited extent to aid the process of finding related works. All research contributions,  
 1240 including the design of experiments, data analysis, and conclusions, are entirely our own. The LLMs  
 1241 were used strictly as writing aids to improve presentation quality, not for generating research content  
 or shaping the substance of our work.

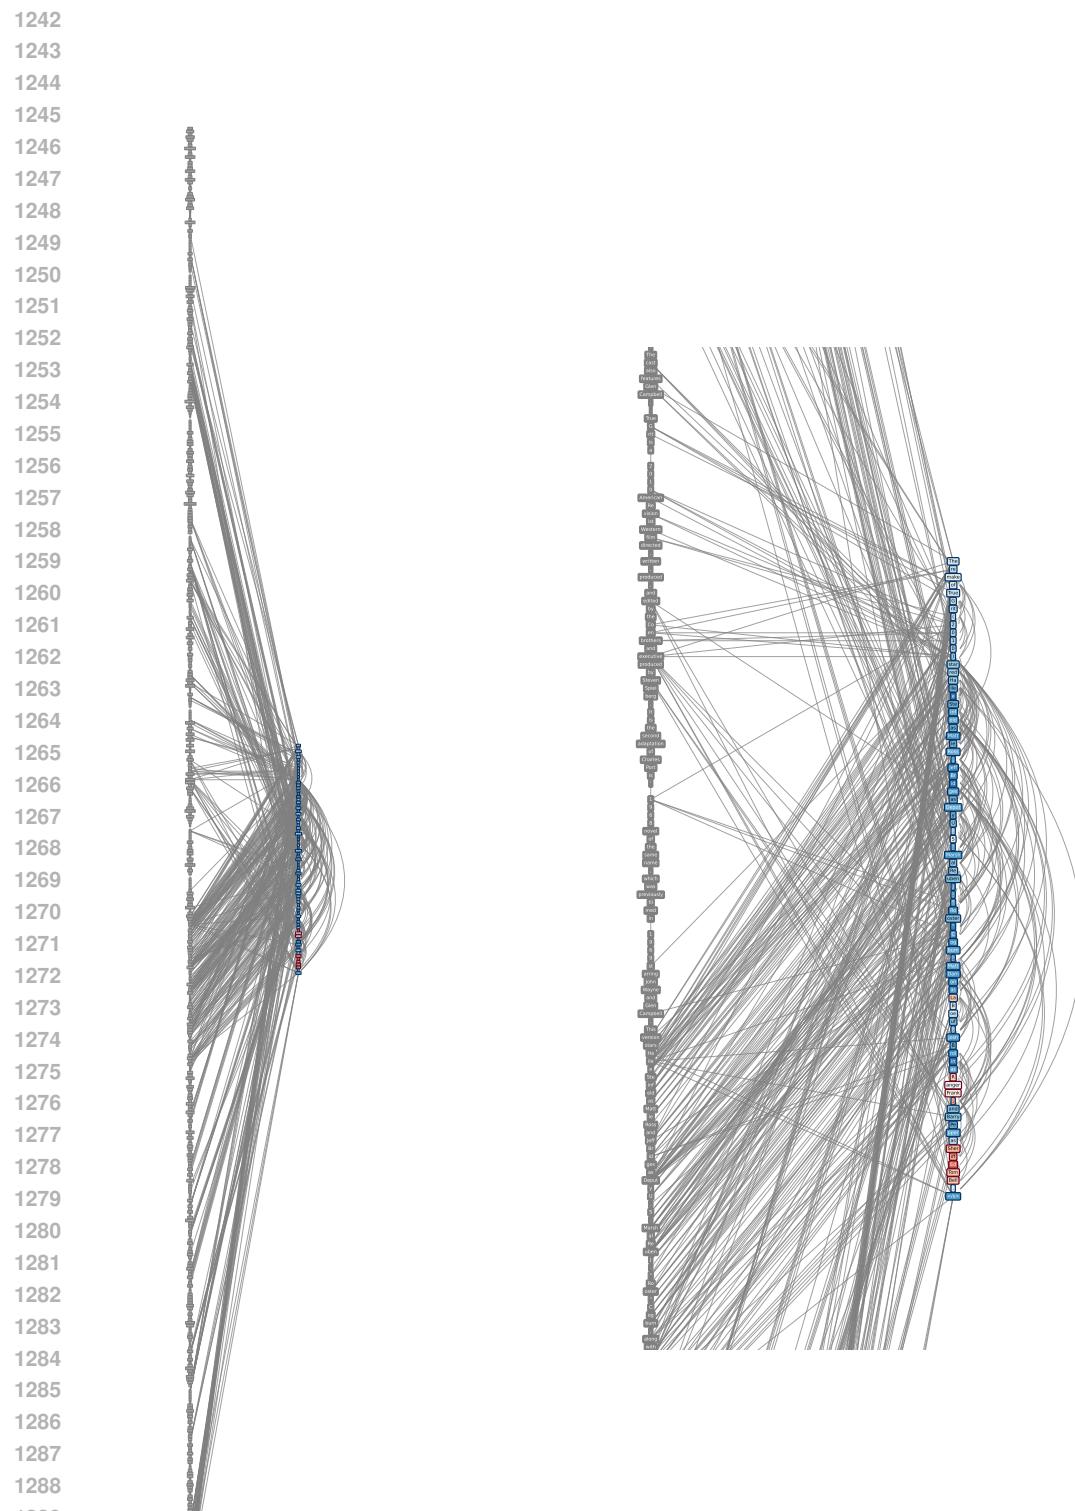


Figure 5: Visualisation of test sample 1590, along with token-wise labels and predictions.

1296  
1297  
1298  
1299  
1300  
1301  
1302  
1303  
1304  
1305  
1306  
1307  
1308  
1309  
1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317  
1318  
1319  
1320  
1321  
1322  
1323  
1324  
1325  
1326  
1327  
1328  
1329  
1330  
1331  
1332  
1333  
1334  
1335  
1336  
1337  
1338  
1339  
1340  
1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349

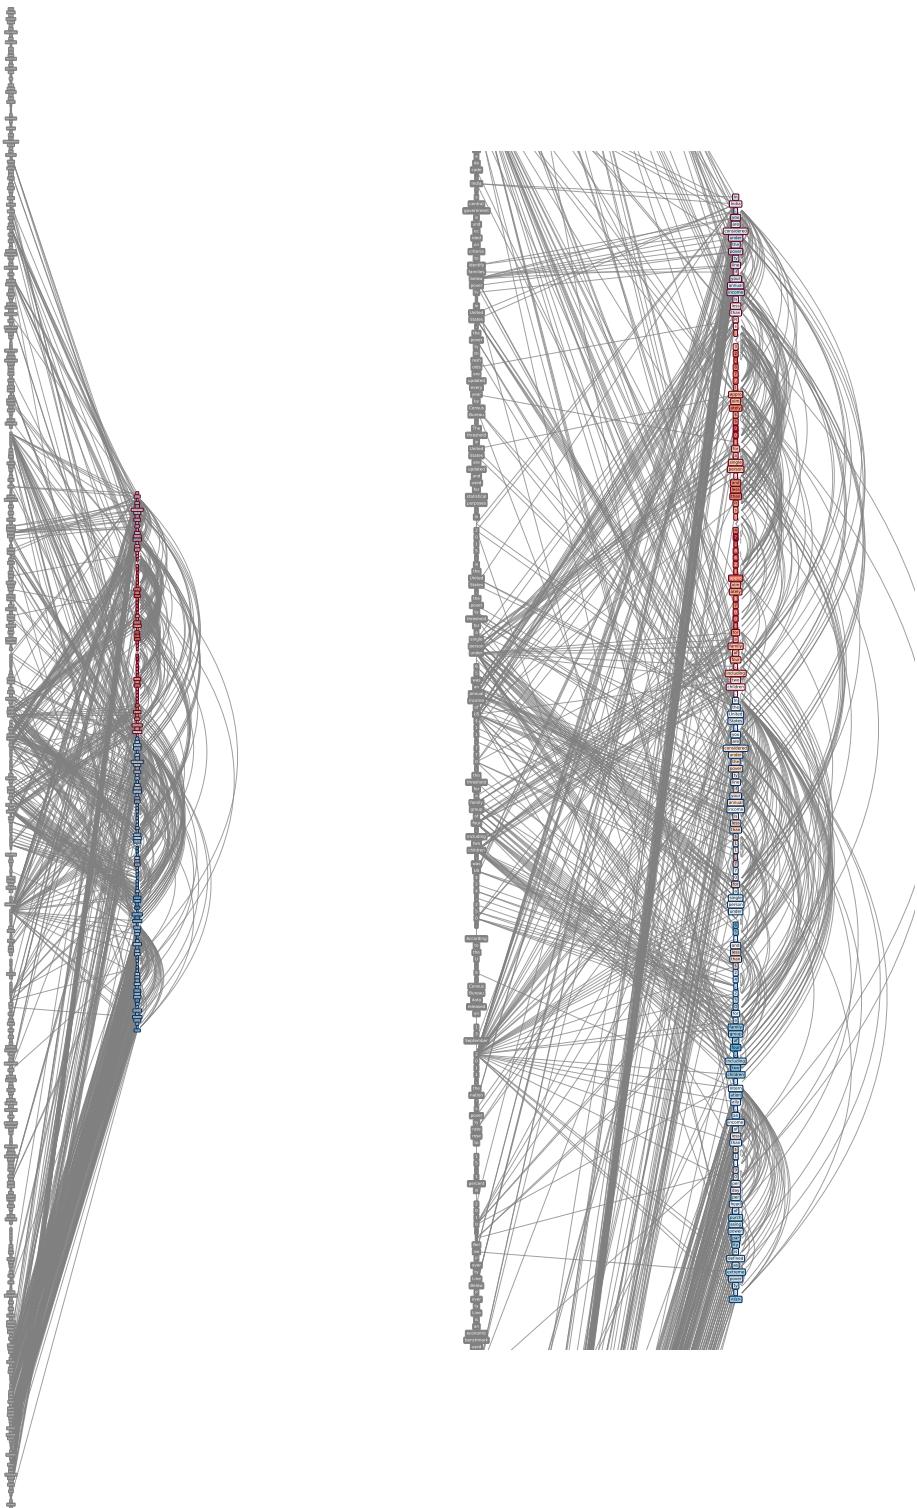


Figure 6: Visualisation of test sample 2354, along with token-wise labels and predictions.