

Multi-Step Knowledge Interaction Analysis via Rank-2 Subspace Disentanglement

Sekh Mainul Islam, Pepa Atanasova, Isabelle Augenstein

University of Copenhagen

{seis, pepa, augenstein}@di.ku.dk

Abstract

Natural Language Explanations (NLEs) describe how Large Language Models (LLMs) make decisions by drawing on external Context Knowledge (CK) and Parametric Knowledge (PK). Understanding the interaction between these sources is key to assessing NLE grounding, yet these dynamics remain under-explored. Prior work has largely focused on i) single-step generation and ii) modeled PK–CK interaction as a binary choice within a rank-1 subspace. This approach overlooks richer interactions and how they unfold over longer generations, such as complementary or supportive knowledge. We propose a novel rank-2 projection subspace that disentangles PK and CK contributions more accurately and use it for the first multi-step analysis of knowledge interactions across longer NLE sequences. Experiments across four QA datasets and three open-weight LLMs demonstrate that rank-1 subspaces struggle to represent diverse interactions, whereas our rank-2 formulation captures them effectively, highlighting PK alignment for supportive interactions and CK alignment for conflicting ones. Our multi-step analysis reveals that hallucinated generations exhibit strong alignment with the PK direction, though context-faithful generations maintain a more balanced alignment between PK and CK¹.

1 Introduction

Large Language Models (LLMs) are employed to generate Natural Language Explanations (NLEs) in a human-readable format, illustrating the underlying decision-making process for predictions in complex reasoning tasks such as Claim Verification (CV) and Question Answering (QA). These NLEs are valuable as they can reveal the utilization of external Context Knowledge (CK) and the Parametric Knowledge (PK) of a model. Consider the

Figure 1 example from a QA task (Cheng et al., 2024), where Llama-3.1-8B-Instruct (Grattafiori et al., 2024) generates the NLE by utilizing both CK and PK, to explain the underlying decision-making process for the final answer prediction. For some tasks, such as CV, we assume the decision-making process and, in turn, the NLE, will rely more on evidence (CK) (Wang and Shu, 2023; Tan et al., 2025), and more on PK in QA tasks with misleading external context. On the other hand, Chain-of-Thought (CoT) prompting (Wei et al., 2022), a widely adopted reasoning methodology, explicitly elicits NLEs of the intermediate reasoning steps and is hypothesized to influence how LLMs integrate PK and CK (Cheng et al., 2024; Su et al., 2024; Tao et al., 2025), potentially improving contextual grounding. However, it remains unclear what the learned PK–CK interaction dynamics in generating the NLEs are, *necessitating a multi-step analysis of PK–CK interactions across longer NLE sequences*.

Prior work (Longpre et al., 2021; Xu et al., 2024; Minder et al., 2025) has primarily focused on uncovering the single-step generation mechanism – typically the final answer, and modeled only conflicting PK–CK interaction as a binary choice in a rank-1 subspace, thereby overlooking richer forms such as complementary or supportive knowledge (Cheng et al., 2024). We thus hypothesize that this *rank-1 subspace is not sufficient to disentangle the individual contributions of PK and CK in all types of knowledge interaction scenarios*. Moreover, they did not characterize the step-by-step PK–CK interaction dynamics over longer sequences, such as NLE generation for different knowledge interactions. To accurately understand the knowledge interaction dynamics during NLE generation, we investigate the following research questions:

RQ1. *Is a rank-1 projection subspace enough for disentangling PK and CK contributions in all types of knowledge interaction scenarios?*

¹Code and data: <https://github.com/copenlu/pk-ck-knowledge-disentanglement>




Figure 1: (Left) Llama-3.1-8B-Instruct model combines parametric (green) and contextual (orange) knowledge to generate NLEs. (Right) Projection of the answer token to a learned low-rank subspace P disentangles its individual knowledge contributions – (a) rank-1 fails with the correct knowledge captured only in its orthogonal subspace, whereas (b) rank-2 succeeds in encoding the correct knowledge sources along the disentangled axes.

RQ2. How do individual PK and CK contributions change over the NLE generation steps for different knowledge interactions?

RQ3. Can we find reasons for hallucinations based on PK–CK interactions?

RQ4. How is the CoT mechanism aligned with the knowledge interaction subspace?

We experiment with four publicly available QA datasets for three open-weight instruction-tuned LLMs. In RQ1 (§4.2), we find that the rank-1 subspace (Minder et al., 2025) fails to disentangle individual knowledge contributions across knowledge interactions. For RQ2 (§4.3), we propose a more accurate rank-2 subspace disentanglement and find that during NLE generation, the model utilizes both knowledge sources with a slight preference for PK; towards the final answer generation, the model aligns closely with the CK direction for PK vs CK conflicting examples; for supportive examples, it aligns more with the PK direction. The learned rank-2 subspace also illustrates PK dominance for sequences with hallucinations (RQ3, §4.4). For RQ4 (§4.5), we observe that the CoT mechanism balances the PK–CK interaction and often reduces PK alignment. Finally, §5 verifies through a subspace intervention that the rank-2 subspace behaviorally disentangles the PK–CK interaction, with PK-axis intervention reducing hallucinations. Overall, this work introduces the first systematic framework for multi-step analysis of knowledge interactions via rank-2 subspace disentanglement, advancing understanding of how LLMs integrate internal and external knowledge.

2 Related Work

Parametric vs. contextual knowledge in LLMs.

A large body of work studies what language models encode as *parametric knowledge* (PK) and how they integrate *contextual knowledge* (CK) at inference time (Petroni et al., 2019; Jiang et al., 2020; Roberts et al., 2020; Hagström et al., 2025, 2026; Marjanovic et al., 2024; Yu et al., 2024; Wangsajaya et al., 2026; Sun et al., 2026a). Early work showed that factual knowledge can be elicited directly from model parameters (Petroni et al., 2019), while probing and representation analyses examined how context and memory are encoded and combined across layers (Brown et al., 2020; Tenney et al., 2019; Bi et al., 2026).

More recent studies focus on PK–CK interaction and conflict. Cheng et al. (2024) analyzed suppression effects between PK and CK under different interaction regimes, while Xu et al. (2024) categorized conflict types (context–memory, inter-context, intra-memory) and their behavioral consequences. Several approaches aim to improve PK–CK balance through training or inference-time interventions, including enhanced pretraining or fine-tuning (Zhang et al., 2024), context-aware representation manipulation (Yuan et al., 2025b), contrastive decoding (Zhao et al., 2024), and lightweight steering that increases context sensitivity without weight updates (Wang et al., 2025). Hagström et al. (2026) empirically compares such context usage manipulation approaches in a joint framework, highlighting the trade-offs between them. Our work complements this line of work by introducing a geometric

framework that disentangles PK and CK in a rank-2 subspace and enables multi-step analysis during generation.

Natural Language Explanations (NLEs) and reasoning. NLEs are widely used to expose or guide model reasoning in complex tasks (Camburu et al., 2018; Atanasova et al., 2020). Explanation-based supervision improves QA performance (Rajani et al., 2019), and Chain-of-Thought (CoT) prompting elicits stepwise reasoning that enhances multi-hop and arithmetic reasoning (Wei et al., 2022; Lampinen et al., 2022). However, multiple studies show that NLEs and CoTs are often unfaithful to the model’s true decision process, functioning as post-hoc rationalizations rather than causal explanations (Atanasova et al., 2023; Turpin et al., 2023; Siegel et al., 2024; Yuan et al., 2025a; Wang and Atanasova, 2025; Sun et al., 2026b). Our work addresses this gap by providing a token-level, representation-based analysis of how PK and CK jointly shape NLE generation, linking explanation faithfulness and hallucination to internal knowledge alignment.

Probing, subspaces, and identifiability. Probing methods reveal that linguistic and factual properties are encoded in low-dimensional subspaces of LM representations (Hewitt and Manning, 2019; Clark et al., 2019). The superposition hypothesis further suggests that multiple features are compressed into overlapping linear subspaces (Elhage et al., 2022). Building on this view, Minder et al. (2025) modeled PK–CK conflict using a single controllable direction (rank-1 subspace). However, when PK and CK are jointly encoded, such a representation is non-identifiable and cannot disentangle individual contributions. We generalize this approach to a rank-2 projection subspace, establishing identifiability and enabling fine-grained, multi-step analysis of PK–CK interaction dynamics during generation.

3 Method

3.1 Task Formulation

Let $\mathcal{Q}$ and $\mathcal{C}$ denote two disjoint sets of queries and contexts, respectively, and consider a QA dataset $\mathcal{D} \subseteq \mathcal{Q} \times \mathcal{C}$. For a question–context pair $(q, c) \in \mathcal{D}$, with $q \in \mathcal{Q}$ and $c \in \mathcal{C}$, a language model p generates both an answer a and an NLE with n tokens as $\mathcal{E} = \{e_i\}_{i=1}^n$. We distinguish between three forms of answer generation: (i) the parametric answer $a(q, \phi)$, produced by recalling PK (independent of c , ϕ denotes no context), (ii) the contextual answer

$a(q, c)$, obtained by leveraging the provided context while disregarding parametric recall, and (iii) the final predicted answer a combining both the provided context and the PK recall guided by the PK–CK knowledge interaction.

During the sequential generation of $\mathcal{E}$, each token e_i is influenced by the interaction between $a(q, \phi)$ and $a(q, c)$. This evolving interaction guides the generation of NLE $\mathcal{E}$, illustrating the decision-making process of the final predicted answer a . To analyze this process, we quantify the contribution of parametric knowledge (α_i^p) and contextual knowledge (α_i^c) at each NLE generation step i , and we define their difference as

$$\Delta_i = \alpha_i^p - \alpha_i^c, i \in [1, n] \quad (1)$$

We normalize individual knowledge contribution $\alpha_i \in [0, 1]$ at each NLE sequence step to compare the influence of PK and CK at each sequence step ($\Delta_i = 0$ indicates balance between PK and CK, $\Delta_i = 1$ indicates PK dominance, and $\Delta_i = -1$ indicates CK dominance) and how this interaction shifts throughout the NLE generation. By tracking Δ_i across all NLE generation steps, we aim to characterize the interaction and the shifting balance between parametric and contextual sources throughout NLE generation.

3.2 Identify Different PK–CK Interactions

To characterize different types of PK–CK interactions guided by individual knowledge contribution, we draw from Minder et al. (2025), who analyze intent-driven answer control, i.e., controlling the model towards a specific answer generation aligned with the intent of following either PK or CK. Let w_c (instruction to follow only CK) and w_p (instruction to follow only PK) denote intents toward predicting $a(q, c)$ and $a(q, \phi)$, respectively, for a given (q, c) .

Minder et al. (2025) restrict intent to conflicting cases, where $a(q, c) \neq a(q, \phi)$ and the final answer a is determined by either w_c or w_p . We extend their conflict-only framing, following Cheng et al. (2024), to encompass a broader set of interactions:

- Supportive: $a(q, c) = a(q, \phi) = a$; PK and CK reinforce the same outcome.
- Complementary: $a(q, c) \neq a(q, \phi)$ but $a = a(q, c) \cup a(q, \phi)$; neither source is sufficient alone, but their combined information yields the final answer.
- Conflicting: $a(q, c) \neq a(q, \phi)$ and $(a = a(q, c)) \oplus (a = a(q, \phi))$; the final answer

aligns exclusively with one knowledge source while rejecting the other.

- Irrelevant: $c \perp q$ such that $a(q, c)$ is null or nonsensical; the context c contains no semantic information relevant to the query q . In this case, the model ignores the context entirely and relies solely on PK ($a = a(q, \phi)$) to formulate a response.
- Knowledge Suppression (noise): $a(q, c) = a(q, \phi) \neq a$; both sources agree on an answer, yet the model fails to utilize either, producing an erroneous or unrelated prediction.

To represent how the model balances contextual (w_c) and parametric (w_p) influences, we extend this intent-driven answer control, and w_b denotes the joint intent of w_c and w_p that illustrates the intrinsic model behavior in generating the final answer a , considering both the PK and CK governed by the underlying interaction scenario. Unlike in Minder et al. (2025), where w_c and w_p are represented using one orthonormal direction in a learned rank-1 projection subspace, we represent w_c and w_p using two orthonormal directions, and $w_b \rightarrow w_c$ and $w_b \rightarrow w_p$ indicate the individual contributions of CK and PK, respectively, in generating a . Then the joint intent $w_b = f(w_c, w_p)$ indicates different knowledge interaction scenarios decided by individual contributions from w_c and w_p . We learn this knowledge interaction function using a rank-2 projection subspace.

3.3 Localize Intent-Guided PK–CK Interactions

In this section, we describe where the intent-guided knowledge interaction emerges in the model space, followed by how the LLM encodes it.

Identifying important layers for rank-1 projection subspace. (showing the individual knowledge contribution towards generating either $a(q, c)$ or $a(q, \phi)$, i.e., w_c or w_p respectively). We follow the activation patching-based mechanistic interpretability approach, Patchscope (Ghandeharioun et al., 2024). Minder et al. (2025) utilizes a Patchscope-based iterative search over all layers to find the best layer corresponding to the maximum likelihood of the PK and CK answers. They construct two minimally different prompts s and t as ‘source’ and ‘target’, with the same q and c , only differing by the intent w_c or w_p , resulting in two different intended answers. To identify the layers capturing the intent present in s , during the forward pass of the target $p(\cdot|t)$, the activation from the hidden state of

a particular layer is replaced by the activations from the same layer with the source $p(\cdot|s)$, resulting in $\tilde{p}(\cdot|t) \approx p(\cdot|s)$. Since they only consider the ‘conflicting’ behaviors in understanding the knowledge interaction in LLMs, they construct two patching datasets capturing opposite directions: $\mathcal{D}_w^{(c \rightarrow p)} = \{s : ((q, c, w_c), a(q, c)), t : ((q, c, w_p), a(q, \phi))\}$ for encoding the CK direction, and $\mathcal{D}_w^{(p \rightarrow c)} = \{s : ((q, c, w_p), a(q, \phi)), t : ((q, c, w_c), a(q, c))\}$ for encoding the PK direction. Layers capturing CK and PK directions, respectively, are selected as:

$$\begin{aligned} \mathbb{L}_{c \rightarrow p} &= \{l \in [1, L] \mid \tilde{p}(a(q, c) \mid q, c, w_p) \geq \tau_c\}, \\ \mathbb{L}_{p \rightarrow c} &= \{l \in [1, L] \mid \tilde{p}(a(q, \phi) \mid q, c, w_c) \geq \tau_p\}, \end{aligned} \quad (2)$$

where τ_c, τ_p are hyperparameters. For more details, please refer to Minder et al. (2025).

Encoding the intent via rank-1 subspace projection. Once we identify the important layers capturing the binary interaction between PK and CK for conflicting behavior, we aim to identify a subset of model weights where the LLM encodes this binary interaction. Minder et al. (2025) hypothesize that LLMs encode knowledge interaction within their parameter space using a low-rank projection subspace. They model this projection subspace $\mathbf{P} \in \mathbb{R}^{d \times d}$ (d is the model hidden representation size) by a rank-1 unit norm direction $\vec{u} \in \mathbb{R}^d$ indicating the unidirectional PK–CK conflicting interaction as $\mathbf{P} = \vec{u}\vec{u}^T$. At any sequence step i , the hidden representation $\vec{h}_i \in \mathbb{R}^d$ for the token e_i in the NLE $\mathcal{E}$ can be linearly decomposed as:

$$\vec{h}_i = (\mathbf{I} - \mathbf{P})\vec{h}_i + \vec{u} \langle \vec{u}^T, \vec{h}_i \rangle \quad (3)$$

$(\mathbf{I} - \mathbf{P})$ component of $\vec{h}_i$ captures other properties in the embedding space, and the $\langle \vec{u}^T, \vec{h}_i \rangle$ captures the PK–CK knowledge contribution in the PK–CK conflicting direction $\vec{u}$. We hypothesize that the rank-1 projection subspace $\mathbf{P}$ cannot disentangle the individual knowledge contributions across different knowledge interaction scenarios. Also, we theoretically argue that this $\mathbf{P}$ fails to satisfy ‘bijective’ properties of mapping from knowledge direction to individual knowledge contribution for all types of knowledge interaction in Theorem A.1 of §A.2. Hence, we propose to model the PK–CK interaction for different knowledge interactions using a rank-2 projection subspace.

Identifying important layers for rank-2 projection subspace (showing the individual knowledge

contribution towards generating a , i.e. $w_b \rightarrow w_c$ and $w_b \rightarrow w_p$). We again follow Patchscope and prepare two patching datasets $\mathcal{D}_w^{(b \rightarrow p)} = \{s : ((q, c, w_b), a), t : ((q, c, w_p), a(q, \phi))\}$ for encoding the contribution of PK interacting with CK in generating a , and $\mathcal{D}_w^{(b \rightarrow c)} = \{s : ((q, c, w_b), a), t : ((q, c, w_c), a(q, c))\}$ for encoding the contribution of CK interacting with PK in generating a . Important layers $\mathbb{L}_{b \rightarrow c}$ and $\mathbb{L}_{b \rightarrow p}$ are selected as:

$$\begin{aligned}\mathbb{L}_{b \rightarrow p} &= \{l \in [1, L] \mid \tilde{p}(a \mid q, c, w_p) \geq \tau_p\}, \\ \mathbb{L}_{b \rightarrow c} &= \{l \in [1, L] \mid \tilde{p}(a \mid q, c, w_c) \geq \tau_c\},\end{aligned}\quad (4)$$

where τ_p, τ_c are hyperparameters (details are described in Figure 10 in §A.1.3).

Encoding the intent via rank-2 subspace projection. Once we identify important layers $\mathbb{L}_{b \rightarrow c}$, and $\mathbb{L}_{b \rightarrow p}$ for $w_b \rightarrow w_c$ and $w_b \rightarrow w_p$ respectively, we learn the joint intent function f encoding how individual knowledge contributions $w_b \rightarrow w_c$ and $w_b \rightarrow w_p$ are mixed to generate the final answer a using the common layers from $\mathbb{L}_{b \rightarrow c}$ and $\mathbb{L}_{b \rightarrow p}$ via a rank-2 projection subspace $\mathbf{P}$ spanned by two orthogonal directions $\vec{\mathbf{u}} \in \mathbb{R}^{d \times 2}$ as $\mathbf{P} = \vec{\mathbf{u}}(\vec{\mathbf{u}}^T \vec{\mathbf{u}})^{-1} \vec{\mathbf{u}}^T$. Since we aim to identify individual contributions from the individual directions, following Minder et al. (2025), we consider those two basis vectors as orthonormal, and hence the rank-2 projection subspace $\mathbf{P}$ is reduced to $\mathbf{P} = \vec{\mathbf{u}}\vec{\mathbf{u}}^T$, since $\vec{\mathbf{u}}^T \vec{\mathbf{u}} = \mathbf{I}_2$. Considering $\vec{\mathbf{u}} = [\vec{\mathbf{u}}_c, \vec{\mathbf{u}}_p]$, at any sequence step i , the hidden representation $\vec{\mathbf{h}}_i \in \mathbb{R}^d$ for the token e_i in the NLE $\mathcal{E}$ can be linearly decomposed as:

$$\vec{\mathbf{h}}_i = (\mathbf{I} - \mathbf{P})\vec{\mathbf{h}}_i + \vec{\mathbf{u}}_c \langle \vec{\mathbf{u}}_c^T, \vec{\mathbf{h}}_i \rangle + \vec{\mathbf{u}}_p \langle \vec{\mathbf{u}}_p^T, \vec{\mathbf{h}}_i \rangle \quad (5)$$

Learning the rank-2 projection subspace. Let $\vec{\mathbf{h}}^{w_p} = (\mathbf{I} - \mathbf{P})\vec{\mathbf{h}} + \vec{\mathbf{u}}_c \langle \vec{\mathbf{u}}_c^T, \vec{\mathbf{h}} \rangle + \vec{\mathbf{u}}_p \langle \vec{\mathbf{u}}_p^T, \vec{\mathbf{h}} \rangle$, and $\vec{\mathbf{h}}^{w_c} = (\mathbf{I} - \mathbf{P})\vec{\mathbf{h}} + \vec{\mathbf{u}}_c \langle \vec{\mathbf{u}}_c^T, \vec{\mathbf{h}} \rangle + \vec{\mathbf{u}}_p \langle \vec{\mathbf{u}}_p^T, \vec{\mathbf{h}} \rangle$ be two patched decompositions of the last-token hidden representation of the NLE $\mathcal{E}$ by replacing the activation from a particular layer $l \in \mathbb{L}_{b \rightarrow p} \cup \mathbb{L}_{b \rightarrow c}$ during the forward pass of $p(\cdot \mid q, c, w_b)$ with the activations from the same layer of the forward pass of $p(\cdot \mid q, c, w_p)$ and $p(\cdot \mid q, c, w_c)$, respectively. Then we learn the projection subspace $\mathbf{P}$ spanned by the two orthonormal basis vectors $\vec{\mathbf{u}} = [\vec{\mathbf{u}}_c, \vec{\mathbf{u}}_p]$ by minimizing the two objectives sequentially, maintaining orthonormal directions

between $\vec{\mathbf{u}}_c$ and $\vec{\mathbf{u}}_p$ ²:

$$\mathcal{J}_p = -\frac{1}{N} \sum_{n=1}^N \log(\tilde{p}(a(q_n, \phi) \mid \vec{\mathbf{h}}_{(n)}^{w_p})), \quad (6)$$

$$\mathcal{J}_c = -\frac{1}{N} \sum_{n=1}^N \log(\tilde{p}(a(q_n, c_n) \mid \vec{\mathbf{h}}_{(n)}^{w_c})). \quad (7)$$

The methodology of assigning the PK–CK direction to the basis vector $\vec{\mathbf{u}}$ as a validation step is described in §A.1.1.

Normalized PK–CK score. At any sequence step i , given the hidden representation $\vec{\mathbf{h}}_i \in \mathbb{R}^d$ for the token e_i in the NLE $\mathcal{E}$, let $c_i = |\langle \vec{\mathbf{u}}_c^T, \vec{\mathbf{h}}_i \rangle|$ and $p_i = |\langle \vec{\mathbf{u}}_p^T, \vec{\mathbf{h}}_i \rangle|$, then the normalized contribution from CK and PK can be computed as $\alpha_i^c = c_i/(c_i + p_i)$ and $\alpha_i^p = p_i/(c_i + p_i)$, and $\alpha_i^c, \alpha_i^p \in [0, 1]$ satisfying the identifiability of individual knowledge contribution (Theorem A.1) under rank-2 subspace.

4 Results

4.1 Experimental Setting

Dataset and Model: We conduct experiments on three open-weight decoder-only instruct-based LLMs: Llama-3.1-8B (Grattafiori et al., 2024), Gemma-2 9B (Team et al., 2024) and Mistral-v0.3 7B (Jiang et al., 2023) using four publicly available QA datasets: BaseFakepedia, MultihopFakepedia (Minder et al., 2025), StrategyQA (Geva et al., 2021), and OpenBookQA (Mihaylov et al., 2018; Cheng et al., 2024). BaseFakepedia is a knowledge conflict dataset containing queries with 23 relations from Wikipedia. MultihopFakepedia is an extension of BaseFakepedia that contains queries that require extra-hop reasoning to generate answers. Both BaseFakepedia and MultihopFakepedia datasets are utilized for the rank-1 subspace analysis in Minder et al. (2025). StrategyQA is a multi-hop reasoning QA dataset that includes implicit reasoning steps in the queries. OpenBookQA is a commonsense reasoning-based QA dataset. Both StrategyQA and OpenBookQA datasets include ‘supportive’, ‘complementary’, ‘conflicting’, ‘irrelevant’, and ‘knowledge suppression’ PK–CK interactions, illustrated in Cheng et al. (2024). Prompt templates for intent-driven answer control, NLE and CoT generations are described in §A.1.2.

²We initialize the projection weight from two orthogonal vectors and `nn.utils.parametrizations.orthogonal()` automatically maintains orthogonality after each optimization step.

4.2 RQ1: Is a Rank-1 Projection Subspace Enough for Disentangling PK and CK Contribution Across All Types of Knowledge Interaction Scenarios?

We verify the sufficiency of rank-1 subspace in disentangling PK–CK contribution in different types of knowledge interactions using two metrics: (a) *Subspace component* (Minder et al., 2025), which indicates the strength of the hidden representation of answer tokens along the direction of the basis vector $\vec{u}$ in the rank-1 subspace. If a rank-1 subspace is sufficient in disentangling PK–CK contributions, then the distribution of subspace components of answer tokens for different interactions should not overlap and have a mean far from zero, indicating a small contribution in the orthogonal $(\mathbf{I} - \mathbf{P})$ direction; (b) *Explained Variance (Cumulative)* (Wall et al., 2003; Lazzaretto et al., 2025), which indicates the minimum rank r required to explain the variance in data distributions for different types of knowledge interactions. It measures the strength of top- r singular values among all possible singular values present in the data distribution, and at rank- r , the fraction should be closer to 1.0 if it is sufficient to explain. More details about the metrics are discussed in Appendix A.1.5. Figure 2 shows the *subspace component* distribution of answer tokens for different knowledge interaction types overlap, and their means converge to 0, indicating insufficiency of rank-1 subspace. Figure 3 also verifies rank-2 subspace as the minimum projection subspace to capture data variance and disentangle PK–CK contribution.

4.3 RQ2: How Do Individual PK and CK Contributions Change Over the NLE Generation for Different Knowledge Interactions?

We investigate the dynamics of individual PK and CK contribution α_i^p and α_i^c at every sequence step i over the NLE $\mathcal{E}$ generation. Among the four datasets, OpenBookQA shows the largest jump in cumulative explained variance from rank-1 to rank-2 across all models (Figure 3). We therefore use OpenBookQA as our development set to (i) locate intent-sensitive layers via activation patching and (ii) learn the rank-2 projection subspace. We then freeze these choices and evaluate the resulting rank-2 subspace on all datasets. StrategyQA yields nearly identical subspaces (directional similarity > 0.9 on both axes; see §A.3). We first




Figure 2: Kernel Density Estimate (KDE) of the PK–CK subspace component $\langle \vec{u}^T, \vec{h}_i \rangle$ across different knowledge interaction types for four QA datasets using Llama-3.1-8B-Instruct model. See similar results for the other two models in Appendix A.2.




Figure 3: Cumulative explained variance (EV_r) at rank(r) from the three models using the four QA datasets. At rank-2, it reaches 1.0 value, indicating sufficiency in capturing different knowledge interaction variants.

localize the layers encoding individual PK and CK contributions. Figure 4 shows that patching from $w_b \rightarrow w_p$ (BOTH→PRI) (L13–18) yields a larger probability gain than $w_b \rightarrow w_c$ (BOTH→CTX) (L15–17), indicating that Llama-3.1-8B-Instruct relies more on PK than on CK for OpenBookQA. We observe similar results for gemma-2-9b-it and Mistral-7B-Instruct-v0.3 (Figures 13 and 14). This PK dominance in a context-rich task suggests that the model tends to recall commonsense information from PK rather than grounding answers in CK, highlighting potential for future work on context-sensitive knowledge control.

To identify the individual knowledge contribution in generating the final answer token a , we project its hidden representation $\vec{h}_a$ from the last layer of $l \in \mathbb{L}_{b \rightarrow p} \cup \mathbb{L}_{b \rightarrow c}$ to the rank-2 projection subspace to compute the p_i and c_i as the





Figure 4: Patchscope on OpenBookQA from Llama-3.1-8B-Instruct. PRI indicates PK, and CTX indicates CK. (a) Activation patching on $\mathcal{D}_w^{(b \rightarrow p)}$. (b) Activation patching on $\mathcal{D}_w^{(b \rightarrow c)}$. Computation details of patchscope are described in Appendix A.1.4.


PK–CK contributions³. Figure 5a indicates an overall higher CK contribution for BaseFakepedia and MultiHopFakepedia, and a higher PK contribution for StrategyQA and OpenBookQA for the Llama-3.1-8B-Instruct model. We further investigate the distribution of different knowledge interactions (following §3.2) in Fig. 5b and find that both BaseFakepedia and MultiHopFakepedia contain more conflicting examples than other knowledge interaction types (§3.2). Our findings match both dataset composition and prior findings: Fakepedia variants often contain *examples with conflicting CK when models are found to suppress PK*; StrategyQA/OpenBookQA have examples that rely on common sense priors and have supportive (i.e., non-conflicting) contexts when models default to parametric recall (Cheng et al., 2024; Tao et al., 2024). Overall, we conclude that *for conflicting examples, the model aligns more with the CK direction and for supportive examples, the model aligns more with the PK direction in the rank-2 projection subspace*.

To understand the knowledge interaction dynamics during NLE generation, We analyze variation in α_i^p and α_i^c over all sequence steps of NLE generation for different knowledge interaction scenarios. Figure 6 shows that for all datasets, *during most of the NLE generations, the model starts with a higher CK, then considers both PK and CK with slight prioritization of PK*. However, for longer NLEs, CK and PK compete with each other with higher fluctuation. Longer NLEs indicate difficult examples with higher depth in multi-hop reason-

³We do not normalize for single token answer.



(a) Individual PK–CK contribution p_i and c_i (§3.3) in generating the answer token for Llama-3.1-8B-Instruct.



(b) Distribution of knowledge interaction types across datasets.

Figure 5: Answer-level PK–CK alignment and distribution of knowledge interaction types across datasets.




Figure 6: PK–CK interaction dynamics over the sequence steps for Llama-3.1-8B-Instruct. The dotted red and blue lines indicate the mean and mode of NLE lengths. Figure 15 contains PK–CK dynamics for different interaction scenarios.

ing and higher token uncertainty (from Figure 16), which force the model to iteratively reconcile PK with CK, resulting in this fluctuating behavior.




Figure 7: PK-CK interaction dynamics over the sequence step from Llama-3.1-8B-Instruct for the two RAG hallucination datasets.




Figure 8: Comparison in individual PK-CK contribution p_i and c_i (§3.3) in generating the answer token a for all the datasets from Llama-3.1-8B-Instruct model between CoT and standard prompting.

4.4 RQ3: Can We Find Reasons for Hallucinations Based on PK-CK Interactions?

To characterize the knowledge interaction dynamics during the NLE generation in terms of context faithfulness, we compare hallucinated vs non-hallucinated responses in the rank-2 projection subspace. Following (Sun et al., 2025), we utilize two commonly used RAG hallucination datasets: RAGTruth (Niu et al., 2024) and Dolly (AC) (Hu et al., 2024), of sizes 18240 and 297 examples, respectively, containing examples from QA, Summarisation and Information Extraction. Each dataset provides high-quality human-annotated spans of hallucinated content across responses from multiple models. Importantly, whether a span is labelled as hallucinated is *model-dependent*: the same RAG input may yield hallucinated text for one model but not for another due to differences in the model’s PK. Due to the limited number of models covered in the two datasets, we find hallucination labels only for the Llama-3.1-8B-Instruct

model in these two datasets. Figure 7 illustrates the PK-CK knowledge interaction dynamics. The gap between PK and CK is much higher for the examples with hallucinated spans than for the other examples across the sequence steps. This result also aligns with similar observations of *positive correlation of PK and hallucination* in Sun et al. (2025).

4.5 RQ4: How is the CoT mechanism aligned with the knowledge interaction subspace?

To verify whether reasoning-based prompting CoT helps the model to stay aligned with the CK and reduces reliance on PK, we compare the PK-CK contribution in generating the final answer between standard prompting and CoT prompting. Figure 8 indicates that *CoT maintains similar CK alignment compared to standard prompting for all the datasets, and also reduces PK alignment except for the OpenBookQA dataset*.

5 Steering Model Behavior via Projection Subspace Intervention

Instead of using the explicit intent instruction w_p (follow prior knowledge only), w_c (follow context knowledge only) and w_b (consider both knowledge sources), we steer the hidden representation of the answer token using the equation:

$$\tilde{\mathbf{h}}_i = (\mathbf{I} - \mathbf{P})\vec{\mathbf{h}}_i + \vec{\mathbf{u}}_c c_i + \vec{\mathbf{u}}_p p_i \quad (8)$$

where $c_i, p_i \in \mathbb{R}^+$ are weights along the CK and PK direction, respectively. We perform this steering in three settings: $p_i > c_i$ (similar to w_p), $p_i < c_i$ (similar to w_c), and $p_i = c_i$ (similar to w_b) on all four datasets and three models using the learned rank-2 projection subspace and measure accuracy on how well the model predicts the same answer without the explicit intent.

From Table 1, the high alignment between steering in the projection subspace and the intent variable that guides the knowledge interaction indeed verifies that the learned projection subspace behaviorally disentangles PK-CK interaction.

For the RAGTruth dataset from Llama-3.1-8B-Instruct, we observe average CK and PK scores of 0.46 ($c_i = 1.8559$) and 0.54 ($p_i = 2.2277$) for hallucinated examples and 0.49 ($c_i = 1.9399$) and 0.51 ($p_i = 2.0927$) for non-hallucinated examples. Using Equation (8) to obtain the hidden representation at each sequence step, following (Sun et al., 2025), we steer 450 examples and regenerate

Dataset	Model	$p_i > c_i$	$p_i < c_i$	$p_i = c_i$
Strategy QA	Llama	0.910	0.940	0.935
	Gemma	0.890	0.910	0.925
	Mistral	0.840	0.875	0.910
Base Fakepedia	Llama	0.865	0.930	0.915
	Gemma	0.830	0.905	0.900
	Mistral	0.835	0.870	0.860
Multihop Fakepedia	Llama	0.870	0.900	0.915
	Gemma	0.845	0.870	0.870
	Mistral	0.827	0.860	0.855
OpenBook QA	Llama	0.940	0.960	0.960
	Gemma	0.917	0.925	0.920
	Mistral	0.890	0.903	0.925

Table 1: Model behavior performance between the answer prediction by intervening in the learned rank-2 subspace and the answer predicted with explicit intent. $p_i > c_i$: $p_i=2.50, c_i=1.95$; $p_i < c_i$: $p_i=2.0, c_i=2.50$; $p_i = c_i$: $p_i=c_i=2.50$.

Steering Condition	Hallucinated / Total
No steering	243 / 450
$p_i = 2.1$ (PK-downweight)	202 / 450 (<0.0001)
$c_i = 1.94$ (CK-upweight)	221 / 450 (0.048)
$p_i = 2.1, c_i = 1.94$ (Both)	193 / 450 (<0.0001)

Table 2: Reduction in hallucinated generation by intervening in the learned rank-2 subspace (p-values from Fisher’s exact test in the bracket).

their responses, keeping the generation parameters unchanged. We then compute the truthfulness of the generated response using Llama-3.1-70B-Instruct with the same prompt template as in (Sun et al., 2025), and report the number of examples identified as hallucinated in Table 2. We compare our hallucination mitigation performance with Sun et al. (2025), which achieves 172 wins, 179 ties, and 99 losses over the non-steered model, resulting in a reduction of (73/450) compared to our reported (50/450) reduction in hallucination generation. However, our result remains statistically significant, indicating a stable reduction in hallucinations compared to the non-steered method. We also note that Sun et al. (2025) computes the per-token residual contribution of the attention and the MLP block and hallucination likelihood based on the positive correlation of the MLP contribution and the negative correlation of the attention contribution with the per-token hallucination gold label. *Overall, our work establishes a causal study of hallucination likelihood with PK–CK interaction via PK directional steering-based truthful generation, whereas Sun et al. (2025) utilizes all possible con-*

cepts relevant for hallucination via residual contribution to detect and mitigate hallucination, resulting in higher but unstable mitigation performance. However, hallucination mitigation performance in (Sun et al., 2025) suggests that the PK axis in the learned rank-2 subspace is neither a unique nor a complete causal mechanism for hallucination. This hallucination mechanism can be explained by another projection subspace apart from the PK–CK subspace. But the high alignment of the PK axis indicates that another subspace should not be orthogonal to the PK direction of the learned rank-2 subspace.

6 Conclusion

This work shows that the interaction between Parametric Knowledge (PK) and Context Knowledge (CK) in LLMs is inherently multidimensional rather than a binary choice. We introduce a rank-2 projection framework that (i) resolves the non-identifiability of rank-1 probes in disentangling individual knowledge contributions, (ii) disentangles token-level PK and CK components, and (iii) enables the first systematic multi-step analysis of PK–CK dynamics during Natural Language Explanation (NLE) generation. Across four QA datasets and three open-weight instruction-tuned LMs, rank-1 subspaces fail to capture diverse knowledge interaction scenarios, whereas rank-2 suffices to explain the observed variance. Our step-wise analysis further reveals consistent patterns: hallucinated generations align strongly with the PK axis, while context-faithful generations maintain a more balanced PK–CK alignment, and Chain-of-Thought prompting shifts generations toward CK by reducing PK reliance. Intervention in the learned rank-2 subspace enables controlling model behavior towards the desired intent and reduces hallucination generation, thereby verifying the behavioral knowledge disentanglement in that subspace. Beyond NLE generation, this subspace-based analysis provides a general, model-internal signal for studying how LLMs balance internal recall and external grounding.

Limitations

Orthonormal assumption of PK–CK directions:

We restrict the subspace to span two orthonormal directions to obtain a pure directional contribution in the LLM knowledge interaction. In principle, the subspace can be spanned by two non-orthogonal

directions, but mapping the PK–CK contribution to those directions would be complicated. We leave it for future work. We verify that these two orthonormal directions correspond to the PK and CK knowledge from the steering-based intervention experiments in Section 5.

Rank-3 or higher rank subspace projection in disentangling PK–CK contribution: We limit our methodology to rank-2 based on the observed cumulative explained variance across all datasets and models, averaged over subspace rank. We also empirically verify the sufficiency of rank-2 in encoding the knowledge interaction concept. As future work, we can examine the orthogonal component of the subspace to understand the model’s sensitivity to it, thereby establishing the need for rank-3.

Comparing rank-2 subspace with PCA, ICA, or alternative representation steering methods:

We find standard PCA and ICA-based methods are invalid baselines for disentangling knowledge interactions within the model. PCA finds directions along the maximum variance in the hidden representation matrix of examples and may not correspond to the actual knowledge directions. ICA finds two statistically independent directions in the representation, which are again causally connected to the knowledge directions. PatchScope-based rank-1 or rank-2 subspace learning enforces the learned subspace direction to the knowledge direction. We theoretically show that rank-1 is non-identifiable in disentangling individual PK and CK contributions, and empirically verify the sufficiency of the proposed rank-2 subspace.

Inspecting triggering effects causing the model to shift from CK to PK: Inspecting triggering effects that cause the model to shift from CK to PK warrants a different set of experiments, which is not the focus of our work, as we aim to disentangle and model the PK–CK contributions. Sun et al. (2025) introduced the external context score (ECS) by calculating the residual contribution from multi-head attention and the parametric knowledge score (PKS) from the FFN blocks’ residual contributions for a particular token, and found a positive correlation between hallucination and PKS. They observe that the model successfully attends to correct context around 65% of the time, but it loses this information over longer sequence generation, leading to lower ECS or higher PKS resulting in

hallucination.

Nonlinearity of LLM representations: We follow the previous findings on the linear decomposition of several concepts encoded in the token’s hidden representation and nonlinearity in the interaction is out of scope for our study.

Evaluation beyond QA: We investigate the effectiveness of the rank-2 projection subspace on three open-weight decoder-only instruction-tuned LLMs across four QA datasets due to the following reasons: 1) QA datasets are the only available annotated datasets with different types of knowledge interactions, 2) QA tasks are traditionally used to explore interaction between the model’s own parametric knowledge (PK) and the external context knowledge (CK) as they naturally rely on both types of knowledge. Investigating diverse sequence generation datasets, such as summarization, translation, or dialogue, would broaden the generalization of our findings. However, translation requires interaction between context and multi-lingual parametric knowledge, which would require a different type of study of how context in one language is linked with PK in another language in the model, and where we do not expect different types of knowledge interactions. Same with both summarization and dialogue generation, where they heavily rely on external context (external document, conversation history). Also, any information-seeking queries in a dialogue are equivalent to a question-answering task.

Ethical Considerations

Use of AI Assistants. The authors used Claude API to assist with language polishing and grammatical corrections, clarity and coherence. All technical contributions, including experimental design choices and final decisions, were made by the authors.

Acknowledgements

  This research was co-funded by the European Union (ERC, ExplainYourself, 101077481) and by VILLUM FONDEN (grant number 40543). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- Pepa Atanasova, Oana-Maria Camburu, Christina Lioma, Thomas Lukasiewicz, Jakob Grue Simonsen, and Isabelle Augenstein. 2023. [Faithfulness Tests for Natural Language Explanations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 283–294, Toronto, Canada. Association for Computational Linguistics.
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. [Generating Fact Checking Explanations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7352–7364, Online. Association for Computational Linguistics.
- Baolong Bi, Shenghua Liu, Yiwei Wang, Yilong Xu, Junfeng Fang, Lingrui Mei, and Xueqi Cheng. 2026. [Parameters vs. Context: Fine-Grained Control of Knowledge Reliance in Language Models](#). In *The Fourteenth International Conference on Learning Representations*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. [e-SNLI: Natural Language Inference with Natural Language Explanations](#). In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Sitao Cheng, Liangming Pan, Xunjian Yin, Xinyi Wang, and William Yang Wang. 2024. [Understanding the Interplay between Parametric and Contextual Knowledge for Large Language Models](#). *Preprint*, arXiv:2410.08414.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What Does BERT Look at? An Analysis of BERT’s Attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. 2022. [Toy models of superposition](#). *Transformer Circuits Thread*.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. [Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies](#). *Transactions of the Association for Computational Linguistics*, 9:346–361.
- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. 2024. [Patchscopes: A Unifying Framework for Inspecting Hidden Representations of Language Models](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 15466–15490. PMLR.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The Llama 3 Herd of Models](#). *Preprint*, arXiv:2407.21783.
- Lovisa Hagström, Youna Kim, Haeun Yu, Sang-goo Lee, Richard Johansson, Hyunsoo Cho, and Isabelle Augenstein. 2026. [CUB: Benchmarking Context Utilisation Techniques for Language Models](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25101–25133, San Diego, California, United States. Association for Computational Linguistics.
- Lovisa Hagström, Sara Vera Marjanovic, Haeun Yu, Arnav Arora, Christina Lioma, Maria Maistro, Pepa Atanasova, and Isabelle Augenstein. 2025. [A Reality Check on Context Utilisation for Retrieval-Augmented Generation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 19691–19730, Vienna, Austria. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. [A Structural Probe for Finding Syntax in Word Representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xiangkun Hu, Dongyu Ru, Lin Qiu, Qipeng Guo, Tianhang Zhang, Yang Xu, Yun Luo, Pengfei Liu, Yue Zhang, and Zheng Zhang. 2024. [Knowledge-Centric Hallucination Detection](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6953–6975, Miami, Florida, USA. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7B](#). *Preprint*, arXiv:2310.06825.

- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How Can We Know What Language Models Know?](#) *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Andrew Lampinen, Ishita Dasgupta, Stephanie Chan, Kory Mathewson, Mh Tessler, Antonia Creswell, James McClelland, Jane Wang, and Felix Hill. 2022. [Can language models learn from explanations in context?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 537–563, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Margherita Lazzaretto, Jonas Peters, and Niklas Pfister. 2025. [Invariant Subspace Decomposition](#). *Journal of Machine Learning Research*, 26(95):1–56.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. [Entity-Based Knowledge Conflicts in Question Answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Sara Vera Marjanovic, Haeun Yu, Pepa Atanasova, Maria Maistro, Christina Lioma, and Isabelle Augenstein. 2024. [DYNAMICQA: Tracing Internal Knowledge Conflicts in Language Models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14346–14360, Miami, Florida, USA. Association for Computational Linguistics.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. [Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium. Association for Computational Linguistics.
- Julian Minder, Kevin Du, Niklas Stoehr, Giovanni Monea, Chris Wendler, Robert West, and Ryan Cotterell. 2025. [Controllable Context Sensitivity and the Knob Behind It](#). In *The Thirteenth International Conference on Learning Representations*.
- Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, KaShun Shum, Randy Zhong, Juntong Song, and Tong Zhang. 2024. [RAGTruth: A Hallucination Corpus for Developing Trustworthy Retrieval-Augmented Language Models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10862–10878, Bangkok, Thailand. Association for Computational Linguistics.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language Models as Knowledge Bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Explain Yourself! Leveraging Language Models for Commonsense Reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942, Florence, Italy. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How Much Knowledge Can You Pack Into the Parameters of a Language Model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics.
- Noah Siegel, Oana-Maria Camburu, Nicolas Heess, and Maria Perez-Ortiz. 2024. [The Probabilities Also Matter: A More Faithful Metric for Faithfulness of Free-Text Explanations in Large Language Models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 530–546, Bangkok, Thailand. Association for Computational Linguistics.
- Xin Su, Tiep Le, Steven Bethard, and Phillip Howard. 2024. [Semi-Structured Chain-of-Thought: Integrating Multiple Sources of Knowledge for Improved Language Model Reasoning](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8597–8613, Mexico City, Mexico. Association for Computational Linguistics.
- Jingyi Sun, Pepa Atanasova, Sagnik Ray Choudhury, Sekh Mainul Islam, and Isabelle Augenstein. 2026a. [Evaluation Framework for Highlight Explanations of Context Utilization in Language Models](#). *Computational Linguistics*, pages 1–35.
- Jingyi Sun, Qianli Wang, Pepa Atanasova, Nils Feldhus, and Isabelle Augenstein. 2026b. [Investigating the Interplay between Contextual and Parametric Chain-of-Thought Faithfulness under Optimization](#). *Preprint*, arXiv:2605.24960.
- ZhongXiang Sun, Xiaoxue Zang, Kai Zheng, Jun Xu, Xiao Zhang, Weijie Yu, Yang Song, and Han Li. 2025. [ReDeEP: Detecting Hallucination in Retrieval-Augmented Generation via Mechanistic Interpretability](#). In *The Thirteenth International Conference on Learning Representations*.
- Xin Tan, Bowei Zou, and Ai Ti Aw. 2025. [Improving Explainable Fact-Checking with Claim-Evidence Correlations](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 1600–1612, Abu Dhabi, UAE. Association for Computational Linguistics.
- Yufei Tao, Adam Hiatt, Erik Haake, Antonie J. Jetter, and Ameeta Agrawal. 2024. [When Context Leads](#)

- but Parametric Memory Follows in Large Language Models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4034–4058, Miami, Florida, USA. Association for Computational Linguistics.
- Yufei Tao, Adam Hiatt, Rahul Seetharaman, and Ameeta Agrawal. 2025. "Lost-in-the-Later": Framework for Quantifying Contextual Grounding in Large Language Models. *Preprint*, arXiv:2507.05424.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, and 179 others. 2024. *Gemma 2: Improving Open Language Models at a Practical Size*. *Preprint*, arXiv:2408.00118.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. BERT Rediscovered the Classical NLP Pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. 2023. Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Michael E. Wall, Andreas Rechtsteiner, and Luis M. Rocha. 2003. *Singular Value Decomposition and Principal Component Analysis*, pages 91–109. Springer US, Boston, MA.
- Haoran Wang and Kai Shu. 2023. Explainable Claim Verification via Knowledge-Grounded Reasoning with Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6288–6304, Singapore. Association for Computational Linguistics.
- Yilin Wang, Heng Wang, Yuyang Bai, and Minnan Luo. 2025. Continuously Steering LLMs Sensitivity to Contextual Knowledge with Proxy Models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 4682–4698, Suzhou, China. Association for Computational Linguistics.
- Yingming Wang and Pepa Atanasova. 2025. Self-Critique and Refinement for Faithful Natural Language Explanations. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 8481–8507, Suzhou, China. Association for Computational Linguistics.
- Nadya Yuki Wangsajaya, Haeun Yu, and Isabelle Augenstein. 2026. Emergence of Context Characteristics Sensitivity in Large Language Models. *Preprint*, arXiv:2606.09525.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. Chain of Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*.
- Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024. Knowledge Conflicts for LLMs: A Survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8541–8565, Miami, Florida, USA. Association for Computational Linguistics.
- Haeun Yu, Pepa Atanasova, and Isabelle Augenstein. 2024. Revealing the Parametric Knowledge of Language Models: A Unified Framework for Attribution Methods. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8173–8186, Bangkok, Thailand. Association for Computational Linguistics.
- Shuzhou Yuan, Jingyi Sun, Ran Zhang, Michael Färber, Steffen Eger, Pepa Atanasova, and Isabelle Augenstein. 2025a. Graph-Guided Textual Explanation Generation Framework. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 29362–29386, Suzhou, China. Association for Computational Linguistics.
- Xiaowei Yuan, Zhao Yang, Ziyang Huang, Yequan Wang, Siqi Fan, Yiming Ju, Jun Zhao, and Kang Liu. 2025b. Exploiting Contextual Knowledge in LLMs through V-usable Information based Layer Enhancement. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31726–31741, Vienna, Austria. Association for Computational Linguistics.
- Hao Zhang, Yuyang Zhang, Xiaoguang Li, Wenxuan Shi, Haonan Xu, Huanshuo Liu, Yasheng Wang, Lifeng Shang, Qun Liu, Yong Liu, and Ruiming Tang. 2024. Evaluating the External and Parametric Knowledge Fusion of Large Language Models. *Preprint*, arXiv:2405.19010.
- Zheng Zhao, Emilio Monti, Jens Lehmann, and Haytham Assem. 2024. Enhancing Contextual Understanding in Large Language Models through Contrastive Decoding. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4225–4237, Mexico City, Mexico. Association for Computational Linguistics.

A Appendix

A.1 Replication Details

A.1.1 Determining PK and CK Directions in the Rank-2 Projection Subspace.

Once we obtain the rank-2 projection subspace $\mathbf{P}$ spanned by the orthonormal basis vectors $\vec{\mathbf{u}} = [\vec{\mathbf{u}}_c, \vec{\mathbf{u}}_p] \in \mathbb{R}^{d \times 2}$, we identify the PK and CK directions as follows:

$$\vec{\mathbf{u}}_p = \arg \max_{\vec{\mathbf{u}}_p, \vec{\mathbf{u}}_c} \left(\vec{\mathbf{u}}^\top \mathbf{H}_{w_p} \right), \quad (9)$$

$$\vec{\mathbf{u}}_c = \arg \max_{\vec{\mathbf{u}}_p, \vec{\mathbf{u}}_c} \left(\vec{\mathbf{u}}^\top \mathbf{H}_{w_c} \right), \quad (10)$$

where $\mathbf{H}_{w_p} \in \mathbb{R}^{n \times d}$, and $\mathbf{H}_{w_c} \in \mathbb{R}^{n \times d}$ are matrices of hidden representation of answer tokens $a(q, \phi)$ and $a(q, c)$ guided by intents w_p and w_c respectively over n examples. If the optimization for learning $\vec{\mathbf{u}}_p$ and $\vec{\mathbf{u}}_c$ (§3.3) converges well based on two pre-selected directions, then we will always get two different directions for the two knowledge sources.

A.1.2 Prompt Template

Prompt template to generate controlled answers $a(q, \phi)$, $a(q, c)$ and a driven by intents w_p , w_c , and w_b are described in Table 3. The prompt template for generating NLEs and CoT-based prompts across datasets and models is described in Figure 9.

A.1.3 Hyperparameters

We identify important layers $\mathbb{L}_{b \rightarrow c}$ and $\mathbb{L}_{b \rightarrow p}$ to capture individual PK and CK contributions, respectively, from the final answer a via Patchscope using the hyperparameters described in Figure 10.

A.1.4 Computation Overhead of PatchScope

We perform patchscope-based layer selection to learn the rank-2 PK–CK projection subspace on a single NVIDIA B200 GPU (180 GB). For both Llama-3.1-8B and Mistral-v0.3-7B with 32 layers, it takes less than 30 minutes for 500 examples; for Gemma-2-9B with 26 layers, it takes slightly over 30 minutes for 200 examples. This observation verifies negligible computational overhead for patchscope-based layer selection and subsequent rank-2 subspace learning, making it computationally feasible for larger models.

A.1.5 Evaluation Metric

To empirically verify the insufficiency of the rank-1 projection subspace in disentangling individual

knowledge contributions, we utilize two specific metrics:

Subspace component: For each knowledge interaction type, we compute the subspace component $\langle \vec{\mathbf{u}}^T, \vec{\mathbf{h}}_a \rangle$ of the hidden representation of the answer token $\vec{\mathbf{h}}_a$, which is equivalent to the scalar component of the basis vector $\vec{\mathbf{u}}$ for the rank-1 projection subspace $\mathbf{P}$ (Minder et al., 2025). If $\mathbf{P}$ adequately captures PK–CK interactions in the dataset, then the distribution of these subspace components should have a mean significantly different from 0. Such a pattern would indicate minimal contribution in the orthogonal complement $(\mathbf{I} - \mathbf{P})$ and reflect distinct, disentangled interaction behavior across different knowledge types.

Explained Variance (Cumulative): Consider the matrix $H = [\{\vec{\mathbf{h}}_{a_j}\}_{j=1}^N]$, $H \in \mathbb{R}^{N \times d}$ as the concatenation of hidden representation $\vec{\mathbf{h}}_a \in \mathbb{R}^d$ of the answer token a over N examples. The d singular values $\sigma_1 \gg \sigma_2, \dots \gg \sigma_d$ from the diagonal matrix $\Sigma \in \mathbb{R}^{d \times d}$ after the Singular Value Decomposition (SVD) of $HH^T = A\Sigma B$ indicates the strength of HH^T in d orthonormal directions. For top- r singular values $\sigma_1 \gg \sigma_2, \dots \gg \sigma_r$, the cumulative explained variance $EV_r: EV_r = \frac{\sum_{j=1}^r \sigma_j^2}{\sum_{j=1}^d \sigma_j^2}$ indicates the sufficiency of rank- r projection subspace in encoding the variance in knowledge interactions present in the dataset (Wall et al., 2003; Lazzaretto et al., 2025). If the rank-1 projection subspace $\mathbf{P}$ adequately captures PK–CK interactions in the dataset, then the cumulative explained variance at rank-1 EV_1 should approximate 1.0.

A.2 Additional Results

Theorem A.1 (Non-identifiability under rank-1). *Let the hidden representation $\vec{\mathbf{h}}_i$ for the input x_i at the sequence step i be decomposed as*

$$\vec{\mathbf{h}}_i = c_i \vec{\mathbf{u}}_{CK} + p_i \vec{\mathbf{u}}_{PK} + \xi_i,$$

where $\vec{\mathbf{u}}_{CK}, \vec{\mathbf{u}}_{PK}$ are orthonormal directions corresponding to context and parametric knowledge, $c_i, p_i \in \mathbb{R}$ are their contributions, and ξ_i is noise orthogonal to their span. A rank-1 probe with vector $\vec{\mathbf{v}}$ observes

$$\alpha_i = \vec{\mathbf{v}}^\top \vec{\mathbf{h}}_i = c_i \langle \vec{\mathbf{v}}, \vec{\mathbf{u}}_{CK} \rangle + p_i \langle \vec{\mathbf{v}}, \vec{\mathbf{u}}_{PK} \rangle.$$

Then (c_i, p_i) are not uniquely identifiable from α_i whenever both coefficients are nonzero.

Proof. Let $a = \langle \vec{\mathbf{v}}, \vec{\mathbf{u}}_{CK} \rangle$ and $b = \langle \vec{\mathbf{v}}, \vec{\mathbf{u}}_{PK} \rangle$. For any (c, p) , choose $(c', p') = (c + \delta, p - \frac{a}{b}\delta)$ with

Intent	Prompt
w_p	<code>< begin_of_text >< start_header_id >system< end_header_id ></code> Answer the following query considering the provided context. Answer with only one word.< eot_id >< start_header_id >user< end_header_id > Context: Daily Grind Cafe is a popular place to eat lunch out. At a Potluck Party, everyone brings their own lunch to share with friends. Instruction: Ignore the context in answering the query. Query: Which action can help a person save money? A. stop eating at Daily Grind Cafe B. attend a Potluck Party C. Unknown< eot_id >< start_header_id >assistant< end_header_id >
w_c	<code>< begin_of_text >< start_header_id >system< end_header_id ></code> Answer the following query considering the provided context. Answer with only one word.< eot_id >< start_header_id >user< end_header_id > Context: Daily Grind Cafe is a popular place to eat lunch out. At a Potluck Party, everyone brings their own lunch to share with friends. Instruction: Only consider the context in answering the query. Query: Which action can help a person save money? A. stop eating at Daily Grind Cafe B. attend a Potluck Party C. Unknown< eot_id >< start_header_id >assistant< end_header_id >
w_b	<code>< begin_of_text >< start_header_id >system< end_header_id ></code> Answer the following query considering the provided context. Answer with only one word.< eot_id >< start_header_id >user< end_header_id > Context: Daily Grind Cafe is a popular place to eat lunch out. At a Potluck Party, everyone brings their own lunch to share with friends. Instruction: Consider the context in answering the query. Query: Which action can help a person save money? A. stop eating at Daily Grind Cafe B. attend a Potluck Party C. Unknown< eot_id >< start_header_id >assistant< end_header_id >

Table 3: Prompt template for intent-driven answer control.

$\delta \neq 0$. Then $ac + bp = ac' + bp'$, so infinitely many (c', p') yield the same observation. Thus the mapping $(c, p) \mapsto \alpha$ is non-injective, and the individual contributions cannot be disentangled.

A.3 Robustness Analysis

We evaluate the robustness of PatchScope localization and the learned rank-2 PK–CK projection subspace along four axes: layer sensitivity, prompt perturbations, context variation, and geometric stability.

PatchScope layer sensitivity. To ensure that PK–CK disentanglement is not tied to a single layer, we evaluate the rank-2 subspace across six PatchScope-identified layers of Llama-3.1-8B-Instruct (L13–L18). Patching accuracy is stable across 6 layers L13-18, $[0.82, 0.84, 0.83, 0.85, 0.87, 0.86]$, with a mean of 0.84 ± 0.018 and peak performance at L17 (0.87). The resulting stability score of 0.98 indicates that disentanglement consistently emerges across this mid-to-late layer block rather than being a layer-specific artifact.

Intent	Prompt
NLE	< begin_of_text >< start_header_id >system< end_header_id > Answer the following query considering the provided context. Generate your final answer with only one word. If you are unable to answer the query, generate your final answer as "Unknown". Also, generate an explanation to determine your final answer. Return your output in JSON format: {"explanation": "your explanation here", "answer": "your final response here"}. Only include the JSON object in your response. < eot_id >< start_header_id >user< end_header_id > Context: Daily Grind Cafe is a popular place to eat lunch out. At a Potluck Party, everyone brings their own lunch to share with friends. Query: Which action can help a person save money? A. stop eating at Daily Grind Cafe B. attend a Potluck Party C. Unknown < eot_id >< start_header_id >assistant< end_header_id >
CoT	< begin_of_text >< start_header_id >system< end_header_id > Answer the following query considering the provided context. Generate your final answer with only one word. If you are unable to answer the query, generate your final answer as "Unknown". Also, generate an explanation to determine your final answer. Return your output in JSON format: {"explanation": "your explanation here", "answer": "your final response here"}. Only include the JSON object in your response. < eot_id >< start_header_id >user< end_header_id > Context: Daily Grind Cafe is a popular place to eat lunch out. At a Potluck Party, everyone brings their own lunch to share with friends. Query: Which action can help a person save money? A. stop eating at Daily Grind Cafe B. attend a Potluck Party C. Unknown < eot_id >< start_header_id >assistant< end_header_id > Give your answer by analyzing step by step.

Figure 9: Prompt templates for NLE generation and CoT-based prompting.

Hyperparameter	Llama	Gemma	Mistral
number of samples	500	200	500
batch size	24	10	10
τ_p	0.65	0.75	0.75
τ_c	0.60	0.85	0.65
margin	0.3	0.3	0.3
eps	0.05	0.05	0.05

Figure 10: Patching hyperparameters for identifying important layers for rank-2 projection subspaces from Llama-3.1-8B-Instruct, Gemma-2-9B-it, and Mistral-7B-Instruct-v0.3.

Prompt perturbation robustness. We test invariance to surface-form variations via paraphrasing, synonym substitution, and word-order perturba-

tions. End-to-end prediction agreement remains high (model-dependent), at 0.92, 0.88, and 0.85 for paraphrase, synonym, and word-order changes,




Figure 11: Kernel Density Estimate (KDE) of the PK-CK subspace component $\langle \vec{u}^T, \vec{h}_i \rangle$ across different knowledge interaction types for four question-answer datasets using the gemma-2-9b-it model.




Figure 12: Kernel Density Estimate (KDE) of the PK-CK subspace component $\langle \vec{u}^T, \vec{h}_i \rangle$ across different knowledge interaction types for four QA datasets using Mistral-7B-Instruct-v0.3.

respectively. To assess geometric robustness, we independently learn projection subspaces from original and perturbed prompts and measure their cosine similarity, obtaining mean consistency scores of 0.95 ± 0.04 , 0.93 ± 0.05 , and 0.89 ± 0.08 . These results indicate that the learned subspace is largely insensitive to superficial prompt variations and primarily reflects semantic knowledge alignment.

Context variation sanity check. If the projection subspace captures genuine knowledge interaction, PK-CK contributions should shift with context availability. Varying context from minimal (no context) to medium (first 50 CK tokens) to maximal (full context) yields: minimal : ($p_i=0.82, c_i=0.24$), medium : ($p_i=0.65, c_i=0.58$), maximal : ($p_i=0.85, c_i=0.87$). The increase in CK contribution from minimal to medium context reflects




Figure 13: Patchscope on OpenBookQA dataset from gemma-2-9b-it. (a) Activation patching on $\mathcal{D}_w^{(b \rightarrow p)}$ results in a higher contribution of PK in generating the final answer, as the probability gap between the source and target is higher. (b) Activation patching on $\mathcal{D}_w^{(b \rightarrow c)}$ results in a lower contribution of CK in generating the final answer, as the probability gap between the source and target is lower. We consider the common layers from both activation patching to learn the rank-2 projection subspace.




Figure 14: Patchscope on OpenBookQA dataset from Mistral-7B-Instruct-v0.3. (a) Activation patching on $\mathcal{D}_w^{(b \rightarrow p)}$ results in a higher contribution of PK in generating the final answer, as the probability gap between the source and target is higher. (b) Activation patching on $\mathcal{D}_w^{(b \rightarrow c)}$ results in a lower contribution of CK in generating the final answer, as the probability gap between the source and target is lower. We consider the common layers from both activation patching to learn the rank-2 projection subspace.

greater reliance on evidence as it becomes available, while maximal context strongly activates both knowledge sources, consistent with supportive and complementary interactions rather than pure suppression of PK.

Geometric stability and separation. The learned parametric and contextual directions are highly stable across resampling, as shown with examples




Figure 15: PK-CK interaction dynamics over the sequence steps for different interaction scenarios for Llama-3.1-8B-Instruct. The dotted red and blue lines indicate the mean and mode of NLE lengths.




Figure 16: Entropy of NLE generation across all datasets from Llama-3.1-8B-Instruct for different NLE lengths grouped in four quartiles.

from the StrategyQA dataset, with average directional similarity exceeding 0.9 for both axes. Moreover, the two directions remain well separated: the mean cosine overlap between PK and CK axes is 0.00 across bootstraps. This consistent orthogo-

nality supports the identifiability motivation for a rank-2 formulation. Overall, these results demonstrate that PK-CK disentanglement is stable across layers, robust to prompt surface perturbations, responsive to context evidence, and geometrically well-conditioned.